

NISS

Bayesian Inference on Network Traffic Using Link Count Data

Claudia Tebaldi and Mike West

Technical Report Number 66
September, 1997

National Institute of Statistical Sciences
19 T. W. Alexander Drive
PO Box 14006
Research Triangle Park, NC 27709-4006
www.niss.org

Bayesian Inference on Network Traffic Using Link Count Data

Claudia Tebaldi & Mike West
ISDS, Duke University

Claudia Tebaldi, PhD, is postdoctoral fellow in the Climate Analysis Section/Geophysical Statistics Project, National Center for Atmospheric Research, Boulder CO 80307. Mike West is Professor and Director, Institute of Statistics and Decision Sciences, Duke University, Durham NC 27708-0251. The research reported here was partially supported by NSF grant DMS-9313013 to the National Institute of Statistical Sciences. This paper is to appear (with invited discussion) in the June 1998 of the Journal of the American Statistical Association.

Abstract

We study Bayesian models and methods for analysing network traffic counts in problems of inference about the traffic intensity between directed pairs of origins and destinations in networks. This is a class of problems very recently discussed by Vardi in a 1996 JASA article, and of interest in both communication and transportation network studies. The current paper develops the theoretical framework of variants of the origin-destination flow problem, and introduces Bayesian approaches to analysis and inference. In the first, the so-called fixed routing problem, traffic or messages pass between nodes in a network, with each message originating at a specific source node, and ultimately moving through the network to a predetermined destination node. All nodes are candidate origin and destination points. The framework assumes no travel time complications, considering only the number of messages passing between pairs of nodes in a specified time interval. The route count, or route flow, problem is to infer the set of actual number of messages passed between each directed origin-destination pair in the time interval, based on the observed counts flowing between all directed pairs of adjacent nodes. Based on some development of the theoretical structure of the problem and assumptions about prior distributional forms, we develop posterior distributions for inference on actual origin-destination counts and associated flow rates. This involves iterative simulation methods, or Markov chain Monte Carlo (MCMC), that combine Metropolis-Hastings steps within an overall Gibbs sampling framework. We discuss issues of convergence and related practical matters, and illustrate the approach in a network previously studied in Vardi's 1996 article. We explore both methodological and applied aspects much further in a concrete problem of a road network in North Carolina, studied in transportation flow assessment contexts by civil engineers. This investigation generates critical insight into limitations of statistical analysis, and particularly of non-Bayesian approaches, due to inherent structural features of the problem. A truly Bayesian approach, imposing partial stochastic constraints through informed prior distributions, offers a way of resolving these problems, and is consistent with prevailing trends in updating traffic flow intensities in this field. Following this, we explore a second version of the problem that introduces elements of uncertainty about routes taken by individual messages in terms of Markov selection of outgoing links for messages at any given node. For specified route choice probabilities, we introduce the concept of a super-network, namely a fixed routing problem in which the stochastic problem may be embedded. This leads to solution of the stochastic version of the problem using the methods developed for the original formulation of the fixed routing problem. This is also illustrated. Finally, we discuss various related issues and model extensions, including inference on stochastic route choice selection probabilities, questions of missing data and partially observed link counts, and relationships with current research on road traffic network problems in which travel times within links are non-negligible and may be estimated from additional data.

Keywords: Bayesian Inference; Network Flow Data; Origin-Destination Flows; Markov Chain Monte Carlo

1. INTRODUCTION

In a recent article, Vardi (1996) addressed problems of inference in traffic intensity counts in networks based on observed traffic counts on directed network links between pairs of nodes. The archetype application context is a communication network. Here traffic counts on an individual directed link represent numbers of messages transmitted across that link. Interest lies, at least in part, in estimating the source, or origin, to destination counts; that is, for each pair of communicating nodes i and j in the network, estimate the number of packets transmitted from i to j based on the observed link counts. This is one of a collection of problems of interest in various areas, including, with extensions to incorporate non-negligible and stochastic link transit times, studies of transportation networks. In this area, there is long-standing interest in estimation of so-called “OD matrices” that represent the actual or expected flows or counts of traffic between all possible pairs of origin (“O”) and destination (“D”) nodes or zones in urban road networks during a specific time interval. These matrices are used as inputs to network equilibrium models that are central tools in transportation modelling and policy studies (Berka and Boyce 1994, Sheffi 1995, West 1994). Our work directly connects with specific road transportation network problems, and we develop an applied study in the area in Section 4 below.

The current article addresses the network count inference problem from a Bayesian perspective. We elaborate on the theoretical structure of the basic network count problem, and develop analysis for the so-called fixed routing problem; that is, inferring flows of traffic between all directed origin-destination pairs based on observed counts on all links in the network. This underlies the development of posterior analysis based on iterative simulation methods, and is presented in Section 2 of the paper following introduction of the network specification, notation, terminology and basic structure. Various difficulties of analysis encountered in non-Bayesian approaches are alleviated by approaching the problem this way. Throughout Section 2 and in Section 3, we illustrate the ideas and present some computational results in an example network studied by Vardi, as cited. Section 2 focusses on the algebraic and statistical structure of the basic network flow problem, and discusses prior modelling and aspects of resulting posterior distributions. Here we discuss and exemplify a general structural feature of the network problem that raises issues and challenges for any statistical analysis. This relates to unavoidable ambiguities in inferences based solely on likelihood or other non-Bayesian methods, and whose solution requires imposition of additional constraints. We discuss and demonstrate how this is naturally achieved by imposition of partial stochastic constraints through informed prior distributions. Section 3 is concerned with issues of implementation and computation. Section 4 discusses an informative real application. This is followed by discussion of extensions to incorporate stochastic route selection, also discussed by Vardi (1996), using Markov models for sequential selection of links to be traversed in moving between a specific origin-destination pair. We develop a simple and natural embedding of such a problem in a larger, artificial network introduced to map the stochastic route choice problem onto the fixed route framework. This implies that a solution to the stochastic version of the problem lies in the fixed route problem already developed. This neat solution is attractive in its conceptual simplicity and rather interestingly reverses the roles of fixed and stochastic networks; whereas the fixed network problem is a special case of stochastic routing, the solution to the latter is achieved by reversing the roles and defining it as a special fixed routing problem. We conclude with discussions about extensions to include inference on stochastic route choice selection probabilities, questions of missing data and partially observed link counts, and relationships with current research on road traffic network problems.

2. THE BASIC NETWORK INFERENCE PROBLEM

2.1 Notation and Structure

We adopt the basic structure and notation of Vardi (1996). Consider a fixed network of n nodes, arbitrarily labelled A, B, \dots . The basic problem considers a fixed period of communication or passage of traffic during which a collection of messages moves through the network; each message originates at one node in the network, its origin or source node, and travels to another, its destination node, along a unique path as long as we deal with a “fixed routing problem”. We are concerned with actual numbers, or counts, of messages travelling between pairs of nodes in the network. For any ordered pair of nodes $a = (i, j)$, write X_a for the number of messages originating at node i and terminating at its destination node j . X_a is the origin-destination count between these two nodes. Given n nodes in the network, there are $c = n(n - 1)$ origin-destination pairs. Interest lies in estimating the full set of such counts based on observations of traffic counts on all individual directed links in the network, i.e. of traffic counts between pairs of nodes that communicate directly, without intervening nodes. Let r be the total number of directed links in the network. Let $s = (i, j)$ represent the directed link from node i to node j , and write Y_s for the traffic count on this link. Then, based on observed link counts $\mathbf{Y} \stackrel{\text{def}}{=} (Y_1, \dots, Y_r)'$, we are interested in inferring origin-destination counts $\mathbf{X} \stackrel{\text{def}}{=} (X_1, \dots, X_c)'$. Note that the number of observed link counts r is typically smaller than the number of origin-destination pairs c , especially in very large networks. Throughout the paper, we implicitly assume $r < c$.

Following Vardi (1996), \mathbf{Y} and \mathbf{X} are related through the $r \times c$ routing matrix $\mathbf{A} = \{A_{s,a}\}$ where $A_{s,a} = 1$ if the directed link s belongs to the directed route through the network between origin-destination pair a , $A_{s,a} = 0$ otherwise. Note that \mathbf{A} is typically singular, with a number of columns larger than the number of rows. With this definition, and with the links implicitly ordered $1, \dots, r$, we have the defining identity

$$\mathbf{Y} = \mathbf{A}\mathbf{X}. \tag{1}$$

This simply expresses each link count Y_s as the sums of the origin-destination counts for all routes that include directed link s . From an algebraic perspective, \mathbf{Y} imposes a set of linear constraints on \mathbf{X} . The routing matrix summarises the network traffic structure in useful ways. For example, $(\mathbf{A}\mathbf{A}')_{a,a}$ counts the number of origin-destination routes passing through link a , and $(\mathbf{A}\mathbf{A}')_{a,b}$ counts the number of routes that pass through both links a and b . Similarly, $(\mathbf{A}'\mathbf{A})_{s,s}$ is the number of links in route s , and $(\mathbf{A}'\mathbf{A})_{s,t}$ the number of links that routes s and t share in common.

A realistic problem will have a routing matrix that has no duplicate columns and in which each column has at least one non-zero entry. This relates to an identification issue in the estimation approach of Vardi (1996, Section 2), though is a practically relevant constraint independent of issues of estimation. A duplicate column means a redundancy in origin-destination pair specification. A zero column is disallowed, as it would imply a non-communicating origin-destination pair. Similarly, a zero row would imply a link lying outside the network, so is also disallowed.

Vardi (1996) gives several example networks. One such is the $n = 4$ node network in Figure 1. This network, analysed further below, has $r = 7$ directed links and $c = 12$ origin-destination pairs; the ordered sequence of nodes comprising these links and origin-destination routes appears in Figure 2. The corresponding 7×12 routing matrix is

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




Figure 1. An example four node network.

- | | |
|-------------------------|--|
| $Y_1 : A \rightarrow B$ | $X_1 : A \rightarrow B$ |
| $Y_2 : B \rightarrow A$ | $X_2 : A \rightarrow C$ |
| $Y_3 : A \rightarrow C$ | $X_3 : A \rightarrow C \rightarrow D$ |
| $Y_4 : B \rightarrow C$ | $X_4 : B \rightarrow A$ |
| $Y_5 : C \rightarrow B$ | $X_5 : B \rightarrow C$ |
| $Y_6 : C \rightarrow D$ | $X_6 : B \rightarrow A \rightarrow C \rightarrow D$ |
| $Y_7 : D \rightarrow C$ | $X_7 : C \rightarrow B \rightarrow A$ |
| | $X_8 : C \rightarrow B$ |
| | $X_9 : C \rightarrow D$ |
| | $X_{10} : D \rightarrow C \rightarrow B \rightarrow A$ |
| | $X_{11} : D \rightarrow C \rightarrow B$ |
| | $X_{12} : D \rightarrow C$ |

Figure 2. The link and route structure for the network in Figure 1.

2.2 Prior Models for \mathbf{X}

Our problem is that of computing and summarising the joint posterior distribution $p(\mathbf{X}|\mathbf{Y})$ for all route counts \mathbf{X} given the observed link counts \mathbf{Y} . This requires a model for the prior distribution $p(\mathbf{X})$ to be tied together with the deterministic expression (1) that implies \mathbf{Y} given \mathbf{X} . Consider first the kinds of models we might assume for $p(\mathbf{X})$.

The statistical framework of Vardi (1996) involves an assumption that the origin-destination traffic \mathbf{X} is generated from a collection of independent Poisson distributions for the elements X_a . Assuming $X_a \sim Po(\lambda_a)$, independently over routes a , Vardi proceeds to explore likelihood and method of moments estimation approaches to inference on the Poisson rates $\Lambda \stackrel{\text{def}}{=} \{\lambda_1, \dots, \lambda_c\}$. Both Poisson and independence assumptions have been common in transportation network flow applications, though both are questionable in many contexts. In developing prior estimates of origin-destination matrices in connection with network equilibrium models, for example (Berka and Boyce 1994, Sheffi 1995) conditional Poisson models are typical although dependencies between flows are often strong and may be modelled through hierarchical models for the rates Λ (e.g. Sen 1987, West 1994). Thus, for example, West (1994) develops Bayesian analyses under the Poisson assumption but with random effects regression models relating the Poisson rates λ_a . His models, estimated based on survey data, incorporate origin and, separately, destination specific effects, together with interaction terms that modify the Poisson rates as a function of distance (in terms of estimated travel times) between origin and destination, and with additional route specific random effects to model residual patterns of interactions and unexplained, extra-Poisson variability. Thus more complex prior models than the independent Poisson forms are available by introducing structure among the Poisson rates through more or less standard hierarchical components. These kinds of models can be incorporated in our development below. For the current study, however, we restrict attention to the independent Poisson structure to simplify presentation of the key developments and to tie in directly with the previous work of Vardi.

Thus we assume, following Vardi (1996), that $X_a \sim Po(\lambda_a)$, independently over a . Prior specification is completed by a prior for Λ , determining a joint model

$$p(\mathbf{X}, \Lambda) = p(\Lambda) \prod_{a=1}^c \lambda_a^{X_a} \exp(-\lambda_a) / X_a!, \quad (2)$$

as the starting point for analysis. Though we are primarily interested in estimating \mathbf{X} , estimation of Λ will be an enabling activity and we will jointly infer \mathbf{X} and Λ from (2) as a result. Begin with the assumption that the Poisson rates are independently drawn from specified marginal prior distributions $p(\lambda_a)$ for each a . As mentioned above, this is consistent with more realistic models, for some applied contexts, in which the priors $p(\lambda_a)$ will themselves be further structured hierarchically in terms of hyperparameters.

2.3 Conditional Posteriors for MCMC Posterior Simulation

Given the prior (2), we now observe the link counts \mathbf{Y} and will condition to deliver the required posterior $p(\mathbf{X}, \Lambda|\mathbf{Y})$. Naturally, posterior computations are difficult analytically in any other than trivial and quite unrealistic networks, so that our approach develops iterative MCMC simulation methods. Consider, in particular, Gibbs sampling, in which we iteratively resample from conditional posteriors for elements of the \mathbf{X} and Λ variables.

First consider Λ . We note that

$$p(\Lambda|\mathbf{X}, \mathbf{Y}) \equiv p(\Lambda|\mathbf{X}) = \prod_{a=1}^c p(\lambda_a|X_a),$$

whose components have the form of the prior density $p(\lambda_a)$ multiplied by the gamma form arising in the Poisson-based likelihood function. Thus, conditional on \mathbf{X} , we can easily simulate new Λ values as a set of independent draws from the implied univariate posteriors. If $p(\lambda_a)$ is gamma, or a mixture of gammas, these draws are trivially made from the corresponding gamma or mixture of gamma posteriors. Otherwise, we may use a rejection method or embed Metropolis-Hasting steps in the MCMC scheme in the by now standard ‘‘Metropolis-within-Gibbs’’ framework (e.g. see the many examples in Gilks *et al* 1996).

Now turn to the conditional posterior $p(\mathbf{X}|\Lambda, \mathbf{Y})$, viewing Λ as fixed. Our data \mathbf{Y} are in the form of a set of linear constraints (1) on the route count vector \mathbf{X} , so that conditioning must be performed directly, algebraically, rather than via the usual application of Bayes’s theorem. In the following sub-section we develop a completely general, and automatic approach to analysis of this conditional posterior. Here, to motivate that development and generate additional insights, we detail the structure of the posterior in a very simple example.




Figure 3. A simple three node network with just two directed links.

Take the network in Figure 3, just two directed links between nodes A,B and C, and with routing matrix

$$\mathbf{A} = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \end{pmatrix}.$$

This has three routes $A \rightarrow B$, $B \rightarrow C$ and $A \rightarrow C$, with unobserved counts X_1, X_2 and X_3 respectively, and links counts Y_1 on $A \rightarrow B$ and Y_2 on $B \rightarrow C$. Trivially, we see that knowledge of X_3 in addition to \mathbf{Y} implies that both X_1 and X_2 are known, i.e. $X_1 = Y_1 - X_3$ and $X_2 = Y_2 - X_3$. Furthermore, under the independent Poisson priors for the X_a it easily follows that the marginal posterior density of X_3 is

$$p(X_3|\Lambda, \mathbf{Y}) \propto p(X_3|\lambda_3) \prod_{a=1,2} p(X_a|\lambda_a) I_{Y_a - X_3}(X_a)$$

where $I_y(x)$ is the indicator of $x = y$ for fixed y . This reduces to

$$p(X_3|\Lambda, \mathbf{Y}) \propto \frac{\lambda_3^{X_3}}{X_3!} \prod_{a=1,2} \frac{\lambda_a^{Y_a - X_3}}{(Y_a - X_3)!}$$

over the support $X_3 = 0, 1, \dots, \min(Y_1, Y_2)$. Hence the full joint posterior $p(\mathbf{X}|\Lambda, \mathbf{Y})$ can be simulated by drawing X_3 from the above distribution, then directly evaluating X_1 and X_2 .

This example, though trivial almost to the point of degeneracy, illuminates the general structure of the problem of conditioning on \mathbf{Y} to deduce $p(\mathbf{X}|\Lambda, \mathbf{Y})$. In general, the r vector \mathbf{Y} and defining equation (1) implies a set of r linear constraints on \mathbf{X} . As a result, the posterior for \mathbf{X} may be reduced to set of r linear equations that deliver precise values of r elements of \mathbf{X} given specified values of the remaining $c - r$; the marginal posterior for this remaining $c - r$ elements may be directly evaluated and used as the centerpiece for posterior simulation in the MCMC framework. We move on with such developments below, but first

describe and exemplify some further aspects of the model, and of resulting issues of ambiguity of inferences in the absence of additional constraints or truly informed priors.

2.4 Structural Ambiguity and the Need for Informative Priors

In raw algebraic form the data \mathbf{Y} simply provides a set of linear constraints on the unobserved flows \mathbf{X} . Under the assumed independent Poisson priors for the X_i , and conditional on the values of the Poisson means λ_i , the data simply constrains the conditional prior $p(\mathbf{X}|\lambda)$ to the support defined by $\mathbf{Y} = \mathbf{A}\mathbf{X}$, so defining “ridges” in the resulting likelihood function for Λ . This is best seen through simple examples, so we continue with the example of Figure 3 above. Here we have constraints $Y_1 = X_1 + X_3$ and $Y_2 = X_2 + X_3$. Assume that the underlying mean OD flows on links one and two are the same, $\lambda_1 = \lambda_2 = \lambda$, and are distinct from $\lambda_3 = \mu$. Clearly, there is an identification problem in that we can write $Y_1 = X_1 + X_3 = (X_1 + a) + (X_3 - a)$ and $Y_2 = X_2 + X_3 = (X_2 + a) + (X_3 - a)$ for any integer a less than or equal to $\min(Y_1, Y_2)$. This impacts the resulting likelihood function for (λ, μ) by inducing a diagonal ridge running North-West to South-East, reflecting the strong negative dependence induced by the constraints. Variation in the likelihood along this diagonal ridge depends on the actual data values \mathbf{Y} observed, and can lean towards either extreme or more highly support central values. Particularly critical are cases in which X_1 and X_2 are rather large and happen to have close values, but X_3 is small. The likelihood function based on large and similar values of (Y_1, Y_2) then tends to favour the central part of the ridge, so apparently over-estimating μ and under-estimating λ . This is partly induced by the relatively higher dispersion of the Poisson densities at higher levels, which leads to higher likelihood values for small means when “close” values of the Poisson counts are observed. As a result, analysis with uniform priors on μ and λ will lead to posteriors that tend to over-estimate low rates, and the corresponding OD flows, and under-estimate high rates and flows, though the extent of these biases is much more marked in the over-estimation of low flows. Thus purely likelihood-based analyses are subject to such biases without additional constraints.

This phenomenon occurs quite generally and is responsible for gross over-estimation of low OD flows and their means in some analyses, such as discussed in our application in Section 4. That application concerns a much more complex network than the simple example above, and several links with rather low OD flows are tied together in collections of OD routes which experience very high flow rates, leading to the gross over-estimation of the low flow rates. This argues against the use of uniform priors and suggests that the standard practice in transportation studies, i.e., that of updating initial or prior estimates of OD flows based on previous data and experience (McNeil and Hendrickson 1985, Maher 1983), is partly helping to constrain the problem to overcome these kinds of identification difficulties. We note further that alternative estimation methods, such as that of Vardi (1996), that do not recognise this problem are immediately subject to the same kinds of biases in resulting inferences.

2.5 Algebraic Structure of General Network

We follow the above example and discussion of Section 2.3 with further exploration of the theoretical structure of a general network. This leads to a theoretical result that underlies a generally applicable method of computing samples from conditional posteriors for route counts in an iterative Gibbs sampling framework. In addition to directly enabling simulation based analysis, the discussion is of independent interest in elaborating further on the structure of the network problem.

Consider Λ fixed and known, and focus attention on the conditional posterior $p(\mathbf{X}|\Lambda, \mathbf{Y})$. The following result, that is simply an algebraic deduction from the network structure and defining relation (1) among the traffic counts, is key and critical to the ensuing inferential development.

Theorem 1. *In the network model (1), assume that \mathbf{A} is of full rank r . Then we can reorder the columns of \mathbf{A} so that the revised routing matrix has the form*

$$\mathbf{A} = [\mathbf{A}_1, \mathbf{A}_2] \quad (3)$$

where \mathbf{A}_1 is a non-singular $r \times r$ matrix. Also, similarly reordering the elements of \mathbf{X} vector and conformably partitioning as $\mathbf{X}' = (\mathbf{X}'_1, \mathbf{X}'_2)$, it follows that

$$\mathbf{X}_1 = \mathbf{A}_1^{-1}(\mathbf{Y} - \mathbf{A}_2\mathbf{X}_2). \quad (4)$$

Proof: The partitioning of \mathbf{A} in (3) is an essentially trivial consequence of the full rank assumption. Matrix \mathbf{A} is $r \times c$ with $r < c$, and so has at least r linearly independent columns; simply reordering so that the first r columns are among the linearly independent subset produces (3). Then reorder the elements of \mathbf{X} to correspond to the new ordering of the columns of \mathbf{A} . This gives $\mathbf{Y} = \mathbf{A}\mathbf{X} = \mathbf{A}_1\mathbf{X}_1 + \mathbf{A}_2\mathbf{X}_2$, which, noting the invertibility of \mathbf{A}_1 , leads easily to (4).

The full rank assumption is satisfied by all practical networks. Otherwise, there is a redundancy in specification and one or more rows of \mathbf{A} that can be deleted to reduce to linear independence of the rows, hence full rank of the reduced matrix. Given the partitioned form (3), and with the implicit reordering of elements of \mathbf{X} , the result (4) implies that, given \mathbf{Y} and assumed values of the $(c - r)$ route counts in \mathbf{X}_2 , we can directly compute the remaining r route flows, simply based on the algebraic structure of the defining routing matrix. In some cases, we may be able to directly identify \mathbf{A}_1 by inspection. More often, especially in larger networks, we need an algorithm to appropriately reorder routes and deduce the resulting routing matrix in the form (3) automatically. This may be done as follows.

Recall the QR decomposition of arbitrary, full-rank, matrices (e.g. Golub and Van Loan 1983, Section 6.2). Applied to the $r \times c$ routing matrix \mathbf{A} , this delivers the QR decomposition in terms of an $r \times r$ orthogonal matrix \mathbf{Q} and an $r \times c$ upper triangular matrix \mathbf{R} such that $\mathbf{A} = \mathbf{Q}\mathbf{R}$. The first r columns of \mathbf{R} correspond to r linearly independent columns of \mathbf{A} ; they are identified by a permutation of column indices that are typically delivered by iterative procedures for developing the QR decomposition. For example, the efficient Householder successive reflection procedure, as implemented in the `qr()` routine in S-Plus (1993), delivers a list of indices of columns of \mathbf{A} in permutation such that the first r elements identify the linearly independent columns of \mathbf{A} . As a result, we can very simply, routinely and automatically, identify a reordering of the columns of \mathbf{A} to achieve the form (3) with \mathbf{A}_1 of full rank: simply apply the QR decomposition and extract the column permutation vector, then reorder the columns of \mathbf{A} according to this vector. In passing, note that \mathbf{A}_1 is simply the first r columns of the matrix \mathbf{R} .

In our example network of Section 2.1, the `qr(A)` call in S-Plus produces, among other things, the vector pivot of the permuted column indices, in this case 1, 2, 3, 4, 5, 7, 10, 11, 12, 9, 8, 6. Reordering columns 1 – 12 accordingly leads to the revised routing matrix

$$\mathbf{A} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 1 \\ 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 \end{pmatrix}.$$

It is easily seen that the first 7×7 block is non-singular with inverse

$$\mathbf{A}_1^{-1} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 & -1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & -1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}.$$

With this general Theorem 1 and the resulting constructive method of applying it, we are now in a position to develop the conditional posteriors for route counts in a Gibbs sampling framework.

2.6 Conditional Posteriors for Route Counts

From the results of the previous section, we see that, for any Λ and fixed link counts \mathbf{Y} , the conditional distribution $p(\mathbf{X}|\Lambda, \mathbf{Y})$ is concentrated in a sub-space of dimension $c - r$ defined by the partition (3) of the routing matrix. Having reordered the columns of \mathbf{A} to the form (3), this posterior has the form

$$p(\mathbf{X}_1|\mathbf{X}_2, \Lambda, \mathbf{Y})p(\mathbf{X}_2|\Lambda, \mathbf{Y})$$

where $p(\mathbf{X}_1|\mathbf{X}_2, \Lambda, \mathbf{Y})$ is degenerate at $\mathbf{X}_1 = \mathbf{A}_1^{-1}(\mathbf{Y} - \mathbf{A}_2\mathbf{X}_2)$, and with $\mathbf{X}_2 = (X_{r+1}, \dots, X_c)'$ defining $\mathbf{X}_1 = (X_1, \dots, X_r)'$ as above,

$$p(\mathbf{X}_2|\Lambda, \mathbf{Y}) \propto \prod_{a=1}^c \frac{\lambda_a^{X_a}}{X_a!} \quad (5)$$

over the support defined by $X_a \geq 0$ for all $a = 1, \dots, c$. This is simply the expression of the product of independent Poisson priors for the X_i constrained by the identity (1) rewritten in the form (4). The utility of this expression is in delivering the set of complete conditional posteriors for elements of the \mathbf{X}_2 vector to form part of the iterative simulation approach to posterior analysis. Consider each elements X_i of \mathbf{X}_2 , ($i = r + 1, \dots, c$), and write $\mathbf{X}_{2,-i}$ for the remaining elements. Then, simply by inspection of equation (5), we see that the conditional distribution $p(X_i|X_{-i}, \Lambda, \mathbf{Y})$ is given by

$$p(X_i|\mathbf{X}_{2,-i}, \Lambda, \mathbf{Y}) \propto \frac{\lambda_i^{X_i}}{X_i!} \prod_{a=1}^r \frac{\lambda_a^{X_a}}{X_a!} \quad (6)$$

over the support defined by $X_i \geq 0$ and $X_a \geq 0$ for each $a = r + 1, \dots, c$; this holds for each $i = r + 1, \dots, c$.

Identifying the support of (6) requires study of the (at most r) linear constraints on X_i defined by $X_a \geq 0$ for all elements X_a of $\mathbf{X}_1 = \mathbf{A}_1^{-1}(\mathbf{Y} - \mathbf{A}_2\mathbf{X}_2)$. Given i in $r + 1, \dots, c$, that is given a single component of the \mathbf{X}_2 vector, this implies a set of linear constraints as functions of the conditioning values of $\mathbf{X}_{2,-i}$ and \mathbf{Y} . The resulting constraints are the form $X_i \geq d_i$ or $X_i \leq e_i$, where the values d_i, e_i are functions of the conditioning values of $\mathbf{X}_{2,-i}$ and \mathbf{Y} . Hence, together with $X_i \geq 0$, we obtain a set of at most $r + 1$ constraints on X_i . By directly evaluating these constraints and identifying their intersection, we may deduce the range of X_i over which (6) is non-zero, and hence we identify the unnormalised conditional posterior distribution. In networks of even moderate size, this process can become computationally very burdensome, however, a point that is relevant in developing efficient computation in the following section. Finally, it can be shown that the support is a bounded and connected subset of non-negative integers; details are given in appendix. This is important as it enables rather direct simulation analysis, as part of the overall posterior simulation analysis to which we now turn.

3. IMPLEMENTATION AND POSTERIOR SIMULATION

3.1 Gibbs and Metropolis-Hastings Algorithms

Iterative simulation of the full posterior $p(\mathbf{X}, \Lambda | \mathbf{Y})$ is now enabled. Fix starting values of the route counts \mathbf{X} and proceed as follows.

- (a) Draw sampled values of the rates Λ from the c conditionally independent posteriors $p(\lambda_a | X_a)$ detailed in Section 2.3.
- (b) Conditioning on these values of Λ , simulate a new \mathbf{X} vector by sequencing through $i = r + 1, \dots, c$, and at each step sampling a new X_i from (6), with conditioning elements $\mathbf{X}_{2,-i}$ set at their most recent sampled values; at each step \mathbf{X}_1 is explicitly re-evaluated via (4) as a function of the most recently sampled elements of \mathbf{X}_2 .
- (c) Return to (a) and iterate.

This is a standard Gibbs sampling set-up in which the scalar elements of both Λ and \mathbf{X} are resampled from the relevant distribution conditional on most recently simulated values of all other uncertain quantities. Sampling steps in (a) are easy. In the illustration in the following section, the λ_a are assigned independent priors uniform over a fixed, specified and finite range, so that the posteriors are simply truncated gamma distributions, that for λ_a having shape parameter $X_i + 1$ and scale parameter 1; these are trivially sampled. Sampling steps in (b) appears to require evaluation of the support of (6), as previously discussed, and subsequent evaluation of the unnormalised posterior (6) at each step. Sampling may be performed directly, treating (6) as a simple multinomial distribution on this relevant range. This is easy to implement and has been explored in some example networks. In larger, more realistic networks, however, the implied evaluation of (6) across what may be a very large support, at each iteration and for each element X_i , leads to a computational burden that may be excessive when compared to alternative approaches. Moreover, to do this requires that we identify the support of (6) which, as mentioned earlier, can become computationally very burdensome in networks of even moderate size. Hence indirect but very much more efficient simulation methods become of interest.

In particular, more efficient algorithms are based on embedding Metropolis-Hastings steps within the Gibbs sampling framework. Here candidate values of the X_i are generated, at each stage, from suitable proposal distributions, and accepted or rejected according to the usual Metropolis-Hastings acceptance probabilities. Specifically, suppose a specified and fixed proposal distribution with probability mass function $q_i(X_i)$ for each element X_i in step (b). A candidate value X_i^* is drawn from $q_i(\cdot)$ and accepted with probability

$$\min \left[1, \frac{p_i(X_i^*)q_i(X_i)}{p_i(X_i)q_i(X_i^*)} \right]$$

where X_i is the current, most recently sampled value, and $p_i(\cdot)$ is the unnormalised conditional posterior in equation (6). Note that the unnormalised density $p_i(\cdot)$ will be evaluated only at candidate draws, so that, for a given proposal distribution $q_i(\cdot)$, it is not necessary to either identify the actual support of $p_i(\cdot)$ nor to evaluate it completely across the support, in contrast to the approach via direct Gibbs sampling. If we use a proposal distribution $q_i(\cdot)$ whose support includes values of X_i lying outside the support of $p_i(\cdot)$, then candidate draws of such values will automatically be rejected as they lead to zero acceptance probabilities.

From the structure of the network equations in (1), it is possible to identify bounds on each X_i so that a suitable range for the proposal distribution can be computed. For element X_a given $\mathbf{X}_{2,-a}$, X_a of \mathbf{X}_2 , $X_a = 0$ is a gross lower bound whatever the values in $\mathbf{X}_{2,-a}$. For an upper bound, $X_a \leq \min_i \{Y_i - \sum_{j \neq a} A_{i,j} X_j\}$ where the index i runs over the set of links whose counts include X_a , i.e. those links i for which $A_{i,j} = 1$.

Our algorithm is an iterative, trial and error search to identify upper and lower bounds for X_a given $\mathbf{X}_{2,-a}$ using these two values to start. Then, based on the specified bounds, the implied vector \mathbf{X}_1 is recomputed and checked for feasibility, i.e. non-negative values. If any element of \mathbf{X}_1 is negative, the trial value of X_a is either incremented, in searching for the lower bound on its range, or decremented, in searching for the upper bound. This process terminates and delivers the resulting bounds once the \mathbf{X}_1 vector has r non-negative entries.

This approach represents a standard application of “Metropolis-within-Gibbs” and, subject to verification of irreducibility of the resulting Markov chain on the space inhabited by (\mathbf{X}, Λ) , leads to an overall simulation scheme that will ultimately generate samples from the required joint posterior (Tierney 1994). We have experimented with various such approaches, concentrating on efficient independence chain methods. One such method uses a proposal distribution for $q_i(\cdot)$ that is uniform over the identified support of (6); an alternative takes $q_i(\cdot)$ as the conditional Poisson prior $X_i \sim Po(\lambda_i)$. In each case, both generation of candidate values and evaluation of the acceptance probability are essentially trivial. In our examples to date, each of these approaches has been satisfactory in terms of reasonable acceptance rates and convergence, with the two being essentially undistinguished in most cases. Illustration in the next section is based on the Poisson proposal version.

Theoretical assurance that the MCMC algorithm so defined converges, i.e. ultimately generates samples from the true joint posterior $p(\mathbf{X}, \Lambda | \mathbf{Y})$, follows if we can determine that the Markov chain is irreducible. This reduces to the question of whether or not a current value (\mathbf{X}, Λ) can “move” to any other point in the joint parameter space following a finite number of iterations of the scheme (a) and (b). For the elements of Λ there is no problem, assuming continuous priors with fixed support. For the \mathbf{X} , however, the support of the key conditional posteriors (6) is dependent on resampled values of elements of \mathbf{X}_2 , and so changes with iterations; this complicates this issue somewhat. It can be shown, however, that in fact \mathbf{X}_2 is free to move arbitrarily around its parameter space in consecutive iterations, in spite of the support constraints and complications (details are omitted here). Thus the resulting chain is irreducible and convergence assured.

3.2. Illustration

In simple networks, such as that of Figure 3, exact posterior distributions may be computed and this has been used to validate software implementing the MCMC analysis. We display here some summary inferences for the larger example network of Figure 1 and 2.

Taking $\Lambda = \{1, \dots, 12\}$, we simulated route counts \mathbf{X} from the resulting independent Poisson distributions, producing $\mathbf{X} = (2, 2, 0, 8, 5, 7, 6, 4, 9, 7, 17, 11)'$; these lead to the data values of link counts $\mathbf{Y} = (2, 28, 9, 5, 34, 16, 35)'$. The MCMC scheme is initialised with values of \mathbf{X}_2 chosen by inspection of the observed link flows, an easy task in this small network but one that will take more work in larger and more complex cases. Generally, identifying suitable initial values will involve coming up with a non-negative solution to an underdetermined system of equations, which can be easily set-up to automatically generate ranges of possible solutions and hence possible starting values. For example by using Mathematica we can automatically find a (possibly negative) solution to the underdetermined system $\mathbf{Y} = \mathbf{A}\mathbf{X}$, and the basis for the null space associated to the matrix \mathbf{A} . Playing with different linear combinations of the first with the last we were able to easily determine starting values for our vector \mathbf{X} , actually a number of them, so that we could start our simulation from a widely ranging set of points. It should be remarked that integer programming methods could provide an approach to assigning initial values too, although we have not, as yet, pursued this formally. Notice that we are assuming observations without error for the link counts, so

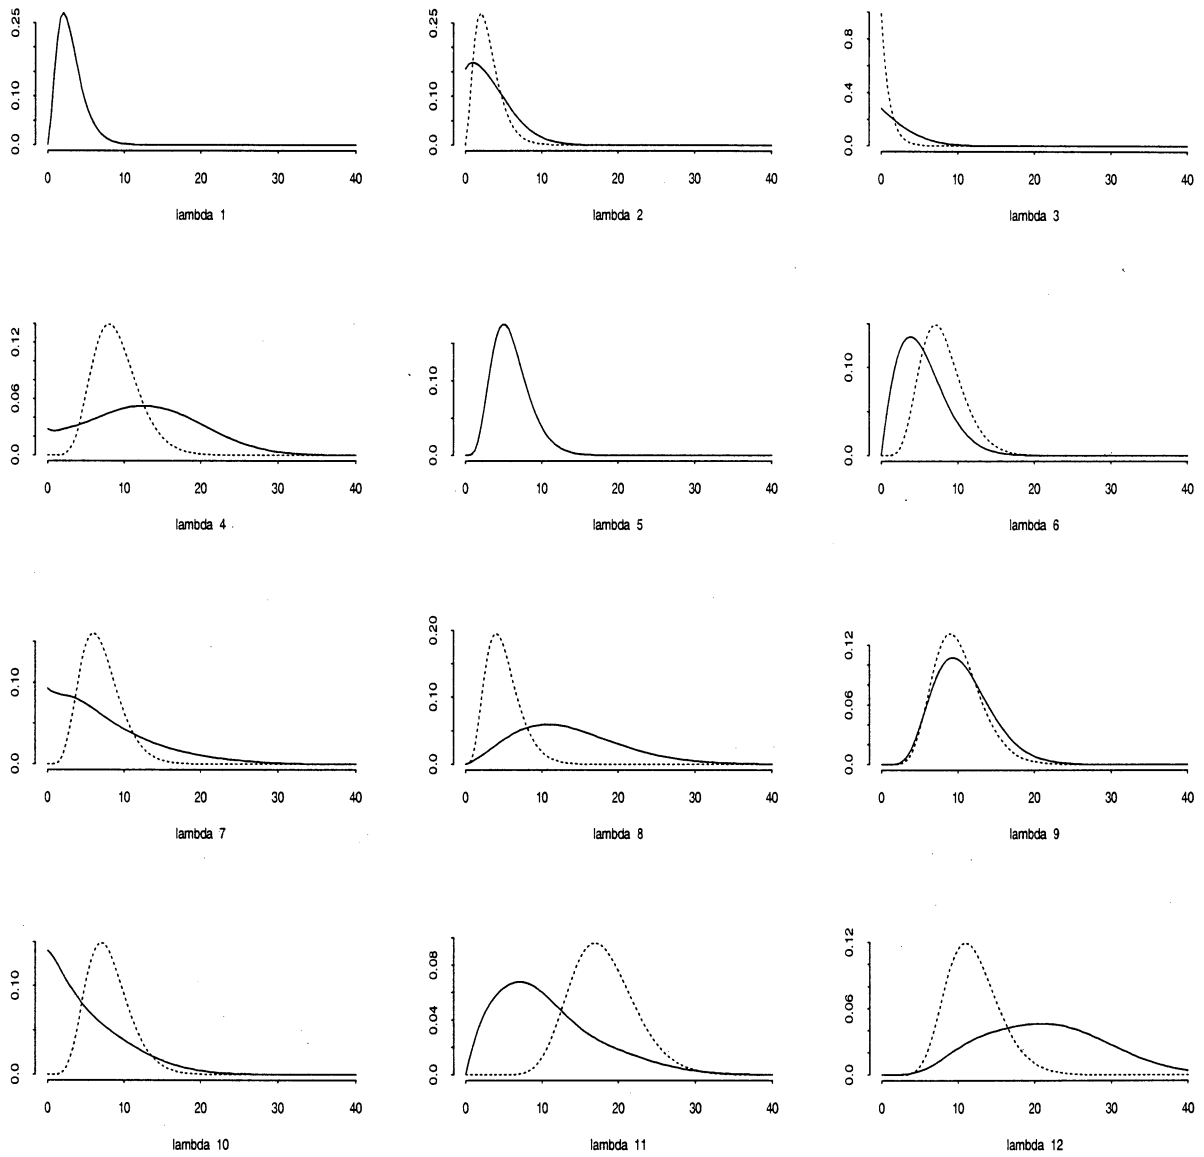


Figure 4. Posterior distributions for the 12 components of Λ , as full lines. The dashed lines are the conditional posteriors were the actual route counts \mathbf{X} known.

we rely on the existence of at list one non negative solution. This ensures, in case other than trivial ones the existence of other possible non negative solutions.

In our analysis, we re-initialised and re-ran the MCMC simulation from these several starting values for \mathbf{X}_2 to assess and validate the rapid convergence, whatever the starting values, experienced. Our reported analysis uses finite uniform priors for the λ_a with common range $(0, L)$ for an appropriately large upper bound L whose value need not be specified. The analysis summarised is based on the Metropolis-Hastings variant with Poisson proposals for the X_a , based on the current values of the respective λ_a . The reported results, and questions of convergence and efficiency, have been explored in various ways. In particular, and most convincing, we have repeated the analysis using several variants of the MCMC scheme; one uses proposal distributions that are uniform on the earlier mentioned ‘large support’, another has uniform proposals on the current support for each X_a conditional on the remaining values $\mathbf{X}_{2,-a}$, and a third reanalysis uses “exact” Gibbs sampling in which the full conditional multinomial distributions of each X_a are evaluated exactly and sampled at each stage. These analyses are in consistent agreement and validate the summary inferences reported using the Poisson proposal. Across all Metropolis-Hastings variant, the rates of acceptance were satisfactorily between 23 and 40 per cent for the different components of vector \mathbf{X}_2 , the differences primarily reflecting the varying degrees of latitude arising in identifying appropriate bounds on the support of proposal distributions.

The above assurances on convergence of the MCMC algorithms are complemented by positive results from various standard convergence investigations as implemented in CODA (Best *et al* 1995). Following this, we summarise analysis by running the MCMC algorithm for a total 50,000 iterations and computing posterior estimates based on only 5,000 values. The full 50,000 iterations are performed rapidly on this small network, running in a real-time of between 2 and 2.5 minutes on an otherwise idle, standard desk-top workstation (DEC AlphaStation 200). How computing time, as well as convergence characteristics, scales up to larger networks is the subject of current investigations. Figures 4 and 5 displays approximate posteriors for Λ and \mathbf{X} respectively. For each λ_a , we compute the posterior density approximation via the Monte Carlo average of known conditional distributions. Since $p(\lambda_a|\mathbf{X}, \mathbf{Y}) = p(\lambda_a|X_a)$ is simply $Ga(X_a + 1, 1)$ truncated to $(0, L)$, we have

$$p(\lambda_a|\mathbf{Y}) = \sum_{X_a} p(\lambda_a|X_a)p(X_a|\mathbf{Y}) \approx m^{-1} \exp(-\lambda_a) \sum_{i=1}^m \lambda^{X_a^{(i)}} / X_a^{(i)!}$$

where $m = 5,000$, the Monte Carlo sample size, and $\{X_a^{(i)}, i = 1, \dots, m\}$ is the set of sampled values of X_a . These densities appear as full lines in Figure 4. As we are dealing with a simulated data set and know the true underlying values of each X_a , we can compute the truncated gammas $p(\lambda_a|X_a)$ for comparison; these appear as the dashed lines in Figure 4, the differences between these and the full posteriors indicative of the uncertainties about the λ_a due to lack of precise knowledge of the X_a . In the cases of $a = 1$ and $a = 5$, the structure of the network implies $X_1 = Y_1 = 2$ and $X_5 = Y_4 = 5$; hence there is no uncertainty about these two X_a values so the posteriors coincide, $p(\lambda_a|X_a) = p(\lambda_a|\mathbf{Y})$. Figure 5 displays histograms of the 5,000 posterior samples of the X_a , representing the marginal posterior $p(X_a|\mathbf{Y})$. Note the degenerate posteriors for X_1 and X_5 as these two values are known. Otherwise, full marginal uncertainties about the route flows are reflected in these histogram approximations to posterior densities; the underlying X_a values from the simulated data are marked as triangles on the lower axis, for comparison.

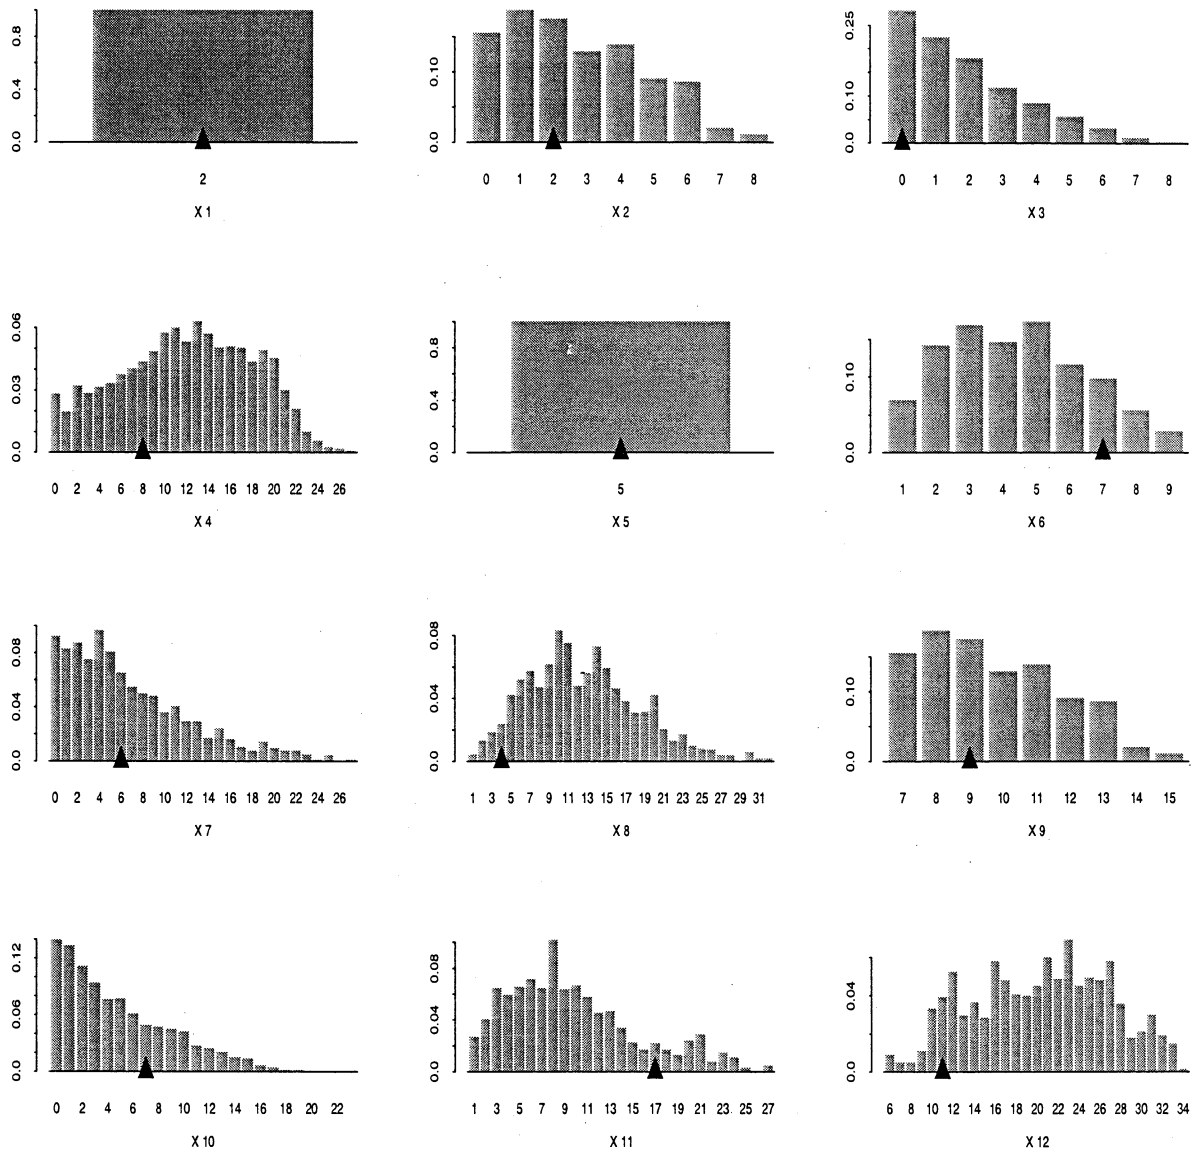


Figure 5. Posterior distributions for the 12 components of X . The actual values are marked by the triangles on the axes.

4. THE MONROE NC TRANSPORTATION NETWORK

Estimation, and prediction, of origin-destination flows in real road networks is a problem central to policy-oriented studies of a range of traffic engineering problems, ranging from real-time traffic management to implementation of congestion-pricing mandates (Sheffi 1985). The archetype context is one of updating estimates of OD flows: based on “today’s” initial OD flow estimates and link traffic since observed, make forecasts of “tomorrow’s” OD flows to feed into traffic management and decision systems, including network equilibrium models that are central tools in transportation modelling and policy studies (Berka and Boyce 1994, Sheffi 1995, Smith 1987, West 1994). This problem has an inherent Bayesian flavour and, though most approaches have been quite ad-hoc, some have adopted basic Bayesian methods (Maher 1983). One interest here is to assess the efficacy of the full Bayesian solution in the context assumed above on a real network. From transportation engineers at North Carolina State University working with the National Institute of Statistical Sciences, we have a section of a road network from the township of Monroe in North Carolina. This network is structured so that all OD pairs have just one route between them, so falls under the fixed routing assumptions of our model. The network is schematically represented in Figure 6. This network has been adapted into the operating “real-time” simulation system *Integration* (Van Aerde et al 1996); this system is extensively used in studying effects on traffic networks of potential network and road traffic policy changes. Based on previous studies of the Monroe road system, OD flows are generated to “load” the network which the runs through a specified time period of the day, and are subject to the usual delays due to traffic congestion, traffic lights, left and right turns, merging of lanes and so on. All aspects of road traffic are continuously monitored and recorded. In particular, observed link counts are recorded on all links in the network. Here we report on some aspects of analyses of observed link counts over a two-hour morning period; this analysis illuminates application of our models, and generates additional insights into the structure and difficulties of the OD estimation problem in a real context but one in which we have the true OD flows to compare with model-based inferences.

In this network, most pairs of the nodes at the boundaries constitute OD pairs; the exceptions are pairs (M,G), (M,L), (G,M), (G,L), (N,L), (L,M), (L,G), (L,N). Observed link counts in the two-hour period appear in Figure 7. After deleting four of the original twenty-four link count observations to eliminate observed linear dependencies, the resulting matrix \mathbf{A} is 20 by 64. The problem thus involves the 44-dimensional vector \mathbf{X}_2 and the 20-dimensional vector \mathbf{X}_1 in the earlier notation.

Our first analysis is exactly as detailed in Sections 2 and 3, using independent uniform priors for the OD rates λ_i . Note that, in a real “on-line” application, observation and analysis from previous days would generate information on which to base more informative priors. This turns out to be absolutely critical in resolving the inherent tendencies towards biased inferences discussed in Section 2.4. As discussed there, this problem is ever-present, but is exacerbated in cases of “disbalanced” flows; notice the wide variability in observed link counts here, indicative of comparable diversity among the underlying OD flows of interest. In such cases, analysis may be subject to significant biases in inferences unless it is appropriately constrained via informed prior distributions.




Figure 6. Physical node-link structure of the Monroe network.

Y_1	:	$A \rightarrow B$:	1980
Y_2	:	$B \rightarrow C$:	1966
Y_3	:	$C \rightarrow D$:	1788
Y_4	:	$D \rightarrow E$:	2600
Y_5	:	$E \rightarrow F$:	2100
Y_6	:	$F \rightarrow E$:	1816
Y_7	:	$E \rightarrow D$:	2880
Y_8	:	$D \rightarrow C$:	2052
Y_9	:	$C \rightarrow B$:	1954
Y_{10}	:	$B \rightarrow A$:	1772
Y_{11}	:	$I \rightarrow B$:	338
Y_{12}	:	$L \rightarrow B$:	284
Y_{13}	:	$H \rightarrow C$:	306
Y_{14}	:	$M \rightarrow C$:	176
Y_{15}	:	$G \rightarrow D$:	68
Y_{16}	:	$N \rightarrow D$:	1000
Y_{17}	:	$O \rightarrow E$:	1104
Y_{18}	:	$B \rightarrow I$:	488
Y_{19}	:	$C \rightarrow M$:	274
Y_{20}	:	$D \rightarrow N$:	926

Figure 7. Link flows for the Monroe network of Figure 6.

We summarise some basic results of this first analysis using uniform priors for the λ_i . The MCMC analysis is run from starting values computed as described in Section 3, and for a large number of iterations in view of the unavoidably high dependencies among OD flows and rates link. After a number of experiments, we ran a final chain for one million iterations, and summarise for posterior inferences here a “dependence-breaking” sub-sample of size 10,000. Across repeat experiments and with different sub-samples the results are consistent. We summarise here marginal posteriors for 16 of the full 64 OD flows, in Figure 8. As, in this case, we actually know the realised OD values, we can assess the accuracy of posterior inferences by comparing the true values, denoted by X_i^* for OD pair i , against the margins; the X_i^* are indicated in the figures. The key point to note here is just how poorly several of the smaller X_i^* values are estimated; they lie way down in the lower tail of their marginal posterior distributions, and so are grossly over-estimated. The higher flows, by comparison, are consistently adequately estimated. This is an example of the general phenomenon discussed in Section 2.4 and is exhibited across repeat analyses of other data from the Monroe network, taken from differing time periods, and also of simulated data. The biased inferences are due to the inherent structural ambiguity in the likelihood function for the Poisson rates.

Now consider re-analysis using more informed priors for the λ_i . We now take independent gamma priors; note that the analysis of Sections 2 and 3 can be developed with minor modification under (conditionally conjugate) gamma priors, and so the easy details are omitted. To mirror an on-line, OD flow “updating” context, we base the priors for “today’s” X_i on the known values X_i^* ; the idea here is that the numbers X_i^* represent estimate based on previous days analyses and observations. Specifically, we choose the prior gamma distribution for λ_i to have shape parameter aX_i^* and scale parameter a for some $a > 0$. Small values of a discount the prior estimate X_i^* and lead to a relatively diffuse prior. One analysis summarised here is based on $a = 0.02$. For each of the 16 OD pairs chosen in the above analysis, we graph the corresponding gamma prior densities (as dashed lines) in Figure 9. Following MCMC-based analysis with these priors, the estimates of corresponding posteriors are computed and also graphed in Figure 9. Figure 10 displays the corresponding marginal posteriors for the OD flows X_i , again with true values X_i^* marked. Evidently, even very weak prior information, roughly “correctly” located based on previous OD estimates, is sufficient to overcome the distortions and biases inherent in “disbalanced” networks. In Figure 11 we summarise all marginal posterior distributions for all 64 OD flows in terms of box plots. We graph box plots for the two analyses: uniform priors and gamma priors on the λ_i , respectively. Superimposing the true X_i^* values, we note uniform consistency of the data with the priors in the latter analysis, and the corresponding uniform correction of the over-estimation bias for smaller flows. We also indicate (with the symbol “V”) the point estimates delivered using the algorithm of Vardi (1996) on this network. As is clear from the figure, this algorithm, though not a direct likelihood-based algorithm, suffers precisely the same problem of over-estimating low flows in the context of dominating high flow rates on subsets of the network, and should be used with caution unless explicitly adjusted to overcome this problem.

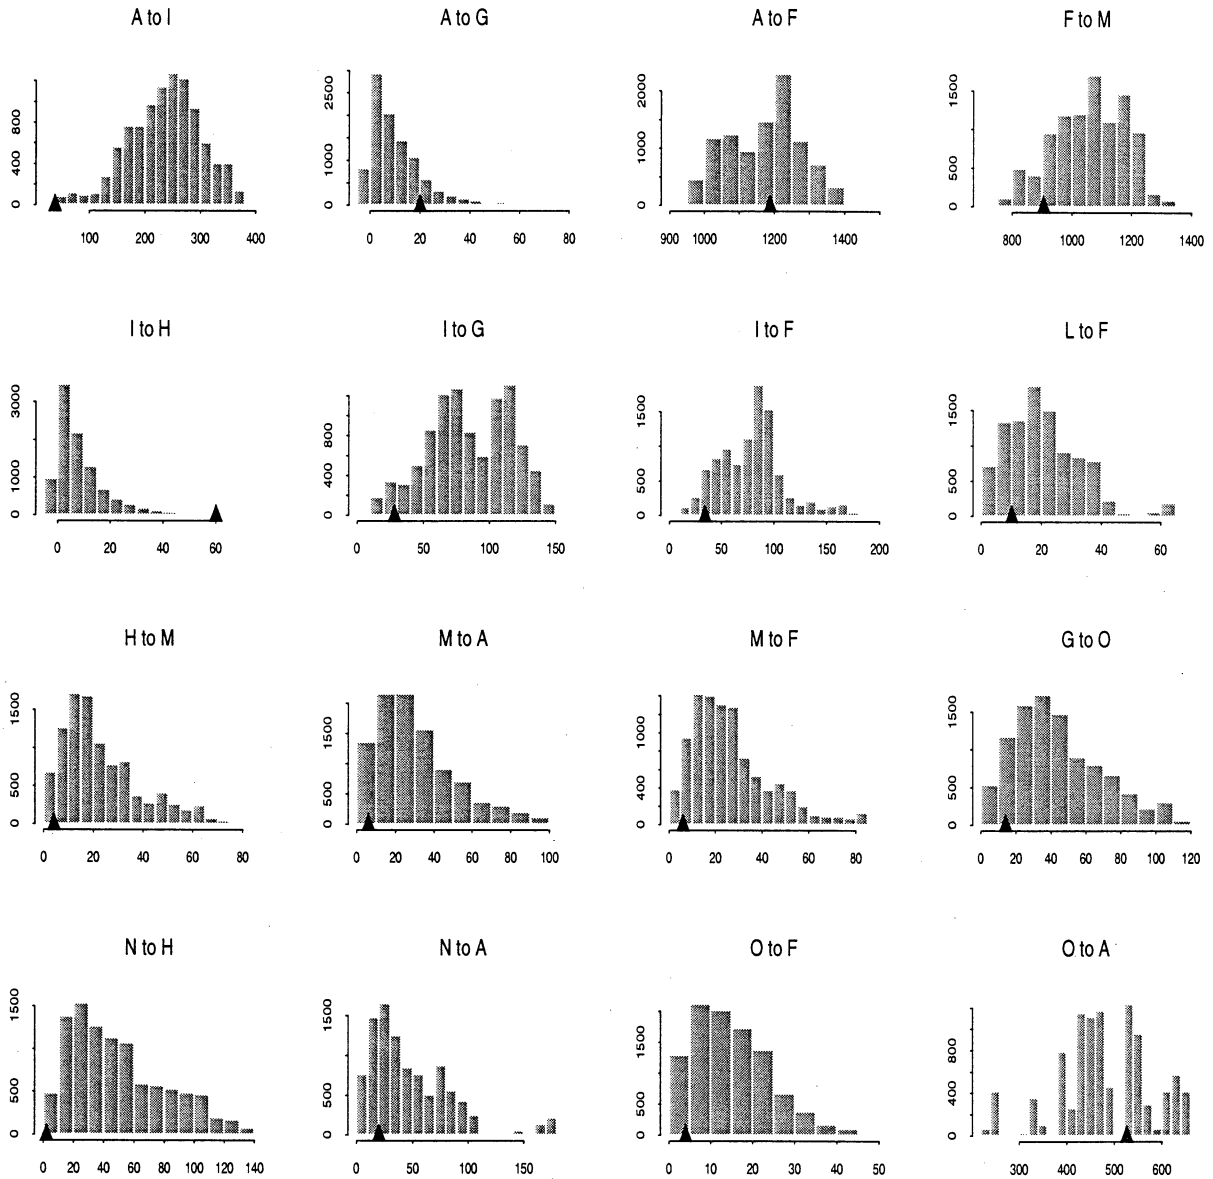


Figure 8. Posterior distributions for 16 components of \mathbf{X} under uniform priors on the λ_i . The true values X_i appear on the x axis.

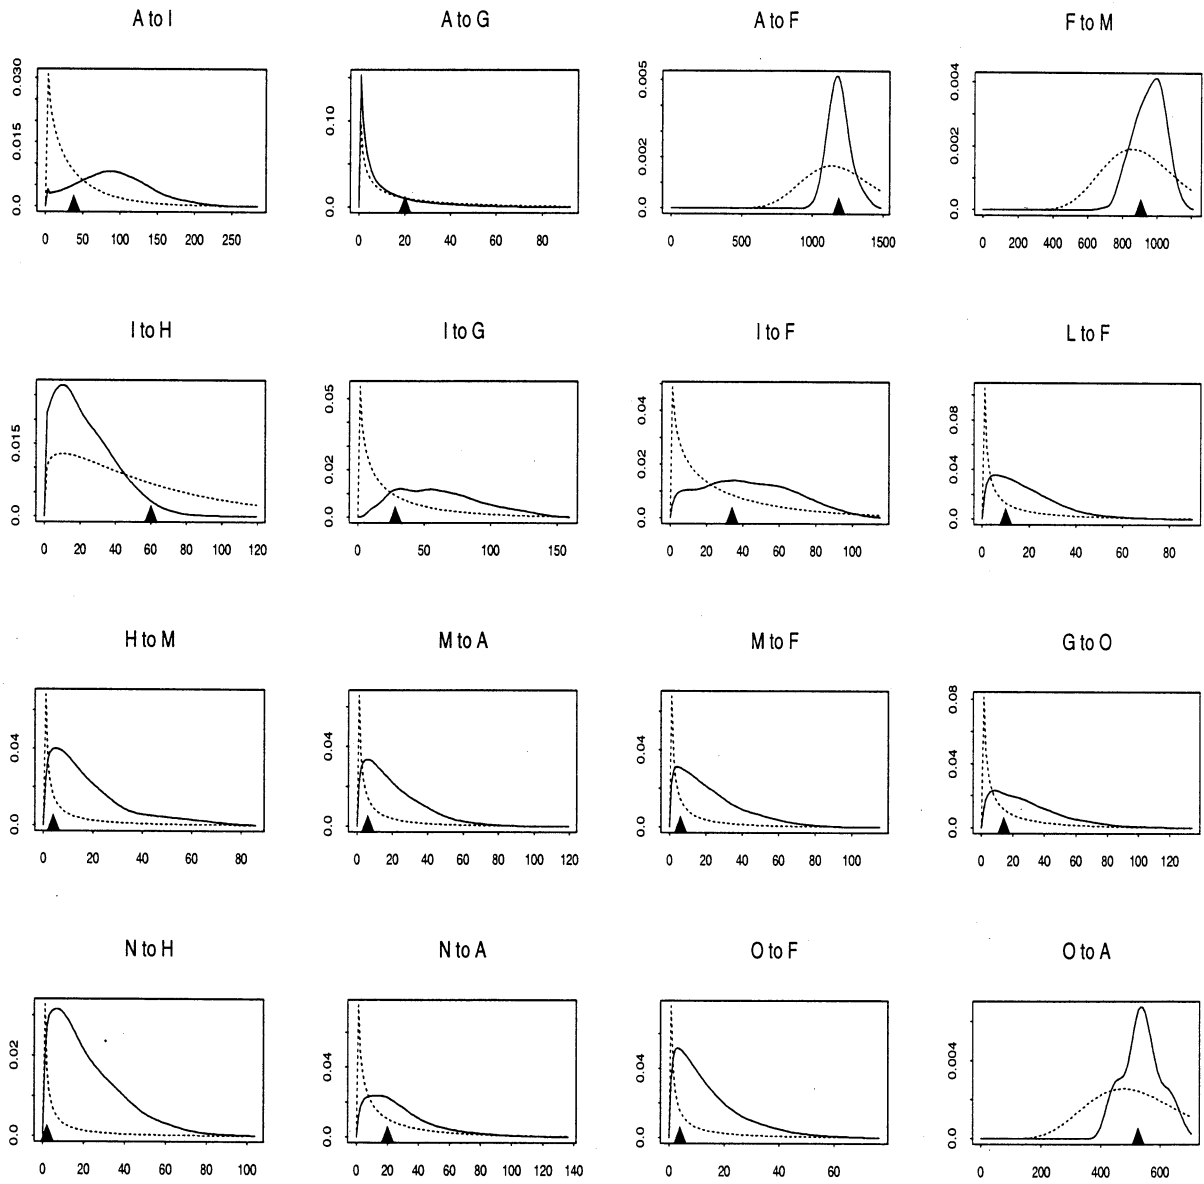


Figure 9. Marginal prior (dashed lines) and posterior (solid lines) densities for 16 components of Λ . The true values X_i appear on the x axis.

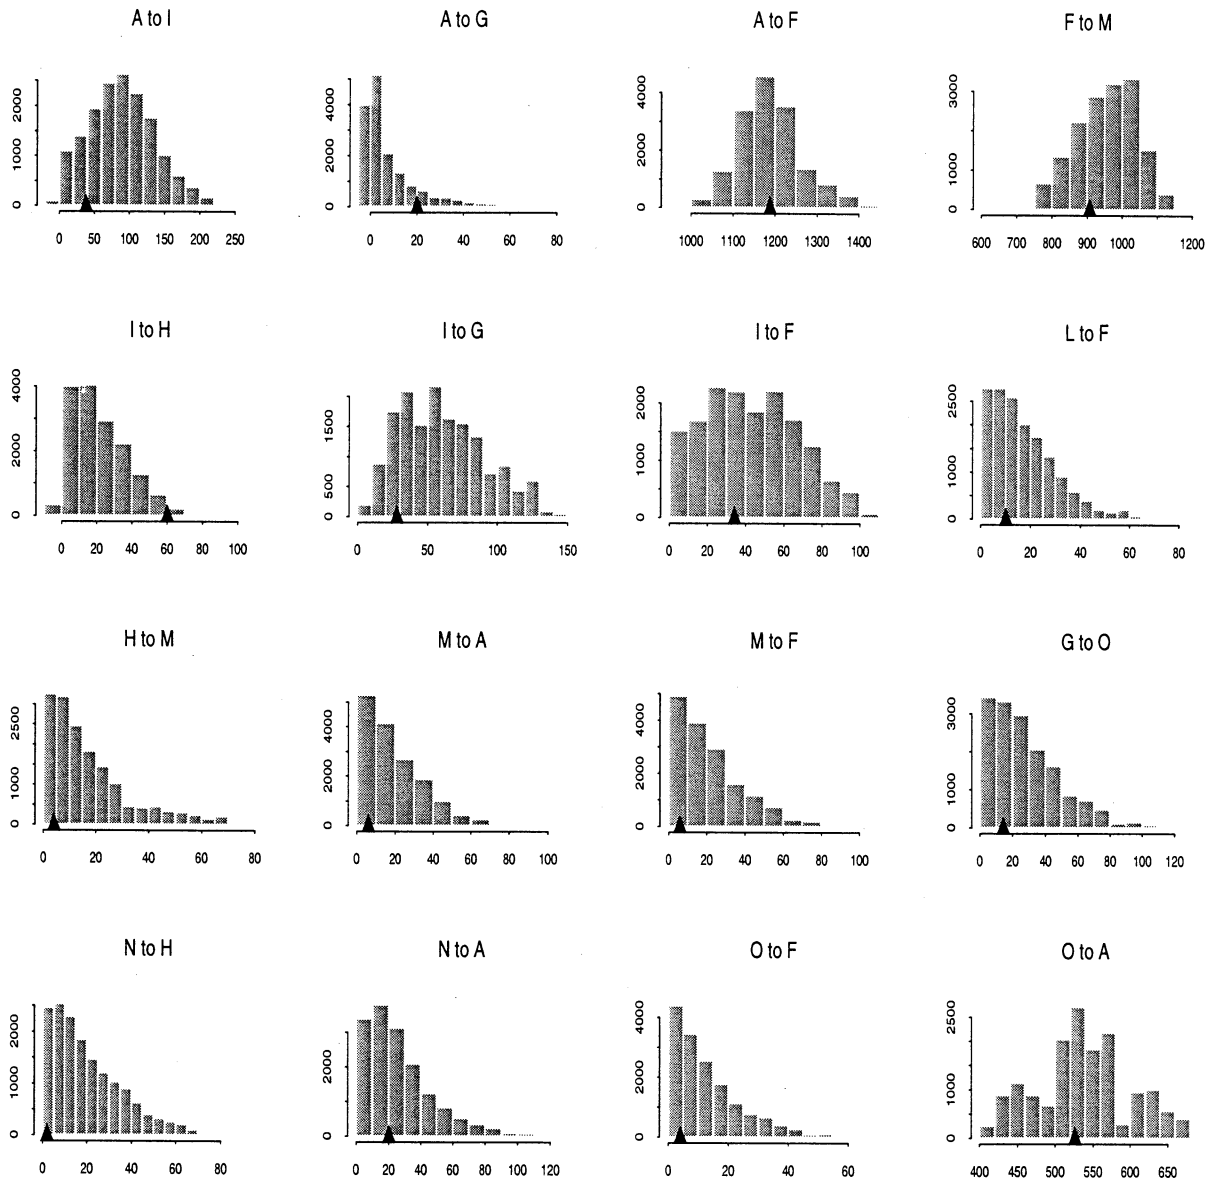


Figure 10. Posterior distributions for 16 components of \mathbf{X} under informed gamma priors on the λ_i . The true values X_i appear on the x axis.

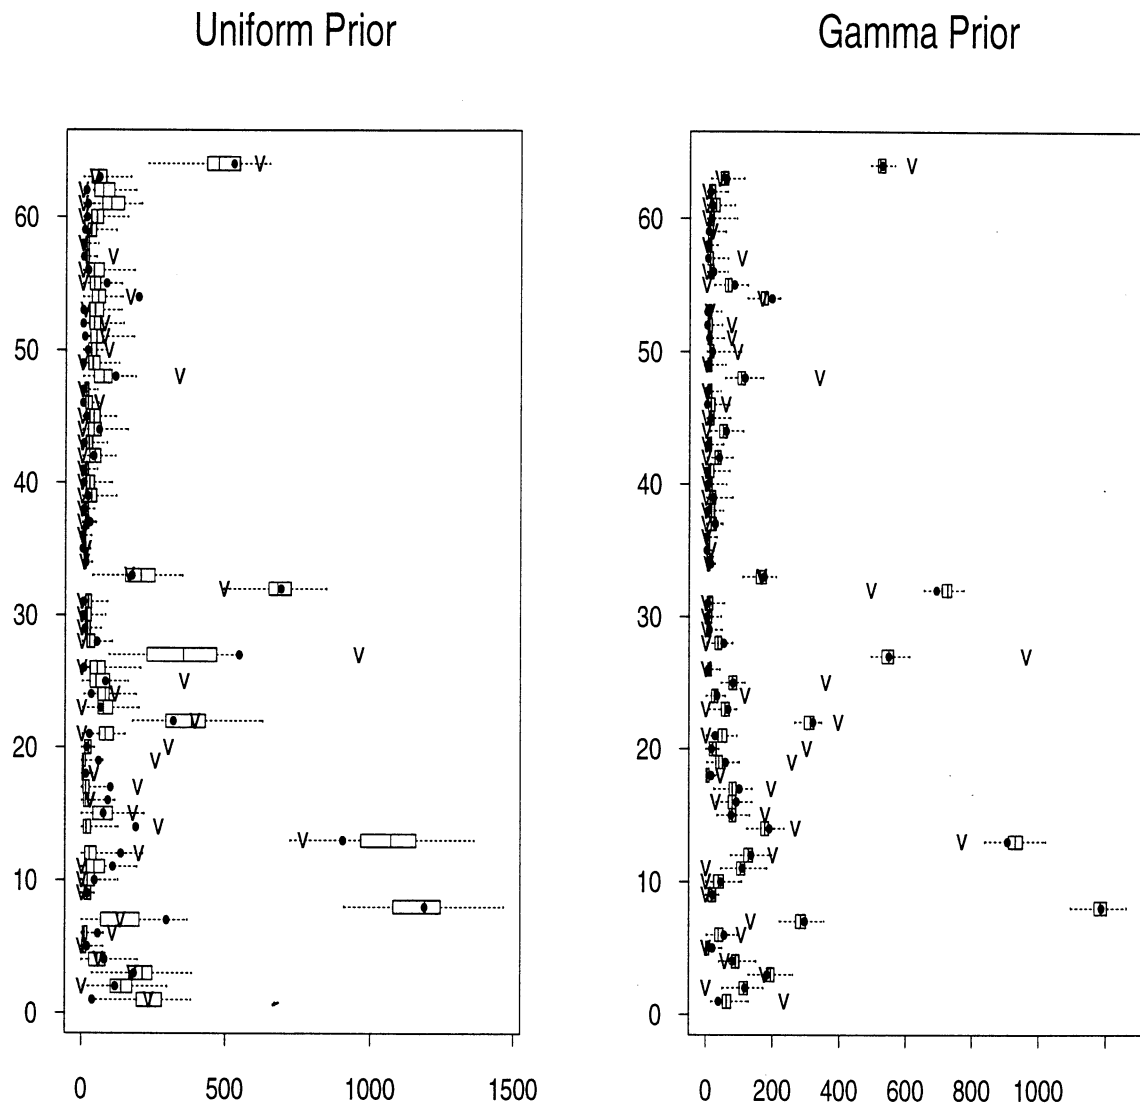


Figure 11. Boxplots of marginal posteriors for all 64 elements of X in the Monroe network. Separate displays are given for analysis based on uniform priors on the λ_i , and based on informed but very diffuse gamma priors. The X_i^* are marked as “•”, and the point estimates of Vardi’s algorithm as “V”.

5. RANDOM (MARKOVIAN) ROUTING

One important model extension relaxes the assumption that all messages or trips between a specified origin-destination pair take the single route specified in the 0/1 routing matrix \mathbf{A} . Vardi (1996) discusses this and develops the case of Markovian routing, in which, a message travelling between a specified origin-destination pair exits its “current” node in the network on one of possibly several links consistent with a route to the specified destination. In communications networks, such random routing may arise as a result of control procedures to avoid or redress queuing questions. In urban road transportation networks, there typically exist one or a small number of primary routes between specified origin-destination pairs, such routes being used by much of the traffic, but with several additional, secondary routes used less frequently either as alternatives in times of congestion on primary routes or as occasional alternatives for other reasons. Primary and secondary routes between a specified origin-destination pair will typically share a common subset of links, with a few additional links being specific to one or a few routes.

A neat extension of the development in Section 2 and 3 provides for analysis using the Bayesian approach developed for the fixed routing network. This is discussed below, following a precise definition and discussion of the random routing model.

Given a specific origin-destination pair $a = (i, j)$, messages leaving node i may now take differing routes to node j . The Markovian routing model simply assumes that, at any node on the way to the destination, each message exits on a link determined by a set of link choice probabilities, independently of the path taken to the current node and independently across messages. In some cases there will be just one exit link possible, in other cases there may be several. As in Vardi (1996), the set-up can be summarised through a modified routing matrix that summaries the link choice probabilities for all possible origin-destination pairs. As in the fixed routing case, the matrix has rows indexing links in the network and columns indexing origin-destination pairs; now, however, the entries are probabilities determining the selection of links (rows) on trips between specified origin-destination pairs (columns). An example matrix from Vardi (1996) is studied further below; with rows and columns labelled by links and origin-destination pairs, this has the probability matrix \mathbf{A} given in Figure 6.

Here we have 4 nodes, $r = 9$ directed links and $c = 12$ origin-destination pairs. Consider, for example, the origin-destination pair AB . Each trip from A to B has an 80% chance of moving directly along link $A \rightarrow B$, and terminating there; with the complementary 20% chance it travels from $A \rightarrow C$. Assuming it follows this latter path, it then travels along $C \rightarrow B$ directly, and terminates, with an 80% chance; otherwise, it moves along the two consecutive links $C \rightarrow D \rightarrow B$ and terminates. By elementary probability computations we can deduce various interesting marginal and conditional probabilities, such as the probability that a message between a specified origin-destination pair j passes through any given link i , or through a specified pair or subset of links.

Now consider the questions of inference about origin-destination counts \mathbf{X} based on observed link counts \mathbf{Y} , as earlier. Retain the same modelling assumptions, so that the X_a are conditionally independent $X_a \sim Po(\lambda_a)$ with specified independent priors on the rates λ_a . Our approach is based on the idea of embedding the random routing problem in a fixed routing problem on an artificial “super-network” and then applying the theory and methods of Sections 2 and 3 to the latter. To motivate this, consider the matrix of Figure 6 and focus on the single origin-destination pair AB . As above, trips from A to B travel one of three routes: $A \rightarrow B$, $A \rightarrow C \rightarrow B$, or $A \rightarrow C \rightarrow D \rightarrow B$, with corresponding marginal probabilities 0.8, 0.16 and 0.04. So we can view the pair AB as comprising three artificial origin-destination pairs corresponding to these three

	1	2	3	4	5	6	7	8	9	10	11	12
	AB	AC	AD	BA	BC	BD	CA	CB	CD	DA	DB	DC
1 $A \rightarrow B$.8	.2	.2	0	0	0	0	0	0	0	0	0
2 $A \rightarrow C$.2	.8	.8	0	1	1	0	0	0	0	0	0
3 $B \rightarrow A$	0	0	0	1	.2	.1	1	0	0	1	0	0
4 $B \rightarrow C$	0	.8	0	0	.8	.1	0	0	0	0	0	1
5 $B \rightarrow D$	0	.2	1	0	0	.8	0	0	1	0	0	0
6 $C \rightarrow B$.8	0	.2	0	0	0	.8	.8	.2	1	1	0
7 $C \rightarrow D$.2	0	.8	0	0	1	.2	.2	.8	0	0	0
8 $D \rightarrow B$	1	0	0	0	0	0	1	1	0	.8	.8	.2
9 $D \rightarrow C$	0	1	0	0	0	0	0	0	0	.2	.2	.8

Figure 12. An example of a probability routing matrix. As in the fixed routing case, rows and columns correspond to directed links and origin-destination pairs, respectively. The entries are now conditional probabilities of traversing the link (row) during a trip between the origin-destination pair (column).

distinct routes. The X_1 trips originating at A and travelling to B can be viewed as initially assigned to one of these three routes according to the marginal route selection probabilities 0.8, 0.16 and 0.04, respectively. Once all assignments are made, the subsets of the X_1 trips travel their allocated routes, and we have a framework of fixed routing.

This leads to the general approach to constructing a super-network with fixed routing. For each column of the probabilistic routing matrix, we identify all possible routes between the corresponding origin-destination pair, and simply list these as a subset of possible fixed routes; the corresponding probabilities assigned to these routes are easily computed. Then we create a new 0/1 fixed routing matrix that has a column for each of these new fixed routes. This can be easily automated (a simple S-Plus function is available from the first author). In the earlier example, the 9×12 probabilistic routing matrix in Figure 2 generates the extended fixed routing matrix in Figure 7. There, in addition to labelling groups of columns with their common origin-destination pair, we add the probabilities of assignment to the columns (i.e. fixed routes) within such subsets. In this case, the original 12 origin-destination pairs generates 27 in the super-network.

Now turn to the issue of modelling and inferring route counts. Consider any origin-destination pair $a = (i, j)$ in the original network, with corresponding counts X_a . The process of creating a corresponding set of fixed routes between i and j generates some number, say k_a of such routes. Given the original probability routing matrix, we can trivially compute the resulting probabilities $\mathbf{p}_a \stackrel{\text{def}}{=} (p_{a,1}, \dots, p_{a,k_a})'$ of selecting each of these fixed routes; note that these k_a probabilities sum to unity. Write $X_{a,t}$ for the number of trips that take fixed route t from i to j . Then, given X_a , the disaggregated counts $X_{a,1}, \dots, X_{a,k_a}$ are conditionally multinomially distributed among the k_a fixed routes, out of the total X_a and with route selection probabilities \mathbf{p}_a . Under the earlier assumption that $X_a \sim Po(\lambda_a)$, this trivially implies that the disaggregated, fixed route counts are themselves marginally independent and Poisson distributed, with $X_{a,t} \sim Po(p_{a,t}\lambda_a)$ for $t = 1, \dots, k_a$. Hence, independently across origin-destination pairs a , we have

$$p(X_{a,1}, \dots, X_{a,k_a} | \lambda_a, \mathbf{p}_a) = \prod_{t=1}^{k_a} \frac{(p_{a,t}\lambda_a)^{X_{a,t}} \exp(-p_{a,t}\lambda_a)}{X_{a,t}!}, \quad (7)$$

	1 AB			2 AC			3 AD			4 BA	5 BC	
	.8	.16	.04	.16	.04	.8	.2	.16	.64	1	.2	.8
1 $A \rightarrow B$	1	0	0	1	1	0	1	0	0	0	0	0
2 $A \rightarrow C$	0	1	1	0	0	1	0	1	1	0	1	0
3 $B \rightarrow A$	0	0	0	0	0	0	0	0	0	1	1	0
4 $B \rightarrow C$	0	0	0	1	0	0	0	0	0	0	0	1
5 $B \rightarrow D$	0	0	0	0	1	0	1	1	0	0	0	0
6 $C \rightarrow B$	0	1	0	0	0	0	0	1	0	0	0	0
7 $C \rightarrow D$	0	0	1	0	0	0	0	0	1	0	0	0
8 $D \rightarrow B$	0	0	1	0	0	0	0	0	0	0	0	0
9 $D \rightarrow C$	0	0	0	0	1	0	0	0	0	0	0	0

Figure 13. Super-network fixed routing matrix for the stochastic network problem of Figure 6. Here we display just the first 12 columns corresponding to the first 5 origin-destination pairs.

	6 BD			7 CA		8 CB		9 CD		10 DA		11 DB		12 DC	
	.1	.1	.8	.8	.2	.8	.2	.2	.8	.8	.2	.8	.2	.2	.8
1 $A \rightarrow B$	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
2 $A \rightarrow C$	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0
3 $B \rightarrow A$	1	0	0	1	1	0	0	0	0	1	1	0	0	0	0
4 $B \rightarrow C$	0	1	0	0	0	0	0	0	0	0	0	0	0	1	0
5 $B \rightarrow D$	0	0	1	0	0	0	0	1	0	0	0	0	0	0	0
6 $C \rightarrow B$	0	0	0	1	0	1	0	1	0	0	1	0	1	0	0
7 $C \rightarrow D$	1	1	0	0	1	0	1	0	1	0	0	0	0	0	0
8 $D \rightarrow B$	0	0	0	0	1	0	1	0	0	1	0	1	0	1	0
9 $D \rightarrow C$	0	0	0	0	0	0	0	0	0	0	1	0	1	0	1

Figure 13 (continued). The final 15 columns, corresponding to the final 7 origin-destination pairs, of the fixed routing matrix for the super-network representation of the stochastic network problem of Figure 6.

which reduces to

$$p(X_{a,1}, \dots, X_{a,k_a} | \lambda_a, \mathbf{p}_a) \propto \lambda_a^{X_a} \exp(-\lambda_a) \prod_{t=1}^{k_a} p_{a,t}^{X_{a,t}} \quad (8)$$

as a function of $(\lambda_a, \mathbf{p}_a)$, noting that $X_a = \sum_{t=1}^{k_a} X_{a,t}$.

The immediate consequence of (7) and (8) is that the original approach to inference about route counts in the fixed routing problem in Sections 2 and 3 applies directly to the super-network. Conditional on the parameters Λ and each of the \mathbf{p}_a vectors, we simply have an expanded network with implied route counts for each of the $\sum_{a=1}^c k_a$ fixed routes. The compatibility of the independent Poisson models for the priors on route counts implies that the construction of conditional posterior distributions for route counts is structurally unchanged. Hence the components of the MCMC analysis for simulating route counts apply and will produce samples of the full set of route counts,

$$\mathbf{X} = \{X_{a,t}, t = 1, \dots, k_a; a = 1, \dots, c\}. \quad (9)$$

From these we can trivially deduce the totals X_a for each origin-destination pair a .

Now turn to inference on the Poisson rates. It is clear that, conditional on all route counts \mathbf{X} in (9), the likelihood function for the underlying rates Λ is just the function (8). This is, again, of exactly the form arising in the original fixed routing problem, and so the construction of posterior samples for Λ follows that development.

This discussion is all conditional on known and fixed routing probabilities, hence known and fixed vectors \mathbf{p}_a for each origin-destination pair a . We note that extensions to incorporate inference on these probabilities are essentially direct. Note that, from (8), we obtain a conditional likelihood function for each of the \mathbf{p}_a vectors given imputed route counts in the super-network and the λ_a . On this basis, the likelihood function factorises into a set of c components of the form in (8), and so the iterative simulation analysis is trivially generalised by linking in a component to sample each \mathbf{p}_a from the corresponding posterior distributions. Assumptions about prior distributions will depend on context, but no essential difficulties arise. Further details will be explored and reported elsewhere.

6. SUMMARY COMMENTS

We have elaborated on the structure of both fixed and random (Markovian) routing problems introduced and analysed by Vardi (1996). In the fixed routing context, our theoretical development has elaborated on aspects of the structure of the network count inference problem that underpin the approach to Bayesian analysis based on MCMC simulation. Our contributions include the isolation of deficiencies in likelihood and other non-Bayesian approaches to inference in these problems due to structural indeterminacies that may lead to systematic and significant biases in inferences in some applications. We demonstrate how this may be remedied through the use of informed prior distributions that provide appropriate ‘‘constraints’’ and note that this strategy is perfectly consistent with current practices in ‘‘updating’’ OD flow estimates in transportation networks (McNeil and Hendrickson 1985, Maher 1983). Other existing algorithms, including that of Vardi, run the unavoidable risk of over-estimating low flows on OD routes that share links with OD routes experiencing high flows. The Bayesian approach here naturally, and satisfactorily, deals with this problem and forms a natural extension of some of the traditional, ad-hoc methods used in traffic engineering solutions.

The embedding of the Markovian routing problem in an elaborated fixed routing context enables similar analyses to be performed there, with evident potential to extend to problems of inference on route choice probabilities. This latter context is very relevant in urban transportation networks, and some developments for application in that area are envisaged. It is clear from the development that the Bayesian view clarifies the structure of these specific network inference problems, and enable relatively direct and easily implementable solutions. Several variations of the specific model assumptions, including the questions of prior models for the underlying Poisson rates Λ , have been mentioned and will represent no substantial additional technical nor computational complications.

The development assumes fully observed link counts in the fixed time period. This is trivially relaxed. Note that, were we to observe counts on only a subset of links, the problem remains essentially the same but with the corresponding rows of the routing matrix \mathbf{A} deleted. In transportation networks, automated traffic detectors, including in-road loop detectors and video cameras, typically monitor traffic on very sparse subsets of key links in large road networks, so this kind of issues is relevant. In large networks, it will be important to focus attention on a subset of likely key routes, i.e. a subset of origin-destination pairs, in addition to a subset of links. Model modifications to partition large scale networks hierarchically into subsets of smaller, manageable networks, and develop inference on route counts and flows based on partial network structure, are then of interest. No such developments exist, currently, to our knowledge.

Additional extensions include questions of monitoring link traffic in several time periods in sequence, and taking into consideration non-negligible transit time over links. We are currently studying link count data from a Seattle highway network as part of a current study of network structure underway at the National Institute of Statistical Sciences (NISS), and will report results of that study elsewhere.

Acknowledgements

The research reported here was partially supported by NSF grant DMS-9313013 to the National Institute of Statistical Sciences (NISS). The authors acknowledge useful discussions with Yehuda Vardi of Rutgers University and Alan F Karr of NISS. The Monroe data sets were provided by Nagui Roupail and Steven Click of North Carolina State University.

APPENDIX

We prove convexity of the support of the uni-dimensional conditional posteriors distributions, and connectedness of the support of the joint posterior, of the vector \mathbf{X}_2 in the fixed routing analysis of Sections 2 and 3.

First, consider the conditional distribution $p(X_{2,a}|\mathbf{X}_{2,-a})$ for any $a = 1, \dots, n-k$. That the corresponding support is convex is proved as follows.

Suppose that two $(n-k)$ -dimensional vectors \mathbf{x}, \mathbf{y} differ in just one of their components, i.e., $\mathbf{x} = (x_1, x_2, \dots, x_i, \dots, x_{n-k})'$ and $\mathbf{y} = (x_1, x_2, \dots, y_i, \dots, x_{n-k})'$. Suppose also that, for any vector \mathbf{Y} , the inequalities $\mathbf{A}_1^{-1}(\mathbf{Y} - \mathbf{A}_2\mathbf{x}) \geq 0$ and $\mathbf{A}_1^{-1}(\mathbf{Y} - \mathbf{A}_2\mathbf{y}) \geq 0$ hold. For any number $\alpha \in [0, 1]$ define the vector

$$\mathbf{z} = \alpha\mathbf{x} + (1 - \alpha)\mathbf{y} = (x_1, x_2, \dots, \alpha x_i + (1 - \alpha)y_i, \dots, x_{n-k})'.$$

Then \mathbf{z} also satisfies $\mathbf{A}_1^{-1}(\mathbf{Y} - \mathbf{A}_2\mathbf{z}) \geq 0$. In fact,

$$\begin{aligned} \mathbf{A}_1^{-1}(\mathbf{Y} - \mathbf{A}_2\mathbf{z}) &= \mathbf{A}_1^{-1}(\mathbf{Y} - \mathbf{A}_2(\alpha\mathbf{x} + (1 - \alpha)\mathbf{y})) \\ &= \mathbf{A}_1^{-1}(\alpha\mathbf{Y} + (1 - \alpha)(\mathbf{Y}) - \mathbf{A}_2(\alpha\mathbf{x} + (1 - \alpha)\mathbf{y})) \\ &= \mathbf{A}_1^{-1}\alpha(\mathbf{Y} - \mathbf{A}_2\mathbf{x}) + \mathbf{A}_1^{-1}(1 - \alpha)(\mathbf{Y} - \mathbf{A}_2\mathbf{y}). \end{aligned}$$

This is the mean of two non-negative quantities by assumption, and so is itself a non-negative quantity. This ends the proof.

We now show that the support of the $n-k$ dimensional vector \mathbf{X}_2 is fully connected, by ruling out the possibility that two or more of its subsets, even if marginally convex, can be separated. The result implies that the MCMC analysis can “smoothly” move around the support, and the chain is irreducible. Begin with the simplified case of a two-dimensional vector $\mathbf{X}_2 = (X_1, X_2)'$. The proof is by contradiction, showing that a configuration such as displayed in Figure 14 is not consistent with our fundamental assumptions: that is, if we assume that both \mathbf{X}_2 and $\mathbf{X}_2 + (1, 1)'$ are feasible, either $\mathbf{X}_2 + (1, 0)$ is feasible as well, or $\mathbf{X}_2 + (0, 1)$ is, if not both.




Figure 14. An example of bi-dimensional support, fully connected in its two marginal components, but not jointly so. In such a case, a Markov chain sampler starting in the bottom-left corner will not visit the region top-right.

In fact, the two pairs of assumptions

$$\begin{aligned} &\{\mathbf{A}_1^{-1}(\mathbf{Y} - \mathbf{A}_2\mathbf{X}_2) \geq 0, \quad \mathbf{A}_1^{-1}(\mathbf{Y} - \mathbf{A}_2(\mathbf{X}_2 + (1, 1)')) \geq 0\} \\ &\{\mathbf{A}_1^{-1}(\mathbf{Y} - \mathbf{A}_2(\mathbf{X}_2 + (1, 0)')) < 0, \quad \mathbf{A}_1^{-1}(\mathbf{Y} - \mathbf{A}_2(\mathbf{X}_2 + (0, 1)')) < 0\} \end{aligned}$$

are incompatible. To begin, note that $\mathbf{A}_1^{-1}(\mathbf{Y} - \mathbf{A}_2(\mathbf{X}_2 + (1, 1)')) \geq 0$ reduces to

$$\mathbf{A}_1^{-1}(\mathbf{Y} - \mathbf{A}_2\mathbf{X}_2 - (A_2^{11}, A_2^{21})' - (A_2^{12}, A_2^{22})') \geq 0$$

where the matrix elements are denoted by double superscript. Now add $\mathbf{A}_1^{-1}(\mathbf{Y} - \mathbf{A}_2\mathbf{X}_2) \geq 0$ by assumption, to lead to

$$\mathbf{A}_1^{-1}(\mathbf{Y} - \mathbf{A}_2\mathbf{X}_2) + \mathbf{A}_1^{-1}(\mathbf{Y} - \mathbf{A}_2\mathbf{X}_2 - (A_2^{11}, A_2^{21})' - (A_2^{12}, A_2^{22})') \geq 0.$$

This reduces to

$$\mathbf{A}_1^{-1}(\mathbf{Y} - \mathbf{A}_2\mathbf{X}_2 + \mathbf{Y} - \mathbf{A}_2\mathbf{X}_2 - (A_2^{11}, A_2^{21})' - (A_2^{12}, A_2^{22})') \geq 0$$

or

$$\mathbf{A}_1^{-1}(\mathbf{Y} - \mathbf{A}_2\mathbf{X}_2 - (A_2^{11}, A_2^{21})' + \mathbf{Y} - \mathbf{A}_2\mathbf{X}_2 - (A_2^{12}, A_2^{22})') \geq 0.$$

This last term cannot then be the sum of two negative quantities, and this establishes the contradiction between the two pairs of inequalities.

The same sequence of steps proves the result for cases when \mathbf{X}_2 and $\mathbf{X}_2 - (1, 1)'$ are feasible, implying then that at least one of $\mathbf{X}_2 - (1, 0)$ or $\mathbf{X}_2 - (0, 1)$ is feasible, (or analogously, \mathbf{X}_2 and $\mathbf{X}_2 + (-1, 1)'$ or \mathbf{X}_2 and $\mathbf{X}_2 + (1, -1)'$ are feasible). As a result, the support is fully connected and irreducibility of the Markov chain sampler is ensured.

In the general multidimensional case we proceed in a similar fashion. By assumption,

$$\mathbf{A}_1^{-1}(\mathbf{Y} - \mathbf{A}_2\mathbf{X}_2) \geq 0$$

and

$$\mathbf{A}_1^{-1}(\mathbf{Y} - \mathbf{A}_2(\mathbf{X}_2 + (1, 1, \dots, 1)')) \geq 0$$

where now the vectors \mathbf{X}_2 , and $(1, 1, \dots, 1)'$ are n -dimensional. Then

$$\mathbf{A}_1^{-1}(\mathbf{Y} - \mathbf{A}_2(\mathbf{X}_2 + (1, 1, \dots, 1)')) \geq 0$$

can be rewritten as

$$\mathbf{A}_1^{-1}(\mathbf{Y} - \mathbf{A}_2\mathbf{X}_2 - (A_2^{11}, A_2^{21}, \dots, A_2^{n1})' - \dots - (A_2^{1n}, A_2^{2n}, \dots, A_2^{nn})') \geq 0$$

and adding $n\mathbf{A}_1^{-1}(\mathbf{Y} - \mathbf{A}_2\mathbf{X}_2) \geq 0$ (by assumption) leads to

$$n(\mathbf{A}_1^{-1}(\mathbf{Y} - \mathbf{A}_2\mathbf{X}_2)) + \mathbf{A}_1^{-1}(\mathbf{Y} - \mathbf{A}_2\mathbf{X}_2 - (A_2^{11}, A_2^{21}, A_2^{n1})' - \dots - (A_2^{1n}, A_2^{2n}, \dots, A_2^{nn})') \geq 0$$

As before, we redistribute the terms to obtain a sum of negative quantities, and hence reach a contradiction: i.e.,

$$\begin{aligned} &\mathbf{A}_1^{-1}(\mathbf{Y} - \mathbf{A}_2\mathbf{X}_2 + \dots + \mathbf{Y} - \mathbf{A}_2\mathbf{X}_2 + \mathbf{Y} - \mathbf{A}_2\mathbf{X}_2 - (A_2^{11}, A_2^{21}, \dots, A_2^{n1})' - \\ &\quad - (A_2^{12}, A_2^{22}, \dots, A_2^{n2})' - \dots - (A_2^{n1}, A_2^{n2}, \dots, A_2^{nn})') = \\ &\mathbf{A}_1^{-1}(\mathbf{Y} - \mathbf{A}_2\mathbf{X}_2 - (A_2^{11}, A_2^{21}, \dots, A_2^{n1})') + \mathbf{A}_1^{-1}(\mathbf{Y} - \mathbf{A}_2\mathbf{X}_2 - (A_2^{12}, A_2^{22}, \dots, A_2^{n2})') + \\ &\quad + \dots + \mathbf{A}_1^{-1}(\mathbf{Y} - \mathbf{A}_2\mathbf{X}_2 - (A_2^{n1}, A_2^{n2}, \dots, A_2^{nn})') \end{aligned}$$

must be non-negative.

This proves irreducibility in the general n -dimensional case.

REFERENCES

- Berka, S., and Boyce, D.E. (1994), "Implementation and solution of a large asymmetric network equilibrium model," *Working Paper*, Urban Transportation Center, University of Illinois at Chicago.
- Best, N.G., Cowles, M.K., and Vines, S.K. (1995), *CODA: Convergence diagnosis and output analysis software for Gibbs sampler output* (Version 0.3), Cambridge UK: Medical Research Council Biostatistics Unit.
- Gilks, W.R., Richardson, S., and Spiegelhalter, D.J. (1996), *Markov Chain Monte Carlo in Practice*, London: Chapman & Hall.
- Golub, G.H., and Van Loan, C.F. (1983), *Matrix Computations*, Baltimore: Johns Hopkins University Press.
- Maher, M. J.(1983), "Inferences on trip matrices from observations on link volumes: A Bayesian statistical approach," *Transportation Research*, 20B, 435-447.
- McNeil, S., and Hendrickson, C. (1985), "A regression formulation of the matrix estimation problem," *Transportation Science*, 19, 278-292.
- Sen, A. (1986), "Maximum likelihood estimation of gravity model parameters," *Journal of Regional Science*, 26, 461-474.
- Sheffi, Y. (1985), *Urban Transportation Networks: Equilibrium Analysis with Mathematical Programming Methods*, New Jersey: Prentice-Hall.
- Smith, T.E. (1987), "Poisson gravity models of spatial flows," *Journal of Regional Science*, 27, 315-340.
- S-Plus (1993), *S-PLUS Guide to Statistical and Mathematical Analysis* (Version 3.2), Seattle: MathSoft, Inc.
- Tierney, L. (1994), "Markov chains for exploring posterior distributions" (with discussion), *Annals of Statistics*, 22, 1701-1762.
- Vardi, Y. (1996), "Network tomography: Estimating source-destination traffic intensities from link data," *Journal of the American Statistical Association*, 91, 365-377.
- West, M. (1994), "Statistical inference for gravity models in transportation flow forecasting," *Discussion paper 94-40*, ISDS, Duke University.
- Van Aerde, M., Hellinga, B., Baker, M., and Rakha, H. (1996), "INTEGRATION: An overview of traffic simulation features," *Technical report*, Department of Civil Engineering, Queen's University, Kingston, Canada.
- Van Zuylen, H. J., and Willumsen, L.G. (1980), "The most likely trip matrix estimated from traffic counts," *Transportation Research*, 14B, 281-293.