Statistical Inverse Problems on Graphs with Application to Flow Volume Estimation in Computer Networks

by

Harsh Singhal

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Statistics)
in The University of Michigan
2009

Doctoral Committee:

Professor George Michailidis, Chair
Associate Professor Anna C. Gilbert
Professor Susan A. Murphy
Professor Vijayan N. Nair
Associate Professor Kerby A. Shedden
Assistant Professor Stilian A. Stoev

To Ma, Papa and Didi.

ACKNOWLEDGEMENTS

TABLE OF CONTENTS

LIST OF FIGURES

CHAPTER I

Introduction

The objective of many physical networks, like computer, road and supply chain networks is to carry flows of different objects; i.e. packets, vehicles and goods in these three types of networks, respectively. Estimating traffic volumes is important for monitoring and provisioning such networks. In this work we look at statistical problems related to estimation of flow volumes, focusing primarily on computer networks.

A computer network, such as the one depicted in Figure 1.1, is comprised of nodes corresponding to network elements such as workstations, routers and switches and links that connect those elements. A network flow contains all the traffic originating at a node and destined for some other node in the network. Each flow can in principle traverse a set of paths connecting its origin and destination, which is determined by the *routing policy*. In computer networks, the flow traffic is carried on packets, whose
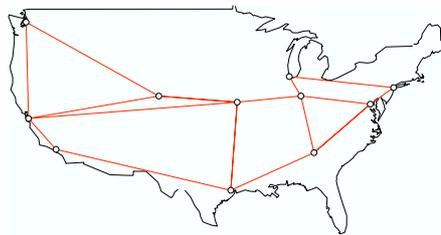
Figure 1.1: Abilene Network Topology

1

payload is expressed in bytes. The volume of traffic refers to either the number of packets and/or the number of bytes in a flow (or on a link) in a given time-interval. An increasing variety of network data is available from modern computer networks. These data differ in their granularity, accuracy, volume and delay [32]. We will primarily be concerned with two kinds of network data. The number of packets and bytes traversing a particular link in a particular measurement interval -typically of the order of a couple of minutes- are available through queries using the Simple Network Management Protocol (SNMP) [37] protocol. The volume of traffic on a link is the sum of volumes of all flows traversing that link. This produces highly aggregate data and the question of interest is to estimate various statistics of the underlying flows. A second kind of measurement is sampled data. Packets of network traffic can be observed (and sampled) at router interfaces. However, during the measurement process sampling is employed due to high flow volumes and resource constraints at routers. It is increasingly common for such measurement infrastructure to be deployed in computer networks [14]. Each packet from the aggregate flow at an observation point is sampled independently with a certain probability (sampling rate) [12]. Typical sampling rates range between .001-.01. Obviously low sampling rates result in large sampling noise. An important issue is how to *select* (design) the sampling rates across the network subject to resource constraints, in order to collect the maximum amount of information on the underlying source-destination flows.

We address several conceptual and practical aspects of the use of such data for flow volume estimation in this work. The insights and results presented are often of general statistical interest in addition to their application in computer networks context.

1.1 Literature Review

We provide a broad review of the literature in this section. Many more references are provided in the following chapters when relevant.

An area of long standing and active interest is modeling of statistical properties of data collected on a single link. Such data has been shown to have interesting structure like long range dependence and heavy tailed distribution [48]. A relatively new area of interest has been estimation of network wide traffic volumes. This has applications for capacity planning and forecasting, routing protocol configuration, provisioning and fault-diagnosis [42].

The term network tomography was introduced in [47] for the problem of estimating source-destination flow volumes from aggregate link measurements. The flow volumes were modeled as Poisson random variables, the difficulties of estimation based on maximum likelihood demonstrated and as an alternative a low complexity method of moments estimator was proposed. Estimability (identifiability) of flow volume distribution was proved using the parametric form of the density of a Poisson random variable. In [6] flow volumes were modeled as being normally distributed with flow variances proportional to their means. The proportionality assumption leads to identifiability of the mean parameters through identifiability of variances. An estimator based on the EM algorithm was proposed. There is some evidence that this method may not estimate accurately enough the distribution of flow volumes in large high speed computer networks [35]. In [31] a computationally efficient pseudo-likelihood method for network tomography was proposed. Recently, a sufficient condition for identifiability of the entire distribution up to mean of flow volumes was established in [7]. Further, an estimator based on the characteristic function of

the aggregate data was proposed. An overview of tomography techniques can be found in [28].

Another class of models imposes other types of constraints for obtaining identifiable (estimable) solutions. For example, gravity models [51] assume that the source and destination of any given packet in the network are independent of each other. Again there is evidence that this assumption is not strictly valid in backbone networks [28]. However, the assumption introduces enough constraints to regularize the problem for a unique solution. A Kalman filter based approach suggested in [41] provides best linear estimates of flow volumes assuming a specific temporal dependence structure with known parameters. The ideas developed in some of the above papers have been employed in [42] and [30] to develop practical traffic volume estimators for continuous monitoring of real networks. A graph-wavelets based approach was developed in [39].

The use of sampled data in networking has become an active area of research. There has been a fair amount of work on ways to augment SNMP data with small amount of sampled data when necessary [30]. The focus of some of the current research is on simultaneously controlling volume and accuracy of such data [12]. For example, [9] discusses some of the considerations regarding sampling error and measurement overhead for some simple sampling schemes. In [15] this issue is investigated for different sampling schemes including threshold sampling, uniform flow sampling, uniform packet sampling and sample and hold. In addition they provide algorithms for dynamic control of sample volume. Another interesting area is the optimal combination of sampled data from across the network [15, 13]. Others [49, 50] study estimation of individual flow distributions through non-parametric techniques based on sampled data. In [53], the problem of combining (possibly dirty) SNMP

and sampled data is investigated.

1.2 Contributions and Organization

In Chapter 2 we study the problem of identifiability of joint distribution of flow volumes in a computer network from (lower dimensional) aggregate measurements collected on its edges. Conceptually, this is a canonical example of a linear statistical inverse problem. In a departure from previous approaches we investigate situations where flow volumes have dependence. We introduce a number of models that capture spatial, temporal and inter-modal (i.e. between packets and bytes) dependence between flow-volumes. These models are fairly general but specific instances that incorporate structural features of network traffic are also investigated. We provide sufficient, sometimes necessary, conditions for the identifiability of the flow volumes distribution (up to mean) under these models. Further, we investigate conditions on network routing that are sufficient for identifiability of flow volumes distribution.

In Chapter 3 we use the results and models developed in Chapter 2 to perform computer network tomography using joint modeling for packet and byte volumes. As usual the goal is to estimate characteristics of source-destination flows based on aggregate link measurements. Specifically, we use two generative models for the relation between packet and byte volumes, establish identifiability of their parameters using results from Chapter 2 and discuss different estimating procedures. The proposed estimators of the flow characteristics are evaluated using both simulated and emulated data. Further, the proposed models allow us to estimate parameters of the packet size distribution, thus providing additional insights into the composition of network traffic.

In Chapter 4 we examine the problem of optimal design in the context of filtering

multiple random walks. Specifically, we define the steady state E-optimal design criterion and show that the underlying optimization problem leads to a second order cone program. The developed methodology is applied to tracking network flow volumes using sampled data, where the design variable corresponds to controlling the sampling rate. The optimal design is numerically compared to a myopic and a naive strategy. Next, we extend the myopic strategy to state space models and numerically investigate several instances of interest for flow volume tracking. Finally, we pose the general problem of steady state optimal design for state space models.

CHAPTER II

Identifiability Results for Network Tomography

2.1   Introduction

Consider a network described by a (directed) graph $G = (V, E)$ with vertex (node) set V and edge (link) set E. Each edge $e \in E$ is an ordered pair of vertices $e = (n_1, n_2) \in E$ that connects vertex n_1 to n_2, $n_1, n_2 \in V$. Flows f_j, $j = 1, \cdots, J$, correspond to ordered pair of vertices and a volume measurement variable X_j is associated with each flow j, with $J \leq |V|^2$. Each flow may traverse several paths. A path P of length L_P is a sequence of nodes connected by edges, i.e. for $P = (n_1, \cdots, n_{L_P+1})$, $(n_i, n_{i+1}) \in E$, for $i = 1, \cdots, L_P$. We say $e_i = (n_i, n_{i+1}) \in P, i = 1, \cdots, L_P$, and n_1 and n_{L_P+1} are the origin and destination vertices of the path P. Let $\mathcal{P}(j)$ denote the set of paths traversed by flow j and $w_j(P)$ the proportion of flow j carried on path P. Note that all paths in $\mathcal{P}(j)$ have the same origin-destination node pair. Hence

$$\mathcal{P}(j) = \{P : w_j(P) > 0\},$$

$$\sum_{P \in \mathcal{P}(j)} w_j(P) = 1.$$

The set of functions $\{\mathcal{P}(j), w_j(P)\}$ determine the *routing policy* of the network.

Observations are made on edges which are a linear combination of the volume measurement variables corresponding to the flows passing through respective links.

The traffic volume on edge e is given by

$$Y_e = \sum_j \sum_{\substack{P \in \mathcal{P}(j) \\ e \in P}} w_j(P) X_j.$$

This can be written in vector notation as:

$$Y = AX,$$

where Y is a $L \times 1$ vector of observations on L edges, X is a $J \times 1$ vector of measurement variables associated with J flows and A is a $L \times J$ *routing* matrix where $[A]_{ij}$ indicates the fraction of the jth flow that traverses the ith link. In certain cases, it will be assumed that A is a binary matrix corresponding to each origin-destination flow traversing through exactly one path; i.e. $w_j(P) = 1$ for a single $P \in \mathcal{P}(j)$. The matrix A is typically not full rank as there are many more flows $(O(n^2)$, where n is the number of nodes in the graph) than links $(O(n)$, since the corresponding graphs are sparse). Our objective is to state assumptions and derive conditions on the routing matrix A under which certain distributional parameters of X are uniquely determined by the distribution of Y which is observed.

For example, consider the network in Figure 2.1 that has 6 nodes and 5 bi-directional links. Let Y_e be the total number of bytes that traverse link e in a time interval. Further, let $X_{(n_1,n_2)}$ be the number of bytes in the flow from node n_1 to node n_2 during the same time interval. Then each Y_e is a sum of $X_{(\cdot,\cdot)}$s corresponding to the flows passing through link e. For example, for $e_1 = (3,4)$ and $e_2 = (4,3)$ we have:

$$Y_{e_1} = X_{(1,5)} + X_{(1,6)} + X_{(2,5)} + X_{(2,6)} + X_{(3,4)}$$

and

$$Y_{e_2} = X_{(5,1)} + X_{(6,1)} + X_{(5,2)} + X_{(6,2)} + X_{(4,3)}.$$

Thus, each Y_e is a linear combination of the $X_{(\cdot,\cdot)}$s. Here the number of links $L = 10$ and the number of flows $J = 30$.
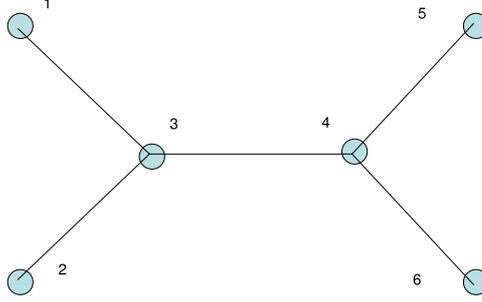
Figure 2.1: Example Topology

Now consider the setup in Figure 2.2, where the network is comprised of 3 nodes and two links. Observations on links 1 and 2 are respectively given by

$$Y_1 = X_1 + X_2,$$

$$Y_2 = X_2 + X_3.$$

As a preview of the basic idea on identifiability, note that if the flow volumes X_i are independent random variables, then their variances are "identifiable" from the joint distribution of observed edge volumes Y_1 and Y_2 as follows:

$$v_y \equiv \begin{pmatrix} \mathrm{Var}(Y_1) \\ \mathrm{Var}(Y_2) \\ \mathrm{Cov}(Y_1, Y_2) \end{pmatrix} = \begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 1 \\ 0 & 1 & 0 \end{pmatrix} \begin{pmatrix} \mathrm{Var}(X_1) \\ \mathrm{Var}(X_2) \\ \mathrm{Var}(X_3) \end{pmatrix} \equiv B v_x.$$
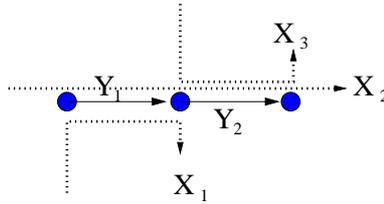
Figure 2.2: Aggregate Volume Measurements

Thus, v_y that contains the variances and the covariance of (Y_1, Y_2), uniquely determines v_x that contains the variances of X_1, X_2 and X_3, since B is a matrix of full rank. For the purpose of this chapter, a matrix C will be called full rank if $Cx = 0$ for a vector x, implies $x = 0$. Now, the matrix B is clearly a function of the routing matrix A given by

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

It can therefore be seen that "identifiability" of variances of the X_i's is related to a matrix function of A being full rank when the X_i's are uncorrelated.

More generally, let $Y(t)$ denote the vector of observations on the links during measurement interval t. These observations may be byte count or packet count as obtained from SNMP data. Further, let $X(t)$ be the (unobserved) vector of flow measurements (packet count or byte count) in the same measurement interval. We will view $X(t)$ (and hence $Y(t)$) as random vectors satisfying some stochastic model. Thus, we can posit the following model:

$$(2.1) \qquad\qquad Y(t) = AX(t), \ t = 1, \cdots.$$

In this formulation the routing matrix A does not change over time. In some cases the dependence on t may be dropped for the sake of notational convenience.

As mentioned earlier, the matrix A is typically not full rank. Thus, (2.1) cannot be solved for $X(t)$. However, under certain distributional assumptions on $X(t)$, the observations $Y(t)$ are sufficient to estimate parameters of the distribution of $X(t)$. The distribution of $X(t)$ can be modeled at different levels of complexity from independent and identically (i.i.d.) Gaussian to long range dependent with cycles induced due to diurnal or weekly patterns. The true structure of network data is quite complex and one needs to balance the need for faithful representation with

analytic tractability and computational feasibility.

In this chapter, we present models for the distribution of X and the routing policy or network structure that result in identifiability of the distribution of X (up to uncertainty in the mean). These conditions are quite often satisfied in computer networks. Notice that in general, means (i.e. $E(X)$) are not identifiable, since adding a constant vector c from the null space of the routing matrix A to X, leaves $Y (= AX)$ unchanged. Let $\mathcal{L}(X)$ denote the distribution of X and $\mathcal{M}$ be a set of possible distributions, i.e. $\mathcal{L}(X) \in \mathcal{M}$. Further, let $\theta(\mathcal{L}(X))$ be a well-defined parameter of the distribution of X. Then, identifiability is formally defined as follows.

Definition II.1. The distribution of a random vector $X \in \mathbb{R}^J$ is identifiable up to mean under model $\mathcal{M}$, from observations of the form $Y = AX$, if for $Y_1 = AX_1$ and $Y_2 = AX_2$, $\mathcal{L}(X_1), \mathcal{L}(X_2) \in \mathcal{M}$, $Y_1 \stackrel{d}{=} Y_2$ (i.e. $\mathcal{L}(Y_1) = \mathcal{L}(Y_2)$) implies that $X_1 \stackrel{d}{=} X_2 + c$ (i.e. $\mathcal{L}(X_1) = \mathcal{L}(X_2 + c)$) for some constant $c \in \mathbb{R}^J$. Similarly, a parameter, $\theta(\mathcal{L}(X))$ is said to be identifiable under model $\mathcal{M}$ if $Y_1 \stackrel{d}{=} Y_2$ (i.e. $\mathcal{L}(Y_1) = \mathcal{L}(Y_2)$) implies that $\theta(\mathcal{L}(X_1)) = \theta(\mathcal{L}(X_2))$.

For the case of independent flow volumes three kinds of identifiability results are known (see Section 2.8.1). These are conditions on routing matrix under which flow volume variances are identifiable, conditions on routing matrix under which entire flow volume distributions are identifiable up to mean and conditions on *routing policy* or network structure that imply that the routing matrix satisfies the required properties for identifiability. In the following sections, we prove similar results when flow volumes are not independent. The techniques are naturally more involved and the independence case can be elegantly recovered as a special case. These results seek to address the question of "how complex can the dependence structure of a linear inverse problem be and still be identifiable". In Section 2.2 we do an empirical inves-

tigation into the nature of dependence observed in computer network flow volumes and propose the independent connections model, which we use as an illustration of subsequent more general models. In sections 2.4, 2.5 and 2.6 we prove results about sufficient (some times necessary) conditions on the routing matrix for identifiability. In Section 2.7 we show that several reasonable instances of routing policy or network structure are sufficient for routing matrices to satisfy the required property for identifiability results. We end with a discussion and possible future directions in Section 2.8.

2.2 Some Empirical Observations and Independent Connections Model

We illustrate some important features of traffic volume data using a publicly available data-set (obtained from `http://www-dirt.cs.unc.edu/ts/`). The data is essentially a 4-variate time series where the 4 variables are packet and byte volumes of forward direction and reverse direction traffic on a link. Each observation represents the traffic traversing that link in a 10 second interval. We limit ourselves to the first 700 observations. The 4 time series are plotted in Figure 2.3.

The temporal dependence is visible in the time series and can be more clearly seen through their auto-correlation functions in Figure 2.4. The auto-correlation function of a time-series $x(t)$, at a given lag l, is the observed correlation between $x(t)$ and $x(t - l)$ over all values of t. For each of the four time-series the auto-correlation functions are significantly greater than zero and decay with increasing lag. The simplest possible model for such time-series is an auto-regressive model. We fit auto-regressive models to each of the 4 time series. The orders of the models chosen through AIC were 4 for forward byte volume, 8 for reverse byte volume and 5 for both forward and reverse packet volumes. The residuals from these models have
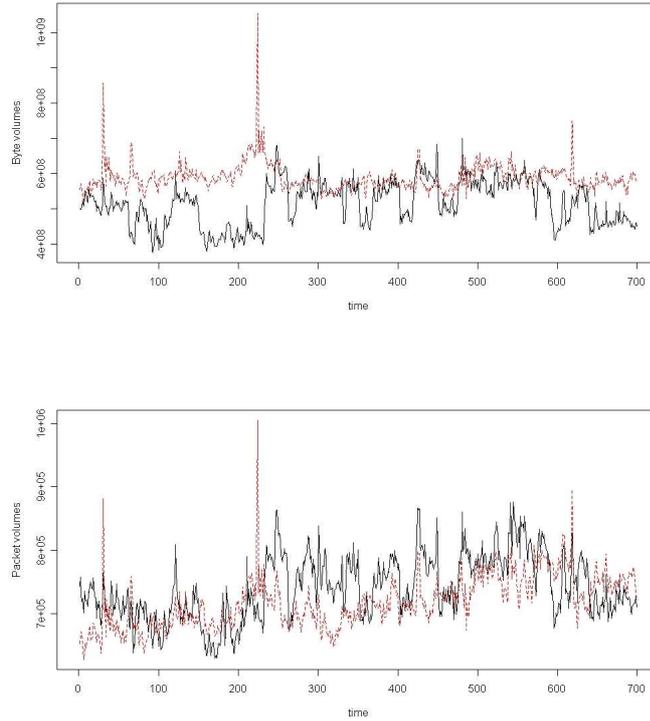
Figure 2.3: Byte (top) and Packet (bottom) time-series

the following correlation matrix (FB = forward byte, FP = forward packet, RB = reverse byte, RB = reverse packet)

	FB	FP	RB	RP
FB	1.00	0.83	0.04	0.22
FP	0.83	1.00	0.24	0.44
RB	0.04	0.24	1.00	0.89
RP	0.22	0.44	0.89	1.00

Thus we clearly see the strong dependence between packet and byte volumes and between forward and reverse flows. Finally we compare the quantiles of the residuals from the AR models to that of a standard normal distribution. Figure 2.5 clearly shows that the quantiles of the observed error distribution are more extreme than that of a normal distribution indicating heavier tails.
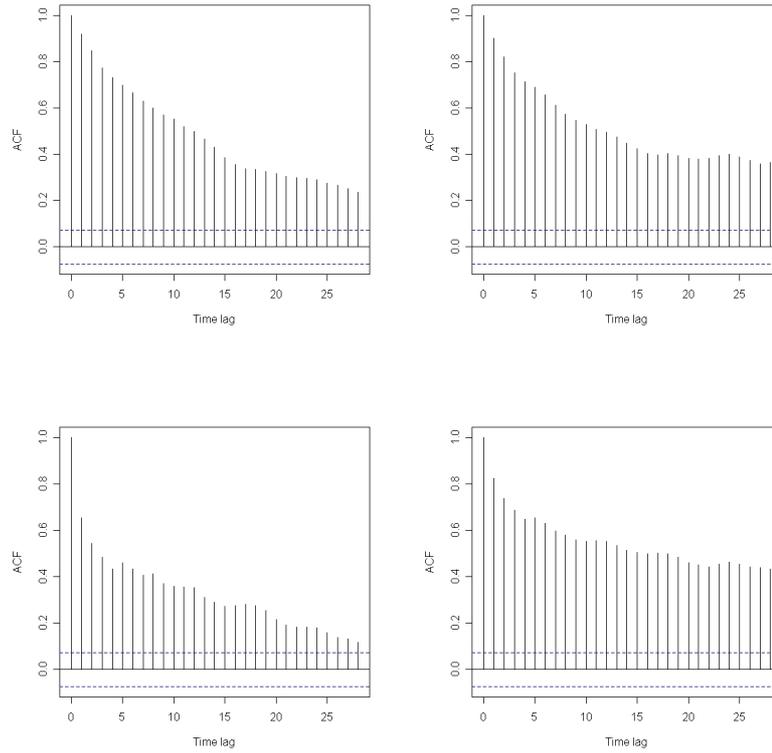
Figure 2.4: Auto-correlation functions: forward byte (top left), forward packet (top right), reverse byte (bottom left) and reverse packet (bottom right)

Next we look at spatial correlation between network-wide flow volumes. The data-set used for this analysis [46] gives byte volumes of all flows in the network for each 15 minute interval over a period of 4 months. As seen in the previous example, we expect correlations between forward and reverse byte volumes to be weaker than for packet volumes. The following analysis shows that these correlations are still substantially stronger than other spatial correlations.

We restrict ourselves to the first 1500 observations and 76 flows with no missing value and comprising the top quarter of flows in terms of average traffic. The time-series for each flow was spline-smoothed to estimate non-stationarities like the well-known diurnal patterns [16, 27]. The residuals from the above step were used to fit an auto-regressive model, with AIC-based order selection, to account for tem-
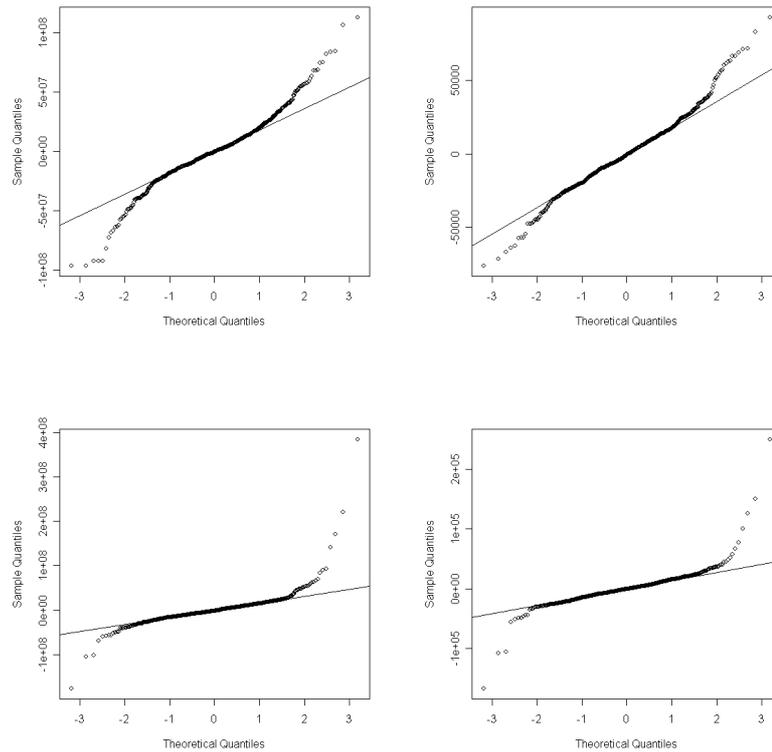
Figure 2.5: QQ plots: forward byte (top left), forward packet (top right), reverse byte (bottom left) and reverse packet (bottom right)

poral dependencies. The pair-wise correlations between residual time-series from the above analysis represent the spatial correlations. These pair-wise correlations can be divided into two sets, the forward-reverse kind and all the others. Figure 2.6 plots the observed densities of these two sets of correlations. Clearly the forward-reverse correlations are stronger than the rest. While, there is a bi modality in both distributions, it is significantly more pronounced for the forward-reverse correlations. Ideally one would like to model all significant spatial correlations. However, in order to have a systematic and parsimonious model we focus on the forward reverse correlations. As mentioned earlier, we believe that such dependence would be stronger and of greater practical interest for packet volumes, as opposed to byte volumes.

Based on these observations and previous studies [16], we outline a useful model
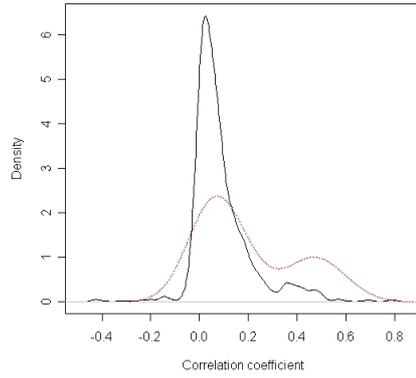
Figure 2.6: Densities of observed correlations: Forward-reverse (dashed) and the rest (solid)
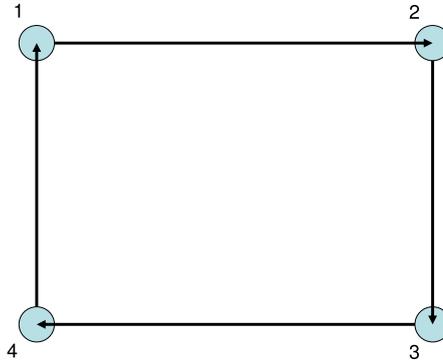
Figure 2.7: Example Topology

for computer networks, which we will refer to as Independent Connections Model. The most significant spatial correlation is the one between the packet counts of a flow and its reverse flow, i.e. for nodes n_1, n_2, the volume of flow from n_1 to n_2 and the volume of flow from n_2 to n_1. Partition the set of flows into two groups F(forward) and R(reverse). Thus, for a particular type of measurement -say packet counts- we have:

$$(2.2) \qquad Y^{(p)} = A_F X_F^{(p)} + A_R X_R^{(p)}.$$

If the number of edges is L and the number of flows is J, then both A_F and A_R are $L \times J/2$ matrices. For example, consider the network in Figure 2.7 comprised of 4

nodes and 4 links, and where all flows follow a clock-wise path. Let $e_1 = (1, 2), e_2 = (2, 3), e_3 = (3, 4)$ and $e_4 = (4, 1)$. The above equation becomes:

$$\begin{pmatrix} Y_{e_1} \\ Y_{e_2} \\ Y_{e_3} \\ Y_{e_4} \end{pmatrix} = \begin{pmatrix} 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 1 & 1 & 1 & 1 & 0 \\ 0 & 0 & 1 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} X_{(1,2)} \\ X_{(1,3)} \\ X_{(1,4)} \\ X_{(2,3)} \\ X_{(2,4)} \\ X_{(3,4)} \end{pmatrix} + \begin{pmatrix} 0 & 0 & 0 & 1 & 1 & 1 \\ 1 & 0 & 0 & 0 & 0 & 1 \\ 1 & 1 & 0 & 1 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 & 1 \end{pmatrix} \begin{pmatrix} X_{(2,1)} \\ X_{(3,1)} \\ X_{(4,1)} \\ X_{(3,2)} \\ X_{(4,2)} \\ X_{(4,3)} \end{pmatrix}.$$

Equation (2.2) can be rewritten as:

$$Y^{(p)} = AX^{(p)},$$

where $A = (A_F, A_R)$ and $X^{(p)} = (X_F^{(p)'}, X_R^{(p)'})'$. In real computer networks, a large part of the traffic is connection oriented. For example, traffic flows transported using the TCP protocol [37], or connections involving Internet (Voice over IP) telephony, lead to packets being exchanged between the two endpoints. In the former case, due to the built-in acknowledgment mechanism of packets in the TCP protocol, while in the latter case due to the bidirectional nature of the connection. Therefore, volumes of flow from node n_1 to node n_2 and vice-versa, are correlated [16]. One of these flows is labeled as a forward flow and the other as a reverse flow and form a flow pair. It is reasonable to assume that flow pairs are independent with possible dependence between forward and reverse flows of a flow pair. In particular, if second moments exist, then the covariance matrix of $X^{(p)}$ is of the form

(2.3)
$$\Sigma_X = \begin{pmatrix} \text{Diag}(\delta_{FF}) & \text{Diag}(\delta_{FR}) \\ \text{Diag}(\delta_{FR}) & \text{Diag}(\delta_{RR}) \end{pmatrix},$$

where each of $\delta_{FF}, \delta_{FR}, \delta_{RR}$ is a vector of length $J/2$ and component wise they correspond to the variances of $X_F^{(p)}$, covariances of $X_F^{(p)}$ and $X_R^{(p)}$ and variances of $X_R^{(p)}$, respectively. Thus, Σ_X is a matrix of dimension $J \times J$.

If $X^{(p)}$ is further assumed to be multivariate Normally distributed the above model corresponds to the following latent variable model:

$$
\begin{aligned}
X_{Fj}^{(p)} &= c_{1j}Z_{1j}, \\
X_{Rj}^{(p)} &= c_{2j}Z_{1j} + c_{3j}Z_{2j}.
\end{aligned}
$$

with Z_{ij} independent Normal with (possibly) different means and unit variances for all i,j. Notice that two independent latent variables, Z_{1j} and Z_{2j}, are associated with flow pair j. The reverse flow in flow pair j is the sum of a component proportional to the forward flow of j and a unique component. This can also be written as

$$
\begin{pmatrix} X_F^{(p)} \\ X_R^{(p)} \end{pmatrix} = \begin{pmatrix} \mathrm{Diag}(c_1) & 0 \\ \mathrm{Diag}(c_2) & \mathrm{Diag}(c_3) \end{pmatrix} \begin{pmatrix} Z_1 \\ Z_2 \end{pmatrix} \equiv CZ.
$$

The above model corresponds to having exactly one type of measurement. Different types of measurements on each flow can be observed in practice as follows:

1. *Bi-modal measurements on each flow.* As mentioned earlier, there are two measurements of interest associated with each flow in computer networks; namely, packet counts and byte counts. We will denote the type of measurement by the superscript, (p) and (b) for packets and bytes, respectively. Since the byte count is the sum of bytes in each packet, there is a strong dependence between these two types of measurements, as seen in the empirical analysis at the beginning of section. Now consider another model, with dependence within flows

and between packet counts and byte counts of the same flow:

$$
\begin{pmatrix} Y^{(p)} \\ Y^{(b)} \end{pmatrix} = \begin{pmatrix} A_F & A_R & 0 & 0 \\ 0 & 0 & A_R & A_F \end{pmatrix} \begin{pmatrix} X_F^{(p)} \\ X_R^{(p)} \\ X_F^{(b)} \\ X_R^{(b)} \end{pmatrix} \equiv \overline{A} X.
$$

Again we assume independence between flow pairs, but not within the forward and reverse flow and packet and byte measurements of the same flow pair. Specifically, if second moments exist, then the covariance of X takes the form:

$$
\Sigma_X = \begin{pmatrix} \mathrm{Diag}(\delta_{Fp,Fp}) & \mathrm{Diag}(\delta_{Fp,Rp}) & \mathrm{Diag}(\delta_{Fp,Fb}) & \mathrm{Diag}(\delta_{Fp,Rb}) \\ \mathrm{Diag}(\delta_{Fp,Rp}) & \mathrm{Diag}(\delta_{Rp,Rb}) & \mathrm{Diag}(\delta_{Rp,Fb}) & \mathrm{Diag}(\delta_{Rp,Rb}) \\ \mathrm{Diag}(\delta_{Fp,Fb}) & \mathrm{Diag}(\delta_{Rp,Fb}) & \mathrm{Diag}(\delta_{Fb,Fb}) & \mathrm{Diag}(\delta_{Fb,Rb}) \\ \mathrm{Diag}(\delta_{Fp,Rb}) & \mathrm{Diag}(\delta_{Rp,Rb}) & \mathrm{Diag}(\delta_{Fb,Rb}) & \mathrm{Diag}(\delta_{Rb,Rb}) \end{pmatrix}.
$$

In the above, $\delta_{Aa,Bb}$ denotes the covariance of $X_A^{(a)}$ and $X_B^{(b)}$ for $a, b \in \{p, b\}$ and $A, B \in \{F, B\}$, each of them a vector of length $J/2$. Thus, Σ_X is a matrix of dimension $2J \times 2J$.

2. *Temporal dependence.* As the empirical analysis shows, network data when viewed over moderate time-scales exhibit not just spatial dependence of the nature captured by previous models but also temporal dependence. This dependence can be modeled as follows:

$$
Y^{(p)}(t) = AX^{(p)}(t),
$$

$$
Y^{(b)}(t) = AX^{(b)}(t),
$$

$$
X^{(p)}(t) = \Phi_{p,1} X^{(p)}(t-1) + \cdots + \Phi_{p,m} X^{(p)}(t-m) + \epsilon^{(p)}(t),
$$

$$
X^{(b)}(t) = \Phi_{b,1} X^{(b)}(t-1) + \cdots + \Phi_{b,m} X^{(b)}(t-m) + \epsilon^{(b)}(t),
$$

where the various $\Phi_{.,.}$ matrices contain the lag coefficients and $\epsilon^{(p)}(t)$, $t = 1, \cdots$, are i.i.d. mean 0 random vectors and so are $\epsilon^{(b)}(t)$, $t = 1, \cdots$. For the purpose of illustration, assume $\Phi_{p,1} = \Phi_p$, $\Phi_{b,1} = \Phi_b$ and $\Phi_{p,k} = \Phi_{b,k} = 0$ for $k > 1$. Assuming stationarity of the above auto-regressive models, it is easy to verify the following:

$$(2.4) \qquad\qquad \Sigma_{X,pp} = \Phi_p \Sigma_{X,pp} \Phi_p' + \Sigma_{pp},$$

$$(2.5) \qquad\qquad \Sigma_{X,pb} = \Phi_p \Sigma_{X,pp} \Phi_b' + \Sigma_{pb},$$

$$(2.6) \qquad\qquad \Sigma_{X,bp} = \Phi_b \Sigma_{X,bp} \Phi_p' + \Sigma_{bp},$$

$$(2.7) \qquad\qquad \Sigma_{X,bb} = \Phi_b \Sigma_{X,bb} \Phi_b' + \Sigma_{bb},$$

$$\text{Cov}(X^p(t), X^{(p)}(t - l)) = \Phi_p^l \Sigma_{X,pp} \equiv \Sigma_{X,pp}^l,$$

$$\text{Cov}(X^{(p)}(t), X^{(b)}(t - l)) = \Phi_p^l \Sigma_{X,pb} \equiv \Sigma_{X,pb}^l,$$

$$\text{Cov}(X^{(b)}(t), X^{(p)}(t - l)) = \Phi_b^l \Sigma_{X,bp} \equiv \Sigma_{X,bp}^l,$$

$$\text{Cov}(X^{(b)}(t), X^{(b)}(t - l)) = \Phi_b^l \Sigma_{X,bb} \equiv \Sigma_{X,bb}^l,$$

where Σ_{pp}, Σ_{pb}, Σ_{bp} and Σ_{bb} are covariances and cross-covariances of the random noise variables $\epsilon^{(p)}(t)$ and $\epsilon^{(b)}(t)$.

Now assume that each of Σ_{pp}, Σ_{pb}, Σ_{bp} and Σ_{bb} are block diagonal matrices of the form (2.3), that captures the spatial correlations between the flows. Further assume that Φ_p and Φ_b are diagonal with each entry less than 1. Thus (2.4- 2.7) imply that Σ_{Xpp}, Σ_{Xpb}, Σ_{Xbp} and Σ_{Xbb} have the same block diagonal form given in (2.3). This in turn implies that the covariances $\Sigma_{X,pp}^l$, $\Sigma_{X,pb}^l$, $\Sigma_{X,bp}^l$ and $\Sigma_{X,bb}^l$ also have the form (2.3).

2.3   Notation and Preliminary Results

Definition II.2. For a $L \times J$ matrix $A = [a_1, \cdots, a_J]$ and $M \times J$ matrix $B = [b_1, \cdots, b_J]$ the $LM \times J$ Khatri Rao product $A \odot B$ is defined as $[a_1 \otimes b_1, \cdots, a_J \otimes b_J]$ where $\otimes$ denotes the Kronecker Product.

Note that rows in $A \odot B$ are element-wise product of a row in A and a row in B. Specifically, row $g(i, j) \equiv (i - 1)L + j$ in $A \odot B$ is the element-wise product of the ith row in A and the $j$the row in B.

Lemma II.3. *[26] If $M \times J$ matrix A has rank J and B has no null columns then $A \odot B$ has rank J.*

Definition II.4. The characteristic function of a J dimensional random vector X is defined as $\psi(t) = E[e^{\iota t' X}]$, for $t \in \mathbb{R}^J$.

A characteristic function is called analytic if it can be represented by a convergent power series in a vicinity of 0 [10]. For identifiability results we will usually assume that relevant characteristic functions are either analytic or have no roots in $\mathbb{R}^J$ (with appropriate J). This allows us to work with log of the characteristic function in the vicinity of 0 or over entire $\mathbb{R}^J$ respectively. In the former case, due to analyticity, the value around 0 uniquely determines the corresponding distribution [10, 26]. A lot of well known distributions, e.g. Gaussian, Exponential, Gamma, have analytic characteristic functions. Some heavy-tailed distribution have characteristic function with no real roots, e.g. α-stable distributions [24, 7].

Definition II.5. For a $L \times J$ matrix $A = [a_1, \cdots, a_J]$, $vec(A) \equiv (a_1', \cdots, a_J')'$.

We will use the same notation, *vec*, for matrices of different dimensions since the meaning is clear from the context.

Lemma II.6. *If $V(\theta) = \theta_1 V_1 + \cdots + \theta_r V_r$ for $\theta \in S$, where S is an open subset of $\mathbb{R}^r$ and V_i is $J \times J$ real matrix, for $i = 1, \cdots, r$, then θ is identifiable from $V(\theta)$ (i.e $\theta_1 \neq \theta_2$ implies $V(\theta_1) \neq V(\theta_2)$) iff $\tilde{V} = [vec(V_1), \cdots, vec(V_r)]$ has rank r.*

Proof: Note that if $V(\theta_1) = V(\theta_2)$ then $vec(V(\theta_1) - V(\theta_2)) = vec(V(\theta_1 - \theta_2)) = 0$. Substituting $\lambda = \theta_1 - \theta_2$

$$0 = vec(\sum_{i=1}^{r} \lambda_i V_i) = \sum_{i=1}^{r} \lambda_i vec(V_i)$$

Thus if $\tilde{V}$ has rank r then $\lambda = 0$. To prove necessity note that if $\tilde{V}\lambda = 0$ for $\lambda \neq 0$, there exists $\theta \in S$ and $\epsilon > 0$ such that $\theta + \epsilon\lambda \in S$. However, then $V(\theta) = V(\theta + \epsilon\lambda)$.

$\square$

Corollary II.7. *If $U = [u_1, \cdots, u_r]$ for $u_i \in \mathbb{R}^J$, $i = 1, \cdots, r$, and $V(\theta) = \theta_1 u_1 u_1' + \cdots + \theta_r u_r u_r'$, for $\theta \in \mathbb{R}^r$ such that $V(\theta) \geq 0$ then θ is identifiable from $V(\theta)$ (i.e $\theta_1 \neq \theta_2$ implies $V(\theta_1) \neq V(\theta_2)$) iff $U \odot U$ has rank r.*

Proof: Follows from the above lemma by noting $vec(u_i u_i') = u_i \otimes u_i$ and that for $\theta \in \mathbb{R}_+^r$, $V(\theta) \geq 0$.

$\square$

Lemma II.8. *(Lemma 1.5.1 [26]) Consider the functional equation*

$$(2.8) \qquad \phi_1(u + b_1 v) + \cdots + \phi_r(u + b_r v) = A(u) + B(v)$$

for $u, v \in \mathbb{R}$, $|u| < \delta$ and $|v| < \delta$. Also assume that the numbers b_j are all distinct (without loss of generality) and non-zero and that complex valued functions $A, B, \phi_1, \cdots, \phi_r$ are continuous. Then, in some neighborhood of the origin, the functions $A, B, \phi_1, \cdots, \phi_r$ are all polynomials of degree $\leq r$.

Lemma II.9. *(Lemma 1.5.2 [26]) Consider the functional equation*

$$(2.9) \qquad \phi_1(\alpha_1' t) + \cdots + \phi_J(\alpha_J' t) = \xi_1(t_1) + \cdots + \xi_L(t_L)$$

defined for $|t_i| < \delta$, $i = 1, \cdots, L$, where $t \in \mathbb{R}^L$ is a column vector of variables $t_1, \cdots, t_L$ and $\alpha_1, \cdots, \alpha_J$ are the column vectors of a given $L \times J$ matrix A. Also assume that functions $\phi_1, \cdots, \phi_J$ and $\xi_1, \cdots, \xi_L$ are continuous. If no column of A is proportional to another column of A or to a column of $I_{L \times L}$, then $\phi_1, \cdots, \phi_J$ are necessarily polynomials of degree $\leq J$.

Lemma II.10. *Consider the functional equation*

$$\phi_1(\alpha_1't) + \cdots + \phi_J(\alpha_J't) = 0 \tag{2.10}$$

defined for $|t_i| < \delta$, $i = 1, \cdots, L$, where $t \in \mathbb{R}^L$ is a column vector of variables $t_1, \cdots, t_L$ and $\alpha_1, \cdots, \alpha_J$ are the column vectors of a given $L \times J$ matrix A, with $L \leq J$. Also assume that functions $\phi_1, \cdots, \phi_J$ are continuous. If $A \odot A$ has rank J, then $\phi_1, \cdots, \phi_J$ are all linear functions.

Proof: $A \odot A$ has rank J implies that no two columns of A are proportional. If $J > L$ then there is at least one column in A not proportional to a column in $I_{L \times L}$. Without loss of generality, let $\alpha_1, \cdots, \alpha_K$ be columns not proportional to any column in $I_{L \times L}$. Then

$$\phi_1(\alpha_1't) + \cdots + \phi_K(\alpha_K't) = \xi_1(t_1) + \cdots + \xi_L(t_L)$$

where if there exists i such that $\alpha_i = c_i e_k$, for e_k the kth column of $I_{L \times L}$, then $\xi_k(t) = -\phi_k(c_i t)$ and $\xi_k(t) = 0$ otherwise. Now using Lemma II.9, $\phi_1, \cdots, \phi_K, \xi_1, \cdots, \xi_L$ are all polynomial of degree $K(\leq J)$ at most.

Now, let

$$\phi_i(u) = \lambda_{iJ} u^J + \cdots + \lambda_{i1} + \lambda_{i0}$$

for $i = 1, \cdots, J$. Thus

$$\phi_1(\alpha_1't) + \cdots + \phi_J(\alpha_J't) = 0$$

implies that

$$\sum_{i=1}^{J}\sum_{k=0}^{J}\lambda_{ik}(\alpha'_i t)^k = 0$$

In the above expression, coefficient of t_k^2 is $\sum_{i=1}^{J}\lambda_{i2}\alpha_{ik}^2$ and coefficient of $t_k t_{k'}$ for $k \neq k'$ is $\sum_{i=1}^{J}\lambda_{i2}2\alpha_{ik}\alpha_{ik'}$. Note that since α_{ik} is the kth element of the ith column in A, $\alpha_{ik} = A_{ki}$. Putting all such equations together (with proper normalization and making copies of certain equations) we get

$$(A \odot A)\lambda_2 = 0$$

Therefore $\lambda_2 = 0$. Similarly, coefficient of $t_{k_1}, t_{k_2}, \cdots, t_{k_p}$ is

$$c(k_1, \cdots, k_p) \sum_{i=1}^{J} \lambda_{ip}\alpha_{ik_1} \cdots \alpha_{ik_p}$$

where $c(k_1, \cdots, k_p)$ is a positive combinatorial factor depending on which of $k_1, \cdots, k_p$ are distinct. Thus

$$(A\odot)^p \lambda_p = 0$$

where $(A\odot)^p = (A\odot)^{p-1} \odot A$ and $(A\odot)^1 = A$. Now from Lemma II.3, $(A\odot)^p$ has rank J for $p > 1$. Thus $\lambda_p = 0$ for $p > 1$.

The case of $J = L$ is handled by noting that if there is a column in A not proportional to a column in $I_{L\times L}$ then the above argument is applicable. Otherwise the only way $A \odot A$ can be rank J is if columns in A are proportional to distinct columns in $I_{L\times L}$. Then using $t = t_0 e_k$ in (2.10) with e_k the kth column of $I_{L\times L}$ for $k = 1, \cdots, L$ one gets $\phi_k(t_0) = 0$ for $k = 1, \cdots, J$.

$\square$

Corollary II.11. *If $Y = AX$ and X_j are distributed independently for $j = 1, \cdots, J$ and the characteristic functions of all X_j are either analytic or have no real roots, then the distribution X is identifiable up to mean, from Y, if $A \odot A$ has rank J.*

Proof: Suppose $X_i^{(1)}$, $i = 1, \cdots, J$ distributed independently and $X_i^{(2)}$, $i = 1, \cdots, J$ distributed independently are such that

$$AX^{(1)} \stackrel{d}{=} AX^{(2)}$$

Thus

$$\log E[e^{\iota t' AX^{(1)}}] = \log E[e^{\iota t' AX^{(2)}}]$$

The above expression is meaningful for $t_i \in \mathbb{R}$, $i = 1, \cdots, J$, if the characteristic functions of all X_i have no real roots or for t_i in a vicinity of 0, $i = 1, \cdots, J$, when the characteristic functions of X_i are analytic. Define

$$\phi_i(t) = \log E[e^{\iota t X_i^{(1)}}] - \log E[e^{\iota t X_i^{(2)}}]$$

for values of t in appropriate range. Thus

(2.11) $$\phi_1(t'\alpha_1) + \cdots + \phi_J(t'\alpha_J) = 0$$

Using Lemma II.10, $\phi_1, \cdots, \phi_J$ are linear functions. Hence the distribution of X is identifiable from distribution of Y up to a mean ambiguity for either of the two conditions on characteristic functions of X_i, $i = 1, \cdots, J$.

$\square$

Lemma II.12. *(Theorem 10.3.5 [26]) Let X_1 and X_2 be J dimensional random vectors whose elements are independent* non-normal *random variables. Further let A_1 and A_2 be $L \times J$ real matrices. If $A_1 X_1 \stackrel{d}{=} A_2 X_2 + c$, for some constant $c \in \mathbb{R}^L$, then every column in A_1 is proportional to some column in A_2 and vice-versa.*

2.4   Spatial Dependence

Intuitively, it seems reasonable that identifiability in the presence of dependence would rely on some notion of sparsity in the dependence structure. In this section we

introduce three distinct but related models such that it is possible to derive conditions under which they are respectively identifiable. The first model is the following:

Covariance Model: Assume $Cov(X) = V(\theta)$ where

$$(2.12) \qquad\qquad V(\theta) = \theta_1 u_1 u_1' + \cdots + \theta_r u_r u_r'$$

with $U = [u_1, \cdots, u_r]$ assumed known.

Note that a completely arbitrary covariance matrix can be modeled as above. Specifically, an arbitrary $J \times J$ covariance matrix can be modeled by using

$$U = U_J \equiv [I_{J \times J}, P_J]$$

where P_J is a *binary* $J \times J(J-1)/2$ matrix with distinct columns, each of which has exactly 2 non-zero entries. It is easy to see that for $i \neq j$, $[V(\theta)]_{ij} = \theta_{J+k(i,j)}$, where $k(i,j)$ is the index of column in P_J with non-zero ith and j element. Further, $[V(\theta)]_{ii} = \theta_i + \sum_{j \neq i} \theta_{J+k(i,j)}$.

In practice one would not use U_J when observations, $Y = AX$, are lower dimensional than X, which is the case of interest. One possibility in such situations is to assume that $Cov(X)$ is block diagonal with J_1 diagonal blocks of size $J_2 \times J_2$ and all off diagonal blocks are 0. This can be easily modeled through using

$$(2.13) \qquad\qquad U = I_{J_1 \times J_1} \otimes U_{J_2}$$

Thus the covariance model (2.12) is quite flexible and, while we do not pursue that direction, one could construct an increasing family of such models to capture hierarchical notions of spatial dependence. Using the identifiability results one could ascertain which of these models are identifiable.

Modeling of covariance alone may not be sufficient in some applications. One may be interested in modeling higher order properties of flow volumes. In fact, for certain

networks flow volumes are known to have a heavy tailed distribution. In this case second-order moments of flow volumes distribution would not exist. Thus we present two separate extensions of the covariance model. Both of these are in the spirit of factor analytic models, and are related to the covariance model through the matrix U. They have "essentially the same sparseness pattern" in spatial dependence as the corresponding covariance model.

Independent Component Model: We assume $X = UZ$ where $Z_1, \cdots, Z_r$ are independent random variables with U assumed known as above. Use of arbitrary distributions for $Z_1, \cdots, Z_r$ allows us to model entire distribution and not just co-variances. To prove identifiability results we will make some assumptions on the characteristic functions of $Z_1, \cdots, Z_r$ as indicated in Section 2.3. When second moments of $Z_1, \cdots, Z_r$ exist the covariance of X is given by equation (2.12). However the coefficients $\theta_1, \cdots, \theta_r$ are restricted to be positive and equal to the variances of $Z_1, \cdots, Z_r$.

Latent Variable Model: Latent variable model corresponding to a covariance model (defined by matrix U) is given as $X = CZ$ where $Z_1, \cdots, Z_J$ are independent random variables and $C \in \mathcal{C}(U)$, where $\mathcal{C}(U)$ is a set of $J \times J$ real matrices defined as follows. If $C \in \mathcal{C}(U)$, then C is a lower triangular matrix with all diagonal entries equal to 1, such that for every vector $d \in \mathbb{R}^J_+$ there is a vector $\theta \in \mathbb{R}^r$ satisfying $CDiag(d)C' = V(\theta) \geq 0$.

The set $\mathcal{C}(U)$ has two important properties. First, if $C \in \mathcal{C}(U)$ then $CDiag(d) \in \mathcal{C}(U)$ for all $d \in \mathbb{R}^J$. In other words, $\mathcal{C}(U)$, is closed under column scaling, a fact which is used crucially in Proposition II.13. Second, when $Z_1, \cdots, Z_J$ have variances, say $d_1, \cdots, d_J$ respectively, then the covariance matrix of $X = CZ$, for $C \in \mathcal{C}(U)$, is equal to $CDiag(d)C'$, and hence equal to $V(\theta)$ for some $\theta \in \mathbb{R}^r$. Thus, we get

the correspondence between the covariance model and the latent variable model. Note that when $V(\theta)$ is positive definite then the Cholesky decomposition gives the corresponding unique coefficient matrix C. A necessary condition for $\mathcal{C}(U)$ to be non-trivial is if U has rank J. (If not then there exists $x \neq 0$ such that $x'u_i = 0$ for all $i = 1, \cdots, r$. Thus $x'V(\theta)x = 0$ for all θ. Thus $V(\theta)$ cannot be positive definite whereas $CDiag(d)C'$ is positive definite whenever all elements of d are positive.) However, this may not be sufficient for $\mathcal{C}(U)$ to be non-trivial. When U corresponds to a block diagonal covariance matrix (e.g. equation (2.13)) then $\mathcal{C}(U)$ contains every matrix obtained from the Cholesky decomposition of $V(\theta)$ (for all $V(\theta) > 0$).

We will refer to C as the *coefficient matrix* and $Z_1, \cdots, Z_J$ as the *latent variables*. We will make certain assumptions, of a nature similar to those in Section 2.3, on the characteristic function of the latent variables.

Example: Independent Connections Model Using equation (2.13), the independent connections model of equation (2.3) is obtained using

$$
U = \begin{pmatrix} I_{J/2 \times J/2} & 0 & I_{J/2 \times J/2} \\ 0 & I_{J/2 \times J/2} & I_{J/2 \times J/2} \end{pmatrix}
$$

Proposition II.13. *Given $Y = AX$, if $AU \odot AU$ has rank r then*

1. *If $Cov(X)$ exists and is equal to $V(\theta)$ given by (2.12) then θ is identifiable from $Cov(Y)$.*

2. *If X satisfies the Independent Component Model with $Z_1, \cdots, Z_r$ such that either their characteristic functions are all analytic or all have no real roots, then the distributions of $Z_1, \cdots, Z_r$ are identifiable up to mean from the distribution of Y.*

3. *If U has rank J and X satisfies the latent variable model with $Z_1, \cdots, Z_J$ all*

non-normal random variables such that either their characteristic functions are all analytic or all have no real roots, then the matrix C and distributions of $Z_1, \cdots, Z_J$ are identifiable up to mean from the distribution of Y.

Proof: Identifiability of covariance model follows from a direct application of Corollary II.7. Identifiability of independent component model follows from Corollary II.11.

For the latent variable model proceed as follows. Suppose $AC_1 Z_1 \stackrel{d}{=} AC_2 Z_2 + \mu$ such that $C_1, C_2 \in \mathcal{C}(U)$ and random vectors Z_1, Z_2 satisfy the said assumptions. Using Lemma II.12, every column in AC_1 is proportional to a column in AC_2 (and vice versa). Suppose $C_1 \neq C_2$. Now consider the following two cases:

1. If columns in AC_1 are proportional to distinct columns in AC_2 and the proportionality constants are all non-zero then

$$AC_1 = AC_2 Diag(d) P$$

   where elements of vector d are nonzero and P is a permutation matrix. Since $P'P = I$,

$$AC_1 C_1' A' = AC_2 Diag(d^2) C_2' A'$$

   interpreting the square of a vector as element-wise second power. From the definition of latent variable model there is θ_1 and θ_2 such that $V(\theta_1) = C_1 C_1'$ and $V(\theta_2) = C_2 Diag(d^2) C_2'$. Since $C_1 \neq C_2$, from the uniqueness of Cholesky decomposition $V(\theta_1) \neq V(\theta_2)$ and hence $\theta_1 \neq \theta_2$.

2. If two columns in AC_1 are proportional to each other or if there are any 0 columns in AC_1 then clearly there exist vectors d_1, d_2 with positive elements such that $d_1 \neq d_2$ but

$$AC_1 Diag(d_1) C_1' A' = AC_1 Diag(d_2) C_1' A'$$

Again from the definition of latent variable model and uniqueness of Cholesky decomposition there is $\theta_1 \neq \theta_2$ such that $V(\theta_1) = C_1 Diag(d_1) C_1'$ and $V(\theta_2) = C_1 Diag(d_2) C_1'$.

In both cases we have $\theta_1 \neq \theta_2$ such that $AV(\theta_1)A' = AV(\theta_2)A'$. However, this is not a possible from the already established identifiability of covariance model. Thus $C_1 = C_2$ and hence the coefficient matrix is identifiable. For identifiability of the distribution of latent variable vector Z up to mean we only need that $AC \odot AC$ have rank J from Corollary II.11. If $AC \odot AC$ has rank less than J, then from Corollary II.7, there exist $d_1, d_2 \in \mathbb{R}^J$ such that $d_1 \neq d_2$ and $ACDiag(d_1)C'A' = ACDiag(d_2)C'A'$. In fact it is easy to see that d_1, d_2 can be chosen to have positive elements. However, due to the uniqueness of Cholesky decomposition this implies that there exist $\theta_1 \neq \theta_2$ such that $AV(\theta_1)A' = AV(\theta_2)A'$, which is again not possible from the (already established) identifiability of the corresponding covariance model.

$\square$

Remark: In the case of Independent Connections Model $AU \odot AU$ having rank $r = 3J/2$ is equivalent to $\overline{B}_c \equiv [A_F \odot A_F, A_F \odot A_R + A_R \odot A_F, A_R \odot A_R]$ having rank $r = 3J/2$. This can be shown to be the case under reasonable conditions on routing and network structure in Section 2.7.

2.5 Spatio-Temporal Dependence

At a high level, spatio-temporal dependence can be handled using the models developed in Section 2.4 as explained in the following. Suppose, $Y_t = AX_t$, for

time-interval $t = 1, \cdots, T$. Define $\overline{Y} = (Y_1', \cdots, Y_T')'$, $\overline{X} = (X_1', \cdots, X_T')'$ and

$$\overline{A} = \begin{pmatrix} A & 0 & \cdots & 0 \\ 0 & A & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & A \end{pmatrix} = I_{T \times T} \otimes A$$

Then $\overline{Y} = \overline{A}\overline{X}$. Suppose $Cov(\overline{X})$ is modeled as $V(\theta)$ as given in equation (2.12) for some U, then from Proposition II.13, θ is identifiable from $Cov(\overline{Y})$ if and only if $\overline{A}U \odot \overline{A}U$ has rank r. However, one would like to explore conditions when the identifiability conditions can be stated in terms of objects simpler than $\overline{A}U \odot \overline{A}U$.

One such condition is as described in the following. Suppose vector θ of length r is partitioned as $\theta = (\theta_{11}', \cdots, \theta_{TT}')'$ such that θ_{ij} is a vector of length r_{ij} for $i = 1, \cdots, T$ and $j = 1, \cdots, i$ and $\sum_{i=1}^{T} \sum_{j=1}^{i} r_{ij} = r$. Further, assume $Cov(X_i, X_j) = V_{ij}(\theta_{ij})$ for $j \le i$ i.e.

$$V(\theta) = \begin{pmatrix} V_{11}(\theta_{11}) & \cdots & V_{T1}'(\theta_{T1}) \\ \vdots & \ddots & \vdots \\ V_{T1}(\theta_{T1}) & \cdots & V_{TT}(\theta_{TT}) \end{pmatrix}$$

Clearly $V_{ii}(\cdot)$ should map into symmetric matrices while there is no such restriction for $V_{ij}(\cdot)$ for $j < i$. We can rewrite $V(\theta)$ as

(2.14) $\qquad V(\theta) = \sum_{i=2}^{T} \sum_{j=1}^{i-1} U_{ij} \otimes V_{ij}(\theta_{ij}) + U_{ij}' \otimes V_{ij}'(\theta_{ij}) + \sum_{i=1}^{T} U_{ii} \otimes V_{ii}(\theta_{ii})$

where U_{ij} is a $T \times T$ matrix with (i,j)th element equal to 1 and all other elements equal to 0. The following Proposition holds for such $V(\theta)$.

Proposition II.14. *If $V(\theta)$ satisfies (2.14) then $V(\theta^{(1)}) = V(\theta^{(2)})$ implies $\theta^{(1)} = \theta^{(2)}$ iff $V_{ij}(\phi^{(1)}) = V_{ij}(\phi^{(2)})$ implies $\phi^{(1)} = \phi^{(2)}$ for all $i = 1, \cdots, T$ and all $j = 1, \cdots, i$.*

Proof: Straightforward since θ_{ij} represents a partition of θ and V_{ij} represents the corresponding partition of V.

□

The advantage of such models is as follows. Suppose $Cov(\overline{X}) = V(\theta)$ specified by equation (2.14). Then $Cov(\overline{Y}) = \overline{A}V(\theta)\overline{A}'$. However,

$$
\begin{aligned}
\overline{A}(U_{ij} \otimes V_{ij}(\theta_i j))\overline{A}' &= (I \otimes A)(U_{ij} \otimes V_{ij}(\theta_{ij}))\overline{A}' \\
&= (U_{ij} \otimes AV_{ij}(\theta_{ij}))\overline{A}' \\
&= (U_{ij} \otimes AV_{ij}(\theta_{ij}))(I \otimes A') \\
&= U_{ij} \otimes AV_{ij}(\theta_{ij})A'
\end{aligned}
$$

Thus $Cov(\overline{Y}) = \overline{A}V(\theta)\overline{A}' = V_Y(\theta)$ where

$$
V_Y(\theta) = \sum_{i=2}^{T}\sum_{j=1}^{i-1} U_{ij} \otimes (AV_{ij}(\theta_{ij})A') + U'_{ij} \otimes (AV'_{ij}(\theta_{ij})A') + \sum_{i=1}^{T} U_{ii} \otimes (AV_{ii}(\theta_{ii})A')
$$

Now using Proposition II.14, the necessary and sufficient condition for identifiability of θ from $Cov(\overline{Y})$ is that $AV_{ij}(\phi^{(1)})A' = AV_{ij}(\phi^{(2)})A'$ implies $\phi^{(1)} = \phi^{(2)}$ for all $i = 1, \cdots, T$ and all $j = 1, \cdots, i$.

Example Continued:

To extend the independent connections model to multiple time intervals, we assume that non-zero covariances are possible only between flow volumes belonging to the same flow pair but possibly different time intervals. This implies that for $j < i$, $Cov(X_i, X_j)$ is of the form

(2.15) $$ V_{ij}(\phi) = \sum_{k=1}^{J} \phi_k D_k + \sum_{k=1}^{J} \phi_{J+k} E_k $$

for $\phi \in \mathbb{R}^{2J}$, D_k a $J \times J$ matrix such that

$$[D_k]_{lm} = \begin{cases} 1 & k = l = m \\ 0 & \text{otherwise} \end{cases}$$

and E_k a $J \times J$ matrix such that

$$[E_k]_{lm} = \begin{cases} 1 & k = l = m \mod J/2 \text{ and } l > m \\ 0 & \text{otherwise} \end{cases}$$

Here D_k, $k = 1, \cdots, J$, serve as a basis for diagonal elements and E_k, $k = 1, \cdots, J$, serve as a basis for the acceptable off diagonal elements. Further, $Cov(X_i, X_i)$ is of the form

$$V_{ii}(\phi) = \sum_{k=1}^{J} \phi_k D_k + \sum_{k=1}^{J/2} \phi_{J+k}(E_k + E_k')$$

for $\phi \in \mathbb{R}^{3J/2}$.

With the above representation of $V_{ij}(\cdot)$, using lemma II.6 and Proposition II.14, θ is identifiable from $Cov(\overline{Y})$ for $T > 1$ if and only if

$$\overline{B} = [vec(AD_1 A'), \cdots, vec(AD_J A'), vec(AE_1 A'), \cdots, vec(AE_J A')]$$

has rank $2J$ and

$$\overline{B}_c = [vec(AD_1 A'), \cdots, vec(AD_J A'), vec(A(E_1 + E_1')A'), \cdots, vec(A(E_{J/2} + E_{J/2}')A')]$$

has rank $3J/2$. Clearly, the latter follows from the former.

Note that $AD_k A' = a_k a_k'$ and $AE_k A' = a_k a_{\{k\}}$ where a_k is the kth column of A and $\{k\} = (k + J/2 - 1 \mod J/2) + 1$. Thus

$$\overline{B} = [vec(a_1 a_1'), \cdots, vec(a_J a_J'), vec(a_1 a_{\{1\}}'), \cdots, vec(a_J a_{\{J\}}')]$$

and

$$\overline{B}_c = [vec(a_1 a_1'), \cdots, vec(a_J a_J'), vec(a_1 a_{\{1\}}' + a_{\{1\}} a_1'), \cdots, vec(a_{J/2} a_{\{J/2\}}' + a_{\{J/2\}} a_{J/2}')]$$

Finally, if A is partitioned as $[A_F, A_R]$ then

$$(2.16) \qquad \overline{B} = [A_F \odot A_F, A_R \odot A_R, A_F \odot A_R, A_R \odot A_F]$$

and

$$(2.17) \qquad \overline{B}_c = [A_F \odot A_F, A_R \odot A_R, A_F \odot A_R + A_R \odot A_F]$$

2.6   Multimodal Tomography and Temporal Dependence

Proposition II.15. *Suppose* $Y_P = AX_P$ *and* $Y_B = AX_B$ *with* $(X_{Pi}, X_{Bi})'$, *distributed independently for* $i = 1, \cdots, J$ *and*

$$(2.18) \qquad B = A \odot A$$

has rank J *and the joint characteristic functions of* $(X_{Pi}, X_{Bi})'$ *are either analytic for all* $i = 1, \cdots, J$ *or have no roots in* $\mathbb{R}^2$ *for all* $i = 1, \cdots, J$ *. Then the distribution of* $(X_P', X_B')'$ *is identifiable from* (Y_P', Y_B') *up to a mean ambiguity.*

Proof: Suppose $(X_{Pi}^{(1)}, X_{Bi}^{(1)})'$, $i = 1, \cdots, J$ distributed independently and $(X_{Pi}^{(2)}, X_{Bi}^{(2)})'$, $i = 1, \cdots, J$ distributed independently are such that

$$\begin{pmatrix} AX_P^{(1)} \\ AX_B^{(1)} \end{pmatrix} \stackrel{d}{=} \begin{pmatrix} AX_P^{(2)} \\ AX_B^{(2)} \end{pmatrix}$$

Thus

$$\log E[e^{\iota(t_P' AX_P^{(1)} + t_B' AX_B^{(1)})}] = \log E[e^{\iota(t_P' AX_P^{(2)} + t_B' AX_B^{(2)})}]$$

for appropriate range of values of t_P and t_B (as in Corollary II.11) Define

$$\phi_i(t, s) = \log E[e^{\iota(tX_{Pi}^{(1)} + tX_{Bi}^{(1)})}] - \log E[e^{\iota(tX_{Pi}^{(2)} + tX_{Bi}^{(2)})}]$$

Thus

$$(2.19) \qquad \phi_1(t_P' \alpha_1, t_B' \alpha_1) + \cdots + \phi_J(t_P' \alpha_J, t_B' \alpha_J) = 0$$

For a fixed t_P, using Lemma II.10, $\phi_1, \cdots, \phi_J$ are linear functions in their second argument. Similarly, they are linear in their first argument. Thus $\phi_1, \cdots, \phi_J$ are bi linear. However, with $t_P = t_B$ and using Lemma II.10, $\phi_1(u, u), \cdots, \phi_J(u, u)$ are linear in u. Thus $\phi_1, \cdots, \phi_J$ are linear functions. Hence $(X'_P, X'_B)'$ is identifiable from (Y'_P, Y'_B) up to a mean ambiguity through an argument similar to Corollary II.11.

$\square$

Remark: The above argument can be easily extended to multimodal tomography and time dependence.

Corollary II.16. Independent Sub-Flow Model

Suppose

$$
\begin{pmatrix} X_{Pi} \\ X_{Bi} \end{pmatrix} = \begin{pmatrix} 1 & \cdots & 1 \\ s_1 & \cdots & s_S \end{pmatrix} \begin{pmatrix} Z_{1i} \\ \vdots \\ Z_{Si} \end{pmatrix}
$$

where $0 < s_1 < \cdots < s_S$ and Z_{ki} are distributed independently for $k = 1, \cdots, S$ and $i = 1, \cdots, J$. Then under the conditions of the above Proposition the log of characteristic functions of $Z_{1i}, \cdots, Z_{Si}$; $i = 1, \cdots, J$ are uniquely defined up to a polynomial of degree $max(S-2, 1)$ in a neighborhood of the origin. Consequently, if the nth order cumulants of $Z_{1i}, \cdots, Z_{Si}$; $i = 1, \cdots, J$ are assumed to exist for $n > max(S-2, 1)$, then these cumulants are identifiable.

Remark: More discussion on this model is provided in Section 3.2.2 and Section 3.7. Note that the above model makes the problem even more "under-constrained" since the model now involves $S \times J$ unobserved random variables but only a $2L$

dimensional observed random vector. Thus, even this "weaker identifiability result" is interesting.

Proof: We have shown that the joint distributions (X_{Pi}, X_{Bi}); $i = 1, \cdots, J$ are uniquely identified up to mean. Suppose $Z_{ki}^{(1)}$ distributed independently for $k = 1, \cdots, S$ and $i = 1, \cdots, J$ and $Z_{ki}^{(2)}$ distributed independently for $k = 1, \cdots, S$ and $i = 1, \cdots, J$ be such that the distributions of

$$\begin{pmatrix} X_{Pi}^{(1)} \\ X_{Bi}^{(1)} \end{pmatrix} = \begin{pmatrix} 1 & \cdots & 1 \\ s_1 & \cdots & s_S \end{pmatrix} \begin{pmatrix} Z_{1i}^{(1)} \\ \vdots \\ Z_{Si}^{(1)} \end{pmatrix}$$

and

$$\begin{pmatrix} X_{Pi}^{(2)} \\ X_{Bi}^{(2)} \end{pmatrix} = \begin{pmatrix} 1 & \cdots & 1 \\ s_1 & \cdots & s_S \end{pmatrix} \begin{pmatrix} Z_{1i}^{(2)} \\ \vdots \\ Z_{Si}^{(2)} \end{pmatrix}$$

differ only in mean.

Thus

$$\phi_i(t, s) = \log E[e^{\iota(tX_{Pi}^{(1)} + tX_{Bi}^{(1)})}] - \log E[e^{\iota(tX_{Pi}^{(2)} + tX_{Bi}^{(2)})}] = a_i t + b_i s$$

Define $\phi_{ki}(t) = \log E \exp \iota t Z_{ki}^{(1)} - \log E \exp \iota t Z_{ki}^{(2)}$. Then

$$\phi_i(t_1, t_2) = \sum_{k=1}^{S} \phi_{ki}(t_1 + t_2 s_k)$$

Therefore

$$\sum_{k=1}^{S} \phi_{ki}(t_1 + t_2 s_k) = a_i t_1 + b_i t_2$$

Using Lemma II.8 ϕ_{ki} are polynomials of degree less than or equal to S.

Now suppose $S \geq 3$. If λ_{jki} is the jth order term in ϕ_{ki} then using the same

argument as in Lemma II.10

$$(A\odot)^j \begin{pmatrix} \lambda_{j1i} \\ \vdots \\ \lambda_{jSi} \end{pmatrix} = 0$$

for

$$A = \begin{pmatrix} 1 & \cdots & 1 \\ s_1 & \cdots & s_S \end{pmatrix}$$

and $j > 1$.

We will show that $(A\odot)^{S-1}$ is full rank and thus λ_{jki} are 0 for $j = S - 1, S$ and $i = 1, \cdots, J$ and $k = 1, \cdots, S$. Note that $(A\odot)^{S-1}$ contains the matrix

$$\begin{pmatrix} 1 & \cdots & 1 \\ s_1 & \cdots & s_S \\ \vdots & \ddots & \vdots \\ s_1^{S-1} & \cdots & s_S^{S-1} \end{pmatrix}$$

which is a Vandermonde matrix and hence full rank. Further, $(A\odot)^S$ is full rank from Lemma II.3.

□

Remark: A constructive proof of the above is given in Section 2.8.3.

Corollary II.17. Compound Model

Suppose

1. *$Prob(X_P \in \mathbb{N}) = 1$ and*

2. *The distribution of X_P is non-trivial i.e. there is no $n \in \mathbb{N}$ such that $Prob(X_P = n) = 1$ and*

3. $X_B = \sum_{i=1}^{X_P} S_i$ where $X_P, S_1, S_2, \cdots$ are distributed independently and $S_1, S_2, \cdots$ are distributed identically and

4. The distribution of S_1 is non-trivial i.e. there is no $s \in \mathbb{R}$ such that $Prob(S_1 = s) = 1$.

Under the conditions of the above Proposition and additionally the above assumptions of a compound model on each $(X_{Pi}, X_{Bi})'$ for $i = 1, \cdots, J$, the distribution of $(X_P', X_B')'$ is fully identifiable.

Remark: More discussion on this model is provided in Section 3.2.1. Note that this is the only model in this chapter where the mean is also identifiable.

Proof: We have already shown that the distributions of $(X_{Pi}, X_{Bi})'$ are identifiable up to mean. We will show that two distributions of $(X_{Pi}, X_{Bi})'$ satisfying the compound model cannot differ in mean.

Suppose

(2.20)
$$\begin{pmatrix} X_P^{(1)} \\ X_B^{(1)} \end{pmatrix} \overset{d}{=} \begin{pmatrix} X_P^{(2)} \\ X_B^{(2)} \end{pmatrix} + \begin{pmatrix} n_P \\ c_B \end{pmatrix}$$

for $n_P \geq 0$ (without loss of generality) and the joint distributions of $(X_P^{(1)}, X_B^{(1)})'$ and $(X_P^{(2)}, X_B^{(2)})'$ satisfying the assumptions of the compound model, specifically

$$X_B^{(1)} = \sum_{i=1}^{X_P^{(1)}} S_i^{(1)}$$

and

$$X_B^{(2)} = \sum_{i=1}^{X_P^{(2)}} S_i^{(2)}$$

From (2.20)

$$X_B^{(1)} | X_P^{(1)} = n \overset{d}{=} X_B^{(2)} + c_B | X_P^{(2)} = n - n_P$$

Since the distributions of $X_P^{(1)}$ and $X_P^{(2)}$ are assumed to be non-trivial from the definition of the compound model, there exist n_1, n_2 with $n_1 < n_2$ such that

(2.21) $$S_1^{(1)} + \cdots + S_{n_i}^{(1)} \overset{d}{=} S_1^{(2)} + \cdots + S_{n_i - n_P}^{(2)} + c_B$$

for $n_i = n_1, n_2$.

Now, clearly $n_1 \geq n_P$. If $n_1 = n_P$ then $n_1 = 0$ since the right hand side of (2.21) is deterministic and the distribution of $S_1^{(1)}$ is non-trivial. This in turn implies that $c_B = 0$ and we are done. Now assume $0 \leq n_P < n_1 < n_2$ and define

$$\phi_k(t) = \log E[e^{\iota t S_1^{(k)}}]$$

in a neighborhood of 0 and for $k = 1, 2$. Now (2.21) can be written as

$$n_i \phi_1(t) = (n_i - n_P)\phi_2(t) + \iota t c_B$$

for $n_i = n_1, n_2$. Since $0 < n_1 < n_2$ we can eliminate $\phi_1(t)$ from the two equations to get

$$\left(\frac{n_1 - n_P}{n_1}\right)\phi_2(t) + \frac{\iota t c_B}{n_1} = \left(\frac{n_2 - n_P}{n_2}\right)\phi_2(t) + \frac{\iota t c_B}{n_2}$$

and thus

$$\frac{n_P(n_1 - n_2)}{n_1 n_2}\phi_2(t) = \frac{\iota t c_B(n_1 - n_2)}{n_1 n_2}$$

If $n_P = 0$ then $c_B = 0$ and we are done. Otherwise, $\phi_2(t)$ is given as linear function in t in a vicinity of 0. If the characteristic function is analytic at 0 then the distribution is determined by the characteristic function in the vicinity of 0. Further, if the log of characteristic function is linear then the corresponding distribution is trivial in the sense not allowed for distribution of X_P in the definition of compound model. Thus we get a contradiction.

$\square$

2.7 Sufficient Conditions on Routing for Identifiability

In this Section, we derive sufficient conditions on the network routing that guarantee full rankness of the matrices appearing in previous identifiability theorems.

2.7.1 Independent Connections Model and Minimum Weight Routing

Recall from Section 2.5 that for $K = 1$, the second order cumulants, Σ_X, of the independent connections model (2.15), are identifiable if and only if $\overline{B}_c$ is full rank. For $K > 1$, the second order cumulants of the independent connections model are identifiable if and only if $\overline{B}$ is full rank.

In the following, assume that the routing matrix, A, is *binary* and that each flow traverses exactly one path (deterministic routing), i.e. $|\mathcal{P}(j)| = 1$ for $j = 1, \cdots, J$. Define the operator $\mathcal{R}(\cdot)$ on paths such that if path $P = (m_1, m_2, \cdots, m_{k-1}, m_k)$ then

$$\mathcal{R}(P) = (m_k, m_{k-1}, \cdots, m_2, m_1).$$

Also, if $\mathcal{P}$ is a set of paths then $\mathcal{R}(\mathcal{P}) = \{\mathcal{R}(P) : P \in \mathcal{P}\}$. A weighted graph has positive weights associated with each edge, $\mathcal{W}(e) > 0$ for all $e \in E$, the edge set. The weight of a path P is defined as the sum of weights of all edges in it, i.e. $\mathcal{W}(P) = \sum_{e \in P} \mathcal{W}(e)$. We call a (directed) graph *symmetric*, if the weight on edge (n_1, n_2) is the same as the weight on edge (n_2, n_1), for all edges (n_1, n_2). A path P from node n_1 to node n_2 is called a *minimum weight* path, if there is no path P' from n_1 to n_2 with $\mathcal{W}(P') < \mathcal{W}(P)$. Also, we will call a (minimum weight) routing scheme *balanced* if the path of the flow from node n_1 to node n_2 is the reverse of the flow from n_2 to n_1. In other words, if the traffic from a node n_1 to a node n_2 is carried on path P, then the traffic from node n_2 to n_1 is carried on $\mathcal{R}(P)$.

Lemma II.18. *Given a symmetric directed graph the following are true:*

1. *Given any non-empty set $\mathcal{P}$ of minimum weight paths, there exist (possibly identical) edges (f_1, f_2) and (l_1, l_2) such that (f_1, l_2) is the unique pair of nodes (k_1, k_2), for which there exists a minimum weight path $P_1 \in \mathcal{P}$ from k_1 to k_2 containing edges (f_1, f_2) and (l_1, l_2). These edges are the first and last edges of a path with maximum weight in the set $\mathcal{P}$.*

2. *Given non-empty disjoint sets $\mathcal{P}_1, \mathcal{P}_2$ of minimum weight paths such that $\mathcal{R}(\mathcal{P}_1) = \mathcal{P}_2$, there exist edges (f_1, f_2) and (l_2, l_1) such that (f_1, l_2) is the unique pair of nodes (k_1, k_2) for which there exist minimum weight paths $P_1 \in \mathcal{P}_1$ and $P_2 = \mathcal{R}(P_1)$ from k_1 to k_2 and from k_2 to k_1, respectively, containing edges (f_1, f_2) and (l_2, l_1) respectively. These edges are the first edges of paths P_M and $\mathcal{R}(P_M)$ respectively where P_M is a path with maximum weight in the set $\mathcal{P}_1$.*

3. *Let (f_1, f_2) and (l_1, l_2) be the (possibly identical) first and last edges of a minimum weight path P. Then, there is no node pair, k_1 and k_2, such that (f_1, f_2) lies in a minimum weight path P_1 from k_1 to k_2 and (l_1, l_2) lies in $\mathcal{R}(P_1)$. Also, there is no node pair k_1 and k_2 such that (f_1, f_2) and (l_2, l_1) belong to a minimum weight path from k_1 to k_2.*

Proof. To prove the first two claims note that if $P_1 = (f_1, f_2, \cdots, l_1, l_2)$ is a minimum weight path then any path P_2 that contains edges (f_1, f_2) and (l_1, l_2) will have weight greater than P_1 unless it is also a (minimum weight) path from f_1 to l_2. This becomes clear when one considers the two possible cases, i.e. if edge (f_1, f_2) precedes edge (l_1, l_2) in P_2 or if edge (l_1, l_2) precedes edge (f_1, f_2) in P_2. In both cases P_2 would have a larger weight than P_1. The first two claims now follow easily.

The third claim can be proved by contradiction. Suppose there exist nodes k_1, k_2, minimum weight path P_1 from k_1 to k_2 and path $P_2 = \mathcal{R}(P_1)$ such that edge (f_1, f_2)

lies in P_1 and edge (l_1, l_2) lies in P_2. This implies that (l_2, l_1) lies in P_1. We will show that P and P_1 cannot both be minimum weight and this proves both assertions of the third claim. In the following "+" represents the concatenation operation where appropriate and $\mathcal{W}(P)$ is the weight of the path P. Clearly $(f_1, f_2) = (l_1, l_2)$ is not possible as that would mean the P_1 contains both (f_1, f_2) and $(f_2, f_1) = (l_2, l_1)$. Let S be the (minimum weight) path from f_2 to l_1 in P. Now P_1 can have 2 possible structures:
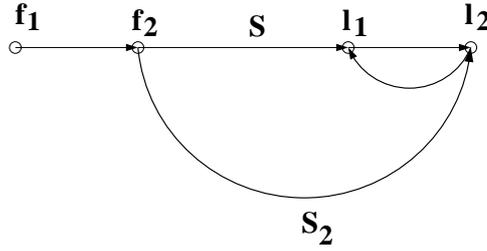
Figure 2.8: Lemma II.18, Case 1
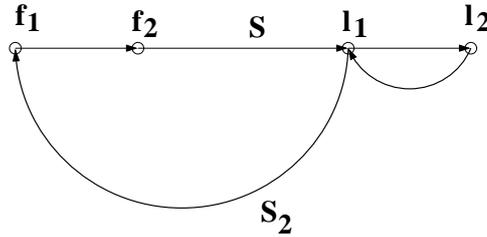
Figure 2.9: Lemma II.18, Case 2

- Case 1. $P_1 = S_1 + (f_1, f_2) + S_2 + (l_2, l_1) + S_3$ (figure 2.8)

  Since both P and P_1 are minimum weight paths, we have that

  $$\mathcal{W}(S_2) = \mathcal{W}(S) + \mathcal{W}(l_1, l_2)$$

  This implies that

  $$\mathcal{W}(S_2) + \mathcal{W}((l_2, l_1)) > \mathcal{W}(S)$$

  which gives that P_1 is not a minimum weight path.

- Case 2. $P_1 = S_1 + (l_2, l_1) + S_2 + (f_1, f_2) + S_3$ (Figure 2.9)

  Since both P and P_1 are minimum weight paths, we get that

  $$\mathcal{W}(S_2) = \mathcal{W}(S) + \mathcal{W}(f_1, f_2)$$

  This implies that

  $$\mathcal{W}(S_2) + \mathcal{W}((f_1, f_2)) > \mathcal{W}(S)$$

  Assuming symmetric weights, the weight of $\mathcal{R}(S)$ is also $\mathcal{W}(S)$. This in turn implies that P_1 is not a minimum weight path.

This proves the third claim and hence the Lemma.

□

One can now establish the following:

Proposition II.19. *Under balanced minimum weight routing on a symmetric graph the matrix $\overline{B}$ (equation (2.16)) is full rank.*

Proof. Trivially re-define $\overline{B}$ as

$$(2.22) \qquad \overline{B} = [A_F \odot A_F, A_F \odot A_R, A_R \odot A_F, A_R \odot A_R]$$

Let

$$v \equiv \overline{B} \begin{pmatrix} v_{FF} \\ v_{FR} \\ v_{RF} \\ v_{RR} \end{pmatrix} = 0,$$

where $v_{FF}, v_{FR}, v_{RF}, v_{RR} \in \mathbb{R}^{J/2}$. We need to show that $v_{FF} = v_{FR} = v_{RF} = v_{RR} = 0$. Let $\mathcal{F}(i, F)$ be the forward flow path for node pair i and $\mathcal{F}(i, R)$ the reverse flow path for the same node pair i, $i = 1, \cdots, J/2$. Here the ordering of node-pairs corresponds

to the ordering of flows assumed in X_F and X_R in equation (2.2). Define operators $\mathcal{P}_F$ and $\mathcal{P}_R$ which map a set of indices to sets of paths as follows:

$$\mathcal{P}_F(I) = \{P : P = \mathcal{F}(i, F) \text{ for some } i \in I\},$$

$$\mathcal{P}_R(I) = \{P : P = \mathcal{F}(i, R) \text{ for some } i \in I\}.$$

Now, define

$$\mathcal{A} = \{\mathcal{F}(i, F) : v_{FF}(i) \neq 0\} \cup \{\mathcal{F}(i, R) : v_{RR}(i) \neq 0\}$$

and

$$I = \{i : v_{FR}(i) \neq 0\} \cup \{i : v_{RF}(i) \neq 0\}.$$

We will show that when $\mathcal{A}$ is non empty, there exists an element in v which is non-zero and when I is non empty there exists another element in v which is non-zero. Use $\mathcal{A}$ as the set of paths in the first part of Lemma II.18 to identify edges (f_1, f_2) and (l_1, l_2) which are traversed by exactly one flow (say) $F_M \in \mathcal{A}$. Now, recall that each ordered pair of link indices (r_1, r_2), corresponds to a row $g(r_1, r_2)$ in $\overline{B}$. Consider the row of $\overline{B}$ corresponding to (f_1, f_2) and (l_1, l_2):

$$r_1 = (r_{FF}^{(1)}, r_{FR}^{(1)}, r_{RF}^{(1)}, r_{RR}^{(1)}).$$

Note that elements of $r_{FF}^{(1)}$ and $r_{RR}^{(1)}$ indicate the forward and reverse flows common to links (f_1, f_2) and (l_1, l_2), elements of $r_{FR}^{(1)}$ indicate node-pairs for which forward flow traverses (f_1, f_2) while reverse flow traverses (l_1, l_2) and elements of $r_{RF}^{(1)}$ indicate node-pairs for which reverse flow traverses (f_1, f_2) while forward flow traverses (l_1, l_2). We then claim the following:

1. $r_{FF}^{(1)}(i) \neq 0$ and $v_{FF}(i) \neq 0$ if and only if $\mathcal{F}(i, F) = F_M$.

2. $r_{RR}^{(1)}(i) \neq 0$ and $v_{RR}(i) \neq 0$ if and only if $\mathcal{F}(i, R) = F_M$.

3. $r_{FR}^{(1)}(i) = r_{RF}^{(1)}(i) = 0$ for all i.

The first two claims follow directly from the first part of Lemma II.18. The third claim follows from the third part of Lemma II.18. Therefore,

$$r_1 \begin{pmatrix} v_{FF} \\ v_{FR} \\ v_{RF} \\ v_{RR} \end{pmatrix} \neq 0$$

and we obtain a contradiction.

Now use $\mathcal{P}_F(I)$ and $\mathcal{P}_R(I)$ as the sets of paths in the second part of Lemma II.18 to identify edges (n_1, m_1) and (n_2, m_2) which are traversed by the forward and reverse flows (or vice versa) of exactly one node pair, say the i_M-th node pair, $i_M \in I$. Consider the row of $\overline{B}$ corresponding to (n_1, m_1) and (n_2, m_2):

$$r_2 = (r_{FF}^{(2)}, r_{FR}^{(2)}, r_{RF}^{(2)}, r_{RR}^{(2)}).$$

Note that $r_{FR}^{(2)}(i) r_{RF}^{(2)}(i) = 0$ for all i. Now we claim the following:

1. $|r_{FR}^{(2)}(i) v_{FR}(i)| + |r_{RF}^{(2)}(i) v_{RF}(i)| \neq 0$ if and only if $i = i_M$.

2. $r_{FF}^2(i) = r_{RR}^2(i) = 0$ for all i.

The first claim follows directly from the second part of Lemma II.18. The second claim follows from the third part of Lemma II.18. Therefore,

$$r_2 \begin{pmatrix} v_{FF} \\ v_{FR} \\ v_{RF} \\ v_{RR} \end{pmatrix} \neq 0.$$

Thus at least one of the rows of v will be non-zero for $\mathcal{A}$ and/or I non-empty. This completes the proof of the result. $\square$

Corollary II.20. *The matrix $\overline{B}_c$ (equation (2.17)) and B (equation (2.18)) are full rank under balanced minimum weight routing on a symmetric graph.*

2.7.2 Independent Connections Model and Hierarchical Graphs

The conditions of minimum cost routing and deterministic routing are not required for proving identifiability in special classes of networks. In one of the early papers on network tomography, Cao et.al [6] proved that under independence, flow volume variances are identifiable if the network has a hierarchical structure. In such a structure, there exists a set of "internal" nodes that neither generate nor sink traffic. Flows exist only between pairs of non-internal (terminal) nodes, which are only connected to internal nodes and not to other non-internal nodes directly. This is a reasonable model if the network under consideration corresponds to a combination of a backbone network and sub networks, with the latter being connected amongst themselves through the backbone network. Hence, the nodes of the backbone network are considered internal nodes.

When there is no dependence between forward and reverse flows and only one type of measurement is considered, identifiability depends on full rankness of $B = A \odot A$. The matrix B can easily be shown to be full rank for hierarchical networks. The proof (see [6]) rests on the fact that for all flows, there exist rows in B which have exactly one non-zero entry occurring at the corresponding indices. For any flow, consider the edge that connects the source node to the first internal node and the edge that connects the last internal node to the destination node. The only flow common to these two edges is the flow under consideration. Thus, the row in B

corresponding to this pair of edges has exactly one non-zero entry occurring at the index corresponding to the flow under consideration. Notice that neither minimum cost routing, nor deterministic routing is required for the argument. In fact, the matrix $\overline{B}$ (equation (2.16)), and hence $\overline{B}_c$ (equation (2.17)), can be shown to be full rank, which implies identifiability, as previously argued, under the independent connections model.

Proposition II.21. *The matrix $\overline{B}$ (equation (2.22)) is full rank for hierarchical networks.*

Proof. We will prove that given $i \in \{1, \cdots, 2J\}$ there exists a row r in $\overline{B}$ such that r is the ith row of a $2J \times 2J$ identity matrix. Index i corresponds to flow pair $i' = ((i-1) \bmod J/2) + 1$. For $i \in \{1, \cdots, J/2\}$ ($i \in \{3J/2, \cdots, 2J\}$) choose ordered pair (i_1, i_2) to be the indices of the first and last edges respectively of the forward (reverse) flow of flow-pair i'. For $i \in \{J/2+1, \cdots, J\}$ ($i \in \{J+1, \cdots, 3J/2\}$) choose ordered pair (i_1, i_2) to be the indices of the first edges respectively of the forward and reverse (reverse and forward) flows of flow-pair i'. Now, choosing r to be the $g(i_1, i_2)$-th row of $\overline{B}$ gives the required result. $\square$

In the following, we use a similar idea to prove full rankness of $B = A \odot A$.

2.7.3 Directed Acyclic Graphs

A directed graph with no cycles is called a Directed Acyclic Graph (DAG). An important example of a DAG is a tree. Clearly there are no reverse (say) flows and $A = A_F$. Thus, identifiability depends on the full rankness of $B = A \odot A$.

Proposition II.22. *For a directed acyclic graph, the matrix B is full rank.*

Proof. Note that all finite DAG have at least one root node. Define $d(n)$ for a node n to be the maximum of length of paths from any root node to n. Also define

$d(n) = 0$ for n being a root node. Note that if there is a path from node n_1 to node n_2 of length l, then $d(n_1) \leq l + d(n_2)$. For flow f, define $\tilde{d}(f) = d(n_2) - d(n_1)$, where n_2 and n_1 are the destination and origin nodes of flow f, respectively.

Now suppose $Bx = 0$ for $x \neq 0$. Consider the set $\mathcal{P}_x$ of paths traversed by flows corresponding to non-zero entries in x. Let P be defined as follows:

$$P = \underset{P' \in \mathcal{P}_x}{\operatorname{argmax}} \ \tilde{d}(P').$$

Let e_1, e_2 be the first and last edges of P, and n_1, n_2 its origin and destination node, respectively. It can be shown that the flow f from n_1 to n_2 is the only flow for which the corresponding entry is non-zero in x and that traverses both e_1 and e_2. If not, let f' be another flow corresponding to a non-zero entry in x that traverses both e_1 and e_2. Let n_1' and n_2' be the origin and destination nodes of flow f'. Since e_1 is traversed by f', there exists a path from n_1' to n_1 and thus $d(n_1') \leq d(n_1)$, with equality if and only if $n_1' = n_1$. Similarly, $d(n_2') \geq d(n_2)$, with equality if and only if $n_2' = n_2$. Thus, for any path P' of f', $\tilde{d}(P') \geq \tilde{d}(P)$, with equality if and only if $n_1' = n_1$ and $n_2' = n_2$. But $P' \in \mathcal{P}$, since f' corresponds to a non-zero entry in x. Thus $f' = f$.

Now, consider the row r in B corresponding to edges e_1 and e_2. There is exactly one index i for which $x_i \neq 0$ and $r_i \neq 0$. Thus, $Bx \neq 0$ which is a contradiction. $\square$

Remark: Note that the above proof does not require deterministic routing. Further, it seems that it does not require minimum cost routing. However, it is easy to construct weights, such that any routing scheme in a DAG is a minimum cost routing scheme. Simply use $d(n_2) - d(n_1)$ as the weight of the edge from n_1 to n_2. A telescoping sum argument implies that any path from a node n_1 to node n_2 has weight $d(n_2) - d(n_1)$

and therefore all paths are minimum cost ones.

Remark: Also note that in general, first moments of flows would not be identifiable in a DAG based on link measurements alone as the matrix A would have more columns than rows. However, A can be shown to be full rank for DAGs under the following conditions:

1. Only flows originating at a root node are present.

2. Only flows terminating at a leaf node are present.

The proof is straightforward. Assume that the first condition is true (the argument for the second condition is analogous). Suppose $Ax = 0$ for a non-zero x. Let $\mathcal{P}$ be the set of paths traversed by flows corresponding to non-zero entries in x. Select $P \in \mathcal{P}$ with maximum weight under the weighting scheme described above. Let r be the row in A corresponding to the last edge in P. Then, $r_i x_i \neq 0$ if and only if flow i corresponds to P. Hence, $rx \neq 0$ and we obtain a contradiction.

2.8   Discussion and Future Work

2.8.1   Connections with Prior Results

Some of the results presented in this chapter are closely related to previous results in literature. Lemma II.10 is a slightly stronger version of Lemma 1.5.4 in [26] under stronger conditions. Corollary II.11 is a slightly stronger version of Theorem 1 in [7]. A somewhat incomplete outline of the equivalent of Proposition II.19 for the independence case appears in [42]. As already noted the equivalent of Proposition II.21 for the independence case is shown in [6]. Finally, the equivalent of Propositions II.19 and II.21 for the independence case for delay/loss tomography are well known [7]. The remaining results do not have a close analogue in existing literature.

2.8.2 Spatial Block Independence

While the spatial dependence models in Section 2.4 are fairly general, there still remains the question of whether it is possible to prove results approaching the generality of those in Section 2.6. Note that spatial dependence is naturally harder to handle than inter-modality dependence or temporal dependence, since it is the spatial domain where the problem is under-constrained. Nevertheless, it is tempting to attempt to use the techniques from Section 2.6 to prove general identifiability results when the flows can be partitioned into blocks, each independent of the other. We refer to this as spatial block independence and the independent connections model of equation (2.3) is a special case. However, in trying to use the technique of Proposition II.15 one quickly realizes that Lemma II.10 is no longer useful. It appears that, what is called for is a version of Lemma II.10 that deals with multivariable functions. Note that in the proof of Proposition II.15 we were able to handle multivariable functions through repeated use of Lemma II.10. However, spatial block independence does not appear amenable to a similar technique.

2.8.3 A Constructive Proof of Corollary II.16

It is straightforward to give a constructive proof of Corollary II.16 in that it naturally suggests a method of moments estimator for the identifiable cumulants. Assume that for some $n > 1$ the n-th order cumulants of Z_{ki} (defined in Corollary II.16) exist for $k = 1, \cdots, S$ and $i = 1, \cdots, J$. We additionally assume that the routing matrix A is binary.

Proof: Let

$$\phi_{ki}(t) = \log E \exp t Z_{ki}$$

for random variables Z_{ki}, $k = 1, \cdots, S$ and $i = 1, \cdots, J$. Then

$$(2.23) \qquad \log E[e^{t_1 Y_{Pj_1} + t_2 Y_{Pj_2}}] = \sum_{k=1}^{S} \sum_{i=1}^{J} \phi_{ki}(t_1 a_{j_1 i} + t_2 a_{j_2 i})$$

for t_1, t_2 in a vicinity of 0 and $a_{ji} = [A]_{ji}$ for $j = 1, \cdots, L$ and $i = 1, \cdots, J$. Also recall that row $g(j_1, j_2) \equiv (j_1 - 1)L + j_2$ in $B \equiv A \odot A$ is equal to the element-wise product of row j_1 and j_2 in A. Define vector $\phi_{PP}^{(n-k,k)}$ of length L^2 as

$$[\phi_{PP}^{(n-k,k)}]_{g(j_1,j_2)} = \frac{\partial^n \log E[e^{t_1 Y_{Pj_1} + t_2 Y_{Pj_2}}]}{\partial t_1^{n-k} \partial t_2^k} \Big|_{(t_1,t_2)=(0,0)}$$

This is the vector of observed cumulants. Further, define $S \times J$ matrix of nth order (unobserved) cumulants Φ_n as

$$[\Phi_n]_{ki} = \frac{\partial^n \phi_{ki}(t)}{\partial t^n} \Big|_{t=0}$$

Using equation (2.23) and after some algebra, it is easy to see that

$$(2.24) \qquad \phi_{PP}^{(n-k,k)} = B \Phi_n' \mathbf{1}$$

where $\mathbf{1}$ is a vector of length S with all entries equal to 1.

Now, define $\phi_{PB}^{(n-k,k)}$ and $\phi_{BB}^{(n-k,k)}$, both of length L^2 as

$$[\phi_{PB}^{(n-k,k)}]_{g(j_1,j_2)} = \frac{\partial^n \log E[e^{t_1 Y_{Pj_1} + t_2 Y_{Bj_2}}]}{\partial t_1^{n-k} \partial t_2^k} \Big|_{(t_1,t_2)=(0,0)}$$

and

$$[\phi_{BB}^{(n-k,k)}]_{g(j_1,j_2)} = \frac{\partial^n \log E[e^{t_1 Y_{Bj_1} + t_2 Y_{Bj_2}}]}{\partial t_1^{n-k} \partial t_2^k} \Big|_{(t_1,t_2)=(0,0)}$$

Similar to equation (2.24) it is easy to show

$$(2.25) \qquad \phi_{PB}^{(n-k,k)} = B \Phi_n' s^{(k)}$$

and

$$(2.26) \qquad \phi_{BB}^{(n-k,k)} = B \Phi_n' s^{(n)}$$

where $s^{(k)}$ is a vector of length S with m-th entry equal to s_m^k, and $s_1, \cdots, s_S$ as defined in Corollary II.16. Thus

$$\Phi_Y^{(n)} \equiv [\phi_{PP}^{(1,n-1)}, \phi_{PB}^{(n-1,1)}, \cdots, \phi_{PB}^{(1,n-1)}, \phi_{BB}^{(1,n-1)}] = B\Phi_n'[\mathbf{1}, s^{(1)}, \cdots, s^{(n-1)}, s^{(n)}]$$

This implies

$$vec(\Phi_Y^{(n)}) = \left(\left[\mathbf{1}, s^{(1)}, \cdots, s^{(n)}\right]' \otimes B\right) vec(\Phi_n') \equiv B_n vec(\Phi_n')$$

Since B is assumed to be full rank, from the property of Kronecker product, B_n is full rank if

$$\left[\mathbf{1}, s^{(1)}, \cdots, s^{(n)}\right]$$

is full rank. The above is a Vandermonde matrix and is full rank for $n > S - 2$. Thus unobserved cumulants $vec(\Phi_n')$ are identifiable from observed cumulants $vec(\Phi_Y^{(n)})$ if $n > max(S - 2, 1)$.

$\square$

2.8.4   State Space Models

One common way to describe spatio-temporal dependence is through state space models and one would like to study identifiability under such models. A state space model corresponding to the current context can be defined as follows. Suppose $X_0, X_1, \cdots$, be J dimensional random vectors each with mean 0. We assume and

$$(2.27) \qquad\qquad\qquad\qquad X_t = CX_{t-1} + \epsilon_t$$

where $C \in \mathcal{C}$ and $Cov(\epsilon_t) = \Sigma \in \mathcal{S}$. Further, we assume

$$(2.28) \qquad\qquad\qquad\qquad Y_t = AX_t$$

We would like to investigate conditions on A with respect to $\mathcal{C}$ and $\mathcal{S}$ under which C and Σ are identifiable from $Cov(Y_t, Y_{t+h})$ for $t = 1, \cdots, T$ and $h = 0, \cdots, T - t$. Note that if $\mathcal{C}$ only includes matrices with spectral radius less than 1, then there exists a Φ, such that for $Cov(X_0) = \Phi$ the above state space model is stationary. In fact,

$$\Phi = C\Phi C' + \Sigma$$

which can be solved as

$$vec(\Phi) = (I - C \otimes C)^{-1} vec(\Sigma)$$

Further, $Cov(X_{t+h}, X_t) = C^h \Phi$. Thus $Cov(Y_{t+h}, Y_t) = AC^h \Phi A'$ and

$$
\begin{aligned}
vec(Cov(Y_{t+h}, Y_t)) &= vec(AC^h \Phi A') \\
&= (A \otimes (AC^h)) vec(\Phi) \\
(2.29) \qquad &= (A \otimes (AC^h))(I - C \otimes C)^{-1} vec(\Sigma)
\end{aligned}
$$

However, the above expression is highly non-linear in the quantities of interest, i.e. C and Σ. Thus the question of identifiability of C and Σ from $Cov(Y_{t+h}, Y_t)$, for $t = 1, \cdots, T$ and $h = 0, \cdots, T - t$, cannot be straightforwardly addressed using techniques of Section 2.4.

2.8.5 Network routing and a counter-example.

In Section 2.7.1, it was shown that minimum cost routing and symmetric weights was sufficient to ensure the full rankness of the $\overline{B}$, and hence B, matrices. The following example shows that absence of these conditions renders the result invalid.

Consider the network shown in Figure 2.7. Here

$$
\begin{pmatrix} Y_{(1,2)} \\ Y_{(2,3)} \\ Y_{(3,4)} \\ Y_{(4,1)} \end{pmatrix} = \begin{pmatrix} 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 \\ 0 & 1 & 1 & 1 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 & 1 & 1 & 1 & 1 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 & 1 \end{pmatrix} \begin{pmatrix} X_{(1,2)} \\ X_{(1,3)} \\ X_{(1,4)} \\ X_{(2,3)} \\ X_{(2,4)} \\ X_{(3,4)} \\ X_{(2,1)} \\ X_{(3,1)} \\ X_{(4,1)} \\ X_{(3,2)} \\ X_{(4,2)} \\ X_{(4,3)} \end{pmatrix}.
$$

Denote by A the matrix above. Thus, neglecting the repetitions in $B = A \odot A$ we get

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

Note that in this case B is a 10×12 matrix and thus cannot be full rank. If symmetric weights are enforced, but not minimum cost routing (or vice-versa) the example would still hold. This example shows that in the absence of minimum cost routing or symmetric weights, then the full rankness of B (and hence $\overline{B}$) is not guaranteed.

2.8.6 Weaker conditions on Routing for Identifiability

The above counter-example shows that the only possibility of relaxing the conditions for Proposition II.19 is to prove the result under non-deterministic routing. To be able to apply the same techniques as in the current proof, given vector $x \neq 0$ we should be able to identify a row r in $\overline{B}$ (or $B = A \odot A$ for the independence case) such that $r_i x_i \neq 0$ for exactly one i. The row r is identified as the row corresponding to the terminal edges of a "maximal" flow. For minimum cost, balanced and deterministic routing, a maximal flow is just the longest flow of a set. For non-deterministic routing a maximal flow P, given a set of flows $\mathcal{P}$ would need to satisfy the following. $P \in \mathcal{P}$ is maximal if there is no pair of node n_1 and n_2 such that there are paths $P_1, P_2 \in \mathcal{P}$ where P_i is from n_1 to n_2 or vice-versa for $i = \{1, 2\}$ and P_1 traverses the first edge of P and P_2 traverses the last edge of P. Simply choosing a path with largest weight in $\mathcal{P}$ would not suffice in this case. It is not clear if such a maximal flow always exists.

In summary, extending Proposition II.19 to the case of non-deterministic routing remains an open problem.

CHAPTER III

Dual Modality Network Tomography

3.1   Introduction

The classical network tomography [47, 6] setup does not consider packet volumes
and byte volumes simultaneously. In this chapter, we use the fact that these two
measures of flow volume are related through the packet size distribution. Motivated
by empirical evidence, we investigate two models that capture the relationship be-
tween packet and byte volumes. In the first model, we assume a compound structure
for the byte volume with the packet volume as the compounding variable. In the
second model, we assume that each flow is made up of independent sub-flows, each
with a fixed packet size. For both models we make some network wide assumptions
in the spirit of classical network tomography. These assumptions attempt to utilize
the structural relationship between packet volume and byte volume of a flow. They
can be viewed as a type of regularization that enables us to estimate flow volume
means from aggregated data. The models introduced in this study try to capture the
main characteristics of the packets-bytes relationship, although the true one may be
more complex as evidenced by the plots of three flows obtained from real network
traces shown in Figure 3.1. Experience suggests that such complex relationships
tend to be present in not highly aggregated flows. The remainder of the chapter is
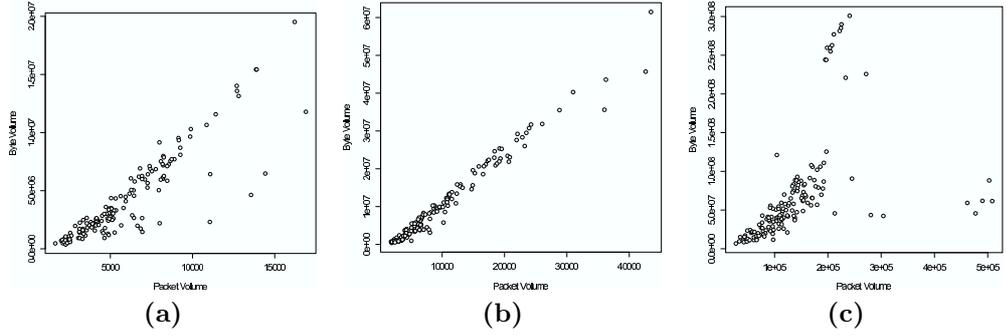
Figure 3.1: Byte volume versus packet volume for 3 observed flows

organized as follows: in Section 3.2 we introduce the proposed flow volumes, while in Section 3.3 we address identifiability issues. In Section 3.4 we study estimation of the models based on a pseudo-likelihood framework and establish consistency and asymptotic normality of the estimators. The performance of the models on simulated and emulated data is assessed in Section 3.5. The issue of estimating characteristics of the packet size distribution is examined in Section 3.6. Finally, some concluding remarks are drawn in Section 3.7.

3.2 The Flow Models

Suppose there are J flows and L directed links in a network. Let A be the $L \times J$ routing matrix such that $A_{ij} = 1$ if flow j traverses link i and 0 otherwise. We assume a deterministic routing policy. Further, let X_t^P and X_t^B be vectors of length J whose elements are packet and byte volumes of the flows in time interval t for $t = 1, \cdots, T$. Define

$$X_t = \begin{pmatrix} X_t^P \\ X_t^B \end{pmatrix}$$

and aggregate SNMP measurements Y_t

$$Y_t = \begin{pmatrix} A & 0 \\ 0 & A \end{pmatrix} X_t \equiv \overline{A} X_t.$$

Further, assume that $\{X_t\}_{t=1}^T$ is a stationary sequence.

3.2.1 Compound Model

We assume

(3.1)
$$X_{t,j}^B = \sum_{k=1}^{X_{t,j}^P} S_{kj}$$

where $\{S_{kj} : k = 1, 2, \cdots\}$ is the size (in bytes) of the kth packet of the jth flow in time interval t (we suppress the time interval indexing for notational convenience). It is assumed that $\{S_{kj} : k = 1, 2, \cdots\}$ are independent and identically distributed (i.i.d.) from some distribution $\mathcal{F}_j$, corresponding to flow j, and *independent* of the packet count $X_{t,j}^P$, for $j = 1, \cdots, J$. Further it is assumed that the packet counts $X_{t,j}^P$, $j = 1, \cdots, J$, are independent across flows. Additionally define the following parameters:

1. Mean packet volume vector $(J \times 1)$, $\mu = E[X_t^P]$.

2. Packet volume variance vector $(J \times 1)$, s_P, i.e. $[s_P]_j = var(X_{t,j}^P)$. Further from our assumption of independence of packet counts, $Cov(X_t^P) = Diag(s_P)$ (this assumption can be relaxed to include the most significant empirically observed spatial correlations in flow volumes, i.e. the ones between forward and reverse flows due to the connection oriented nature of Internet traffic [40]).

3. Mean packet size vector $(J \times 1)$, ψ, i.e. ψ_j is the mean of $\mathcal{F}_j$.

4. Packet size variance vector $(J \times 1)$, v, i.e. v_j is the variance of $\mathcal{F}_j$.

From (3.1), we have $E[X_{t,j}^B | X_{t,j}^P] = \psi_j X_{t,j}^P$ and $Var[X_{t,j}^B | X_{t,j}^P] = v_j X_{t,j}^P$. Now

$$Cov(X_{t,j}^B, X_{t,j}^P) = Cov(E[X_{t,j}^B | X_{t,j}^P], X_{t,j}^P) = \psi_j s_{Pj}$$

and

$$Var(X_{t,j}^B) = Var(E[X_{t,j}^B | X_{t,j}^P]) + E[Var(X_{t,j}^B | X_{t,j}^P)] = \psi_j^2 s_{Pj} + \mu_j v_j$$

Thus

$$\Sigma_X \equiv Cov(X_t) = \begin{pmatrix} Diag(s_P) & Diag(s_P)Diag(\psi) \\ Diag(\psi)Diag(s_P) & Diag(\psi)Diag(s_P)Diag(\psi) + Diag(s) \end{pmatrix}$$

where

$$(3.2) \qquad\qquad\qquad\qquad s_j = \mu_j v_j$$

is the excess variance in byte distribution "not explained by the variance in packet distribution". Collecting the parameters, if $\theta = (s_P, \psi, s)$, then Σ_X is parametrized by θ, i.e. $\Sigma_X(\theta) = \Sigma_X$ and thus

$$\Sigma_Y(\theta) \equiv Cov(Y_t) = \overline{A}\Sigma_X(\theta)\overline{A}'.$$

In the usual setup for traffic demand tomography one assumes a functional relationship between flow volume means and variances for all flows of the type $[s_P]_j = \phi\mu_j^c$. This is a way to get identifiability of mean flow volumes from identifiability of flow volume variance (see for example [6]). In [6], $c = 1$ has been assumed and we will refer to this as *classical tomography*. For comparison purposes, the true relationship between the means and variances of flow volumes in the data examined (see Section 3.5) is shown in Figure 3.2 (a). Joint modeling of packet and byte volume allows us to estimate mean flow volumes under a different assumption on all flows. We will assume (except in Section 3.6) $\psi_j = \psi_0$ and $v_j = v_0$ for all flows j. In this case the mean packet volume μ is identifiable (as shown in section 3.3).

3.2.2 Independent Sub-Flow Model

Another way to jointly model packet and byte flow volumes is to assume that each flow is comprised of independent sub-flows, each with a characteristic packet size. Empirical evidence suggests that just a few packet sizes account for most of the traffic
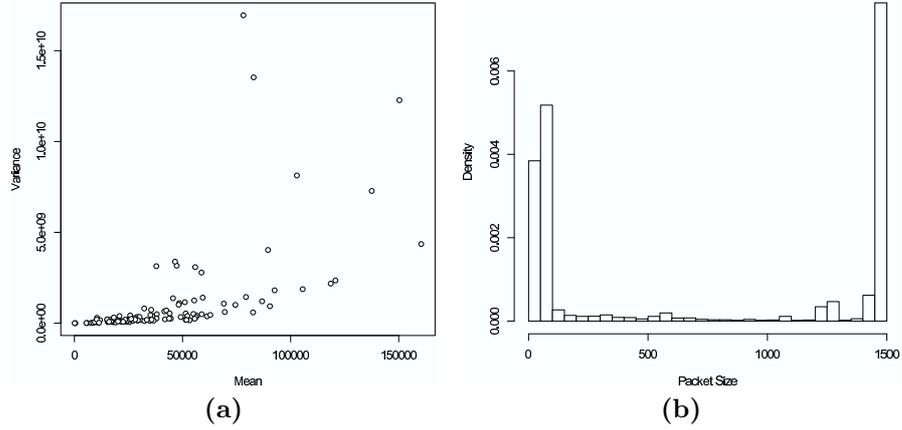
Figure 3.2: Variance versus mean of flow volumes (a) and Observed packet size distribution (b)

in a network. The histogram of the observed packet sizes recovered from header trace of the data described in Section 3.5 is shown in Figure 3.2 (b). These packet sizes are determined by the dominant protocols in the flows. These are typically TCP for web browsing and file transfers and UDP for streaming traffic (e.g. audio and video applications). Other empirical studies have found dominant packet sizes at 40, 576 and 1500 [23]. Thus, we assume that each class of traffic, such as bulk transfers versus streaming traffic, results in one or more sub-flows each with a fixed packet length.

Assume each origin destination flow is made up of S sub-flows each of different "type". Further, all packets of sub-flow of type k have size s_k, for $k = 1, \cdots, S$. Also let $X_{ik}(t)$ be the number of packets in sub-flow k of flow i in a time interval. Let $\mathbf{1}$ be a vector of length S with each element equal to 1 and $s^{(1)}$ be a vector of length S with kth element equal to s_k. Finally, let $\overline{X}_t$ be a $J \times S$ matrix whose (i, k)th element is $[\overline{X}_t]_{ik} = X_{ik}(t)$. Now, the packet volumes vector X^P can be written as

$$X_t^P = \overline{X}_t \mathbf{1},$$

while the byte volume vector X^B as

$$X_t^B = \overline{X}_t s^{(1)}.$$

We will assume that $X_{ik}(\cdot)$ are *independent* for all i and k.

Now let Γ and Φ be $J \times S$ matrices with (i,k)th elements, $[\Gamma]_{ik} = E[X_{ik}(t)]$ and $[\Phi]_{ik} = Var(X_{ik}(t))$, respectively. Under this model we have

$$Var(X_{t,j}^P) = \sum_{k=1}^{S} Var(X_{jk}(t))$$

$$Cov(X_{t,j}^P, X_{t,j}^B) = \sum_{k=1}^{S} s_k Var(X_{jk}(t))$$

and

$$Var(X_{t,j}^B) = \sum_{k=1}^{S} s_k^2 Var(X_{jk}(t)).$$

Then,

$$\Sigma_X \equiv Cov(X_t) = \begin{pmatrix} Diag(\Phi \mathbf{1}) & Diag(\Phi s^{(1)}) \\ Diag(\Phi s^{(1)}) & Diag(\Phi s^{(2)}) \end{pmatrix}$$

where $s^{(2)}$ is a vector formed of element-wise squares of $s^{(1)}$ and

$$E \begin{pmatrix} X_t^P \\ X_t^B \end{pmatrix} = \begin{pmatrix} \Gamma \mathbf{1} \\ \Gamma s^{(1)} \end{pmatrix}.$$

We set $\theta = vec(\Phi)$ and parametrize Σ_X by θ; i.e. $\Sigma_X(\theta) = \Sigma$ and $\Sigma_Y(\theta) = \overline{A}\Sigma_X(\theta)\overline{A}'$. As a regularizing constraint for tomography we will assume $S = 2$ and $\Gamma = \Phi Diag(\alpha)$, as elaborated in Section 3.3. Some comments on the case of $S \geq 3$ are given in Section 3.5.

3.2.3 Equivalence Under Poisson Model

If the packet volumes $X_{t,j}^P$ are distributed as independent Poisson random variables with parameter λ_j for all j and t and all packet size distributions $\mathcal{F}_j$ have finite support such that $\text{Prob}(\{S_{kj} = s_1\} \bigcup \cdots \bigcup \{S_{kj} = s_S\})=1$, then the distribution of X_t is identical to one under independent sub-flow model with Poisson sub-flows. Note

that in this case equation (3.1) can be re-written as

$$X_{t,j}^B = \sum_{i=1}^{S} s_i \sum_{k=1}^{X_{t,j}^P} I(S_{kj} = s_i) \equiv \sum_{i=1}^{S} s_i X_{ji}(t)$$

The independence of $X_{ji}(\cdot)$ for all j and i follows from the independence property of thinned Poisson processes.

3.3 Identifiability and Regularizing Assumptions

In this section, we address the issue of identifiability of the parameters of the two proposed models; i.e. we show that the parameters of interest are uniquely determined by the observed data distribution (or statistics thereof). The strategy for proving identifiability of the parameters in our models has two steps. First, we establish identifiability of parameters associated with the covariance Σ_X and subsequently prove the identifiability of the remaining parameters. The former is based on an identifiability result from the previous chapter.

Let $\mathcal{S}$ be a set of symmetric positive definite matrices, of the form

$$\Sigma_X = \begin{pmatrix} Diag(s_P) & Diag(s_{PB}) \\ Diag(s_{PB}) & Diag(s_B) \end{pmatrix}.$$

where s_P, s_{PB} and s_B are length J vectors of the variance of packet volumes, co-variance of packet and byte volumes and variance of byte volumes, respectively. The following Lemma from proves useful for establishing identifiability.

Lemma III.1. *Under balanced minimum weight routing on a symmetric graph and assuming flow volume covariances $\Sigma_X \in \mathcal{S}$, Σ_X (alternatively (s_P, s_{PB}, s_B)) are iden-tifiable from the covariance of cumulative link measurements, $\Sigma_Y = \overline{A}\Sigma_X\overline{A}'$.*

Proof: The result follows easily from Proposition II.15 and Proposition II.19. While Proposition II.15 is concerned with identifiability of entire distributions (up to

mean), assuming X_P and X_B to be distributed jointly as multivariate normal implies the identifiability of Σ_X from Σ_Y. Viewing this as purely a result on covariance matrices, one can see that joint normality of X_P and X_B is not required.

$\square$

As noted in Section 2.7, the conditions in the above Lemma are usually true in realistic networks. For the rest of the chapter we assume that the conditions in the lemma are met. Thus if $\Sigma_{X1}, \Sigma_{X2} \in \mathcal{S}$ then $\overline{A}\Sigma_{X1}\overline{A}' = \overline{A}\Sigma_{X2}\overline{A}'$ implies $\Sigma_{X1} = \Sigma_{X2}$. Further given a p dimensional *one to one* parametrization $\Sigma_X(\cdot) : \mathbb{R}^p \to \mathcal{S}$, and $\theta_1, \theta_2 \in \mathbb{R}^p$, we have $\overline{A}\Sigma_X(\theta_1)\overline{A}' = \overline{A}\Sigma_X(\theta_2)\overline{A}'$ implies $\theta_1 = \theta_2$.

3.3.1 The Compound Model

As mentioned before, in order to establish identifiability of this model we require the following regularizing assumption. We assume the packet size distribution $\mathcal{F}_j$ is the same for all flows j. As mentioned earlier this implies $\psi_j = \psi_0$ and $v_j = v_0$.

Lemma III.2. *Under balanced minimum weight routing on a symmetric graph and assuming all flows have identical packet size distributions, the parameters of the compound model are identifiable from cumulative link measurements.*

Proof: With $\theta = (s_P, \psi_0, s)$, it is clear that $\Sigma_X(\theta)$ is a one-to-one map. Thus, based on the previous result, θ is identifiable. Identifiability of μ from θ follows from the fact that s is a non-zero vector with non-negative entries and that no non-trivial vector with non-negative entries can lie in the null space of A. This is because all entries in A are non-negative. Thus, $As \neq 0$ and v_0 can be identified from the relation $E[Y_t^P] = As/v_0$. Finally, we get $\mu = s/v_0$.

$\square$

3.3.2 Independent Sub-Flows Model

There may be many different regularizing assumptions that lead to identifiability for tomography under the independent sub-flows model. However, we focus on the following as it works well in practice. First note that in Figure 3.2 (b) that the packet size distribution is concentrated on 2 support points roughly corresponding to streaming traffic ($\sim$ 40 byte payloads) and bulk transfers (1500 byte payloads). Thus, we assume $S = 2$ and $s_1 = 40$ and $s_2 = 1500$ (for identifiability purposes, we only need that $s_1 \neq s_2$). Further, we assume that $\Gamma = \Phi Diag(\alpha)$ for $\alpha = (\alpha_1, \alpha_2)$. This is similar to the assumption of proportionality of means and variances in classical tomography except that we allow for separate proportionality constants, α_1, α_2 for the two sub-flows.

Lemma III.3. *Under balanced minimum weight routing on a symmetric graph and assuming two sub-flows with $s_1 \neq s_2$ and $\Gamma = \Phi Diag(\alpha)$, the parameters of the independent sub-flow model, Φ and Γ, are identifiable from cumulative link measurements.*

Proof: With $\theta = vec(\Phi)$, $\Sigma_X(\theta)$ can be seen to be one-to-one since $(\mathbf{1}, s^{(1)})$ is full rank and

$$\Phi(\mathbf{1}, s^{(1)}) = (s_P, s_{PB})$$

Thus, Φ is identifiable.

Now,

$$(3.3) \qquad E[Y_t^P, Y_t^B] \;=\; A\Gamma(\mathbf{1}, s^{(1)})$$

$$(3.4) \qquad\qquad\qquad =\; A\Phi Diag(\alpha)(\mathbf{1}, s^{(1)})$$

Thus,

$$E[Y_t^P, Y_t^B](\mathbf{1}, s^{(1)})^{-1} = (\alpha_1 w_1, \alpha_2 w_2)$$

where $w_1 = A\Phi_{(\cdot,1)}$ and $w_2 = A\Phi_{(\cdot,2)}$. As before, we have that $w_1 \neq 0$ and $w_2 \neq 0$, since a non-trivial vector with non-negative entries can not be in the null space of A. Since w_1 and w_2 are identifiable, so are α_1 and α_2. Thus, Γ is identifiable.

$\square$

3.4   Estimation Procedure and its Properties

We adopt a pseudo-likelihood approach [21] for estimation purposes. Specifically, we will obtain the estimates through maximizing a function which is not the likelihood of the available data, but rather the likelihood of a normal distribution that has the same mean and covariance as the distribution of the data. There are several computational advantages to using a normal likelihood and in reality the departures from regularizing assumptions tend to have a greater impact than other misspecifications to the likelihood. For a given parametrization of the mean $\eta(\theta)$ and covariance matrix $\Sigma_Y(\theta)$ of a random vector Y_t the normal likelihood is given by:

$$(3.5) \qquad l(\theta) = -\frac{1}{2} tr\left(\Sigma_Y(\theta)^{-1} \sum_{t=1}^{T} (Y_t - \eta(\theta))(Y_t - \eta(\theta))'\right) - \frac{T}{2} \log|\Sigma_Y(\theta)|$$

However, optimizing the above likelihood to get an estimate of θ was found to have quite slow convergence for both second order and EM type algorithms. The intuitive reason for that is that certain parameters appear both in $\Sigma_Y(\theta)$ and $\eta(\theta)$ and that makes the likelihood surface ill-conditioned. The condition number of the information matrix for the normal approximation of the compound model described in Section 3.5 was found to be 7.9×10^{21}. Hence, we propose the following "hybrid" estimator.

Suppose Σ_Y is parametrized as a one-to-one function of the parameter vector θ, $\Sigma_Y(\theta)$. This is true for both of the proposed models with $\theta = (s_P, \psi_0, s)$ for the compound model and $\theta = vec(\Phi)$ for the independent sub-flows model. For estimation, we follow a two-step strategy. In the first step we obtain a consistent

estimate of θ. The covariance only pseudo-likelihood for θ is expressed in terms of the covariance of Y_t. Since $\Sigma_Y(\theta)$ is a one-to-one function, by definition we have that if $\theta_1 \neq \theta_2$, then $\Sigma_Y(\theta_1) \neq \Sigma_Y(\theta_2)$. Further,

$$\hat{\Sigma}_Y = \frac{1}{T} \sum_{t=1}^{T} (Y_t - \overline{Y})(Y_t - \overline{Y})'$$

is a consistent estimate of $\Sigma_Y(\theta)$ under fairly general conditions (specifically temporal independence is not required [43]). Thus,

$$(3.6) \qquad l_Y(\theta) = -\frac{T}{2} tr(\Sigma_Y(\theta)^{-1}\hat{\Sigma}_Y) - \frac{T}{2} \log |\Sigma_Y(\theta)|$$

defines a pseudo-likelihood function. Maximizing the likelihood function $l_Y(\theta)$ in equation (3.6) can be accomplished through the EM algorithm presented in Section 3.4.1. Therefore, at the end of the first step, a consistent estimate $\hat{\theta}$ has been obtained.

The second step proceeds as follows. In both models $E[Y_t] = \tilde{A}\tilde{\mu}$ for some matrix $\tilde{A}$ and vector $\tilde{\mu}$. Specifically, for the compound model we have

$$\tilde{\mu} = \begin{pmatrix} \mu \\ \psi_0 \mu \end{pmatrix}$$

and $\tilde{A} = \overline{A}$. For the independent sub-flows model we have $\tilde{\mu} = vec(\Gamma)$ and

$$\tilde{A} = \begin{pmatrix} A & A \\ s_1 A & s_2 A \end{pmatrix} = \begin{pmatrix} \mathbf{1} & s^{(1)} \end{pmatrix}' \otimes A$$

Further, in both models $\tilde{\mu} = \Theta b$ where Θ is a matrix (or vector) identifiable from $\Sigma_Y(\theta)$ or in other words a function of θ estimated in the first step. The form of Θ and b is given by

$$\Theta = \begin{pmatrix} s \\ \psi_0 s \end{pmatrix}$$

and $b = 1/v_0$ for the compound model and by

$$\tilde{\mu} = \begin{pmatrix} \Gamma_{(\cdot,1)} \\ \Gamma_{(\cdot,2)} \end{pmatrix} = \begin{pmatrix} \Phi_{(\cdot,1)} & 0 \\ 0 & \Phi_{(\cdot,2)} \end{pmatrix} \alpha \equiv \Theta b$$

for the independent sub-flows model.

Next, let the QR decomposition of $\tilde{A}'$ be given by

$$\tilde{A}' = \begin{pmatrix} Q_1 & Q_2 \end{pmatrix} \begin{pmatrix} R' \\ 0 \end{pmatrix}.$$

Now, $\tilde{\mu}$ can be re parametrized as

$$\tilde{\mu} = Q_1 R^{-1} \mu_Y + Q_2 \mu_\perp$$

where $\tilde{A}\tilde{\mu} = \mu_Y$. Note that for $\tilde{A}$ being a $2L \times 2J$ matrix of rank $2L$, then $\mu_\perp \in \mathbb{R}^{2J-2L}$. Further, it is easy to get a consistent estimate of μ_Y; e.g. the sample mean $\hat{\mu}_Y = \frac{\sum_{t=1}^T Y_t}{T}$.

Finally, a consistent estimate of $\mu_\perp$ is obtained by solving

(3.7) $$(\hat{\mu}_\perp, \hat{b}) = \arg \min_{\mu_\perp, b} ||Q_1 R^{-1} \hat{\mu}_Y + Q_2 \mu_\perp - \hat{\Theta} b||_2^2$$

Since $\tilde{\mu}$ has non-negative entries, in practice the above optimization would be done subject to the constraint $Q_1 R^{-1} \hat{\mu}_Y + Q_2 \mu_\perp \geq 0$ which is a quadratic program.

For the purpose of deriving the asymptotic distribution of the estimator that maximizes (3.6) we make explicit the dependence on T:

(3.8) $$l_{Y,T}(\theta) = l_Y(\theta) = -\frac{T}{2} tr(\Sigma_Y(\theta)^{-1} \hat{\Sigma}_Y) - \frac{T}{2} \log |\Sigma_Y(\theta)|$$

We refer to the true value of parameters as θ_0 and to the estimate as $\hat{\theta}_T$.

Proposition III.4. *For X_t, $t = 1, 2, \cdots$ (defined in Section 3.2) being a stationary sequence and whose fourth moments exist the pseudo-likelihood estimator $\hat{\theta}_T$ satisfies*

$$\sqrt{T}(\hat{\theta}_T - \theta_0) \overset{D}{\Rightarrow} J(\theta_0)^{-1} Z_{\theta_0} \sim N(0, J(\theta_0)^{-1} I(\theta_0) J(\theta_0)^{-1})$$

where

$$[I(\theta_0)]_{ij} = E[tr(G_i(Y - EY)(Y - EY)')tr(G_j(Y - EY)(Y - EY)')]$$
(3.9)
$$- tr(G_i \Sigma_Y(\theta_0)) tr(G_j \Sigma_Y(\theta_0))$$

for

(3.10)
$$G_i = \frac{1}{2} \Sigma_Y(\theta_0)^{-1} \frac{\partial \Sigma_Y(\theta)}{\partial \theta_i} \Big|_{\theta = \theta_0} \Sigma_Y(\theta_0)^{-1}$$

and $[J(\theta)]_{ij} = \frac{1}{2} tr(\Sigma_Y(\theta)^{-1} \frac{\partial \Sigma_Y(\theta)}{\partial \theta_i} \Sigma_Y(\theta)^{-1} \frac{\partial \Sigma_Y(\theta)}{\partial \theta_j})$

Corollary III.5. *Under the conditions of Proposition III.4, the hybrid estimator $\hat{\mu}_T$ is also asymptotically normal.*

Remarks:

1. The computational complexity of this estimator is determined by the first step, which involves an EM algorithm. The computational complexity of each EM step is $O(L^4)$ as argued in Section 3.4.1.

2. The pseudo-likelihood (3.6) does not take into account that $E[X_{t,j}^B | X_{t,j}^P] = \psi_j X_{t,j}^P$ and $Var[X_{t,j}^B | X_{t,j}^P] = v_j X_{t,j}^P$ and hence X_t^B given X_t^P is heteroskedastic as opposed to the case of joint normality. An alternative is to assume that X_t^P is normally distributed and X_t^B is normal given X_t^P with mean and variance given by the relation above. In this case the distribution of Y_t does not correspond to any well known distribution and the likelihood of Y_t cannot be written explicitly. In this case the obvious way to obtain estimates would be to use an

EM algorithm where the "full-data" is X_t. Since, this likelihood more closely reflects the compound model it can be expected to be statistically more efficient. However, it was found to have several drawbacks. First, the E-step can no longer be carried out analytically and one has to resort to MCMC methods. Second, the MCMC E-step has to be carried out individually for each time interval which makes it computationally quite expensive. Finally, the gains in statistical efficiency were found to be marginal at best. Thus, we do not pursue that direction in this work.

3. The computational complexity of the first step in the estimation can be reduced by using a method of moments estimator for θ instead of maximum (pseudo-) likelihood estimation. For both models, the elements of the covariance matrix, $\Sigma_Y(\theta)$ can be written as a linear combination of elements of θ. For the compound model this requires an additional estimation step where ψ is estimated and is treated as a known constant in the method of moments step. Thus we get

$$vec(\hat{\Sigma}_Y) = B\theta + \epsilon$$

Hence a consistent estimate of θ can be obtained by minimizing $||vec(\hat{\Sigma}_Y) - B\theta||_2^2$ subject to the constraint $\theta > 0$. This corresponds to solving a quadratic program.

Proof of Proposition: Let Z_T, $T = 2, 3, \cdots$ be a sequence of random vectors, defined as

$$[Z_T]_i = \frac{1}{\sqrt{T}} \left. \frac{\partial l_{Y,T}(\theta)}{\partial \theta_i} \right|_{\theta=\theta_0}$$

and J_T be a sequence of random matrices, defined as

$$[J_T]_{ij} = -\frac{1}{T} \left. \frac{\partial^2 l_{Y,T}(\theta)}{\partial \theta_i \partial \theta_j} \right|_{\theta=\theta_0}$$

We will establish that $Z_T \overset{D}{\Rightarrow} Z_{\theta_0}$ and $J_T \overset{p}{\rightarrow} J(\theta_0)$, for some random vector Z_{θ_0} and constant matrix $J(\theta_0)$.

In the following all functions and derivatives are evaluated at $\theta = \theta_0$. Its easy to show that

$$
\begin{aligned}
\frac{1}{\sqrt{T}}\frac{\partial l_{Y,T}(\theta)}{\partial \theta_i} &= \frac{1}{\sqrt{T}}\left(\frac{\partial l_{Y,T}(\theta)}{\partial \theta_i} - E\frac{\partial l_{Y,T}(\theta)}{\partial \theta_i}\right) \\
&= \frac{1}{2}tr\left[\Sigma_Y(\theta)^{-1}\frac{\partial \Sigma_Y(\theta)}{\partial \theta_i}\Sigma_Y(\theta)^{-1}\left(\sqrt{T}(\hat{\Sigma}_Y - \Sigma_Y(\theta))\right)\right] \\
(3.11) \qquad &= (vec(G_i))'vec(\sqrt{T}(\hat{\Sigma}_Y - \Sigma_Y(\theta)))
\end{aligned}
$$

Define

$$
G = \begin{pmatrix} vec(G_1)' \\ \vdots \\ vec(G_p)' \end{pmatrix}
$$

Thus $Z_T = Gvec(\sqrt{T}(\hat{\Sigma}_Y - \Sigma_Y(\theta)))$.

From CLT $\sqrt{T}(\hat{\Sigma}_Y - \Sigma_Y(\theta))$ converges in distribution to a random matrix with all entries jointly normal distributed. Thus Z_{θ_0} has a multivariate normal distribution. The mean of Z_{θ_0} is 0 and the covariance matrix is given by $I(\theta_0)$.

On the other hand

$$-\frac{1}{T}\frac{\partial^2 l_{Y,T}(\theta)}{\partial\theta_i\partial\theta_j}$$

$$= \frac{1}{2}\frac{\partial}{\partial\theta_j}tr(-\Sigma_Y(\theta)^{-1}\frac{\partial\Sigma_Y(\theta)}{\partial\theta_i}\Sigma_Y(\theta)^{-1}\hat{\Sigma}_Y) + \frac{1}{2}\frac{\partial}{\partial\theta_j}tr(\Sigma_Y(\theta)^{-1}\frac{\partial\Sigma_Y(\theta)}{\partial\theta_i})$$

$$= -\frac{1}{2}tr(-\Sigma_Y(\theta)^{-1}\frac{\partial\Sigma_Y(\theta)}{\partial\theta_j}\Sigma_Y(\theta)^{-1}\frac{\partial\Sigma_Y(\theta)}{\partial\theta_i}\Sigma_Y(\theta)^{-1}\hat{\Sigma}_Y)$$

$$-\frac{1}{2}tr(\Sigma_Y(\theta)^{-1}\frac{\partial^2\Sigma_Y(\theta)}{\partial\theta_i\partial\theta_j}\Sigma_Y(\theta)^{-1}\hat{\Sigma}_Y)$$

$$-\frac{1}{2}tr(-\Sigma_Y(\theta)^{-1}\frac{\partial\Sigma_Y(\theta)}{\partial\theta_i}\Sigma_Y(\theta)^{-1}\frac{\partial\Sigma_Y(\theta)}{\partial\theta_j}\Sigma_Y(\theta)^{-1}\hat{\Sigma}_Y)$$

$$+\frac{1}{2}tr(-\Sigma_Y(\theta)^{-1}\frac{\partial\Sigma_Y(\theta)}{\partial\theta_j}\Sigma_Y(\theta)^{-1}\frac{\partial\Sigma_Y(\theta)}{\partial\theta_i}) + \frac{1}{2}tr(\Sigma_Y(\theta)^{-1}\frac{\partial^2\Sigma_Y(\theta)}{\partial\theta_i\partial\theta_j})$$

$$\xrightarrow{p} \frac{1}{2}tr(\Sigma_Y(\theta)^{-1}\frac{\partial\Sigma_Y(\theta)}{\partial\theta_i}\Sigma_Y(\theta)^{-1}\frac{\partial\Sigma_Y(\theta)}{\partial\theta_j}) = [J(\theta)]_{ij}$$

Thus $\sqrt{T}(\hat{\theta}_T - \theta_0) \xrightarrow{D} J(\theta_0)^{-1}Z_{\theta_0} \sim N(0, J(\theta_0)^{-1}I(\theta_0)J(\theta_0)^{-1})$.

$\square$

Clearly consistent estimates of G_i, $i = 1, \cdots, p$ can be obtained by replacing θ_0 in (3.10) by a consistent estimate like $\hat{\theta}_T$. Now $I(\theta_0)$ can be consistently estimated by replacing G_i and G_j by their consistent estimates and expectations by their empirical means in (3.9). Also $J(\theta_0)$ is consistently estimated by $J(\hat{\theta}_T)$.

Proof of the Corollary: For the asymptotic distribution of the hybrid estimator, note that neglecting the positivity constraints, the objective function 3.7 is maximized for

$$\begin{pmatrix} \hat{\mu}_\perp \\ \hat{b} \end{pmatrix} = -(\hat{P}'\hat{P})^{-1}\hat{P}'Q_1R^{-1}\hat{\mu}_Y$$

where $\hat{P} = (Q_2, -\hat{\Theta})$. Now

$$\hat{P}'Q_1R^{-1} = \begin{pmatrix} Q_2' \\ -\hat{\Theta}' \end{pmatrix} Q_1R^{-1} = \begin{pmatrix} 0 \\ -\hat{\Theta}'Q_1R^{-1} \end{pmatrix}$$

Further,

$$\hat{P}'\hat{P} = \begin{pmatrix} Q_2'Q_2 & -Q_2'\hat{\Theta} \\ -\hat{\Theta}'Q_2 & \hat{\Theta}'\hat{\Theta} \end{pmatrix} = \begin{pmatrix} I & -Q_2'\hat{\Theta} \\ -\hat{\Theta}'Q_2 & \hat{\Theta}'\hat{\Theta} \end{pmatrix}$$

Hence,

$$[\hat{P}'\hat{P}]_{12}^{-1} = -Q_2'\hat{\Theta}(\hat{\Theta}'Q_2Q_2'\hat{\Theta} - \hat{\Theta}'\hat{\Theta})^{-1}$$

Thus

$$\hat{\mu}_\perp = Q_2'\hat{\Theta}(\hat{\Theta}'Q_2Q_2'\hat{\Theta} - \hat{\Theta}'\hat{\Theta})^{-1}\hat{\Theta}'Q_1R^{-1}\hat{\mu}_Y$$

So, finally

(3.12)
$$\hat{\mu} = \left(I + Q_2Q_2'\hat{\Theta}(\hat{\Theta}'Q_2Q_2'\hat{\Theta} - \hat{\Theta}'\hat{\Theta})^{-1}\hat{\Theta}'\right) Q_1R^{-1}\hat{\mu}_Y$$

Or making explicit the dependence on T in (3.12):

$$\hat{\mu}_T = \left(I + Q_2Q_2'\hat{\Theta}_T(\hat{\Theta}_T'Q_2Q_2'\hat{\Theta}_T - \hat{\Theta}_T'\hat{\Theta}_T)^{-1}\hat{\Theta}_T'\right) Q_1R^{-1}\hat{\mu}_{Y,T} \equiv M(\hat{\theta}_T)\hat{\mu}_{Y,T}$$

Now from the proposition $\sqrt{T}(\hat{\theta}_T - \theta) \overset{D}{\Rightarrow} Z_\theta$, a mean 0 normal random variable. Similarly, $\sqrt{T}(\hat{\mu}_{Y,T} - \mu_Y) \overset{D}{\Rightarrow} Z_Y$, another mean 0 normal random variable. Thus a simple application of delta method suggests an asymptotic distribution given by

$$\sqrt{T}(\hat{\mu}_T - \mu) \overset{D}{\Rightarrow} \Psi(\theta, \mu_Y) \begin{pmatrix} Z_\theta \\ Z_Y \end{pmatrix}$$

where Z_θ, Z_Y are jointly normal distributed with

$$
\begin{aligned}
\Sigma_Z \; &\equiv \; E\left[\begin{pmatrix} Z_\theta \\ Z_Y \end{pmatrix} (\; Z_\theta' \;\; Z_Y \;) \right] \\
&= \; E\left[\begin{pmatrix} Gvec((Y-EY)(Y-EY)') \\ Y \end{pmatrix} (\; vec((Y-EY)(Y-EY)')'G' \;\; Y' \;) \right] \\
&\quad -E \begin{pmatrix} Gvec((Y-EY)(Y-EY)') \\ Y \end{pmatrix} E(\; vec((Y-EY)(Y-EY)')'G' \;\; Y' \;)
\end{aligned}
$$

and

$$
\Psi(\theta, \mu_Y) = \left(\frac{\partial M(\theta)}{\partial \theta_1} \mu_Y, \cdots, \frac{\partial M(\theta)}{\partial \theta_n} \mu_Y, M(\theta) \right).
$$

The partial derivatives in the above expression can be written more explicitly as

$$
\frac{\partial M(\theta)}{\partial \theta_i} = M_{1i}(\theta) + M_{2i}(\theta) + M_{3i}(\theta)
$$

where for

$$
\tilde{\Theta} = (\Theta' Q_2 Q_2' \Theta - \Theta'\Theta)^{-1}
$$

$$
M_{1i}(\theta) = Q_2 Q_2' \frac{\partial \Theta}{\partial \theta_i} \tilde{\Theta}\Theta' Q_1 R^{-1}
$$

$$
M_{2i}(\theta) = Q_2 Q_2' \Theta \tilde{\Theta} \frac{\partial (\Theta' Q_2 Q_2' \Theta - \Theta'\Theta)}{\partial \theta_i} \tilde{\Theta}\Theta' Q_1 R^{-1}
$$

$$
M_{3i}(\theta) = Q_2 Q_2' \Theta \tilde{\Theta} \frac{\partial \Theta'}{\partial \theta_i} Q_1 R^{-1}
$$

□

3.4.1 EM Algorithm for Covariance Only Pseudo Likelihood

A very simple EM algorithm can be derived to maximize the pseudo-likelihood in (3.6). Assume $\tilde{Y}_t = Y_t - \overline{Y}$. Then

$$
l_Y(\theta) = -\frac{T}{2} tr\left(\Sigma_Y(\theta)^{-1}\left(\sum_{t=1}^{T} \tilde{Y}_t \tilde{Y}_t'\right)\right) - \frac{T}{2} \log |\Sigma_Y(\theta)|
$$

The above would be the true likelihood of $\tilde{Y}_1, \cdots, \tilde{Y}_T$ if

$$
\tilde{X}_t = \begin{pmatrix} \tilde{X}_t^P \\ \tilde{X}_t^B \end{pmatrix}
$$

were distributed i.i.d. $N(0, \Sigma_X(\theta))$ and $\tilde{Y}_t = \overline{A}\tilde{X}_t$. We use this model to derive the EM algorithm. Let $l_{\tilde{X}}(\theta)$ be the likelihood function based on $\tilde{X}_1, \cdots, \tilde{X}_T$. Thus

$$
\begin{aligned}
l_{\tilde{X}}(\theta) &= -\frac{1}{2}tr\Sigma_P^{-1}(\sum_{t=1}^{T}\tilde{X}_t^P(\tilde{X}_t^P)') - \frac{T}{2}\log|\Sigma_P| \\
&\quad -\frac{1}{2}\sum_j\left((\sum_{t=1}^{T}\frac{(\tilde{X}_{t,j}^B - \psi_j\tilde{X}_{t,j}^P)^2}{s_j}) + T\log s_j\right) \\
&= -\frac{1}{2}tr\Sigma_P^{-1}(\sum_{t=1}^{T}\tilde{X}_t^P(\tilde{X}_t^P)') - \frac{T}{2}\log|\Sigma_P| \\
&\quad -\frac{1}{2}\sum_j\left(\frac{\sum_{t=1}^{T}(\tilde{X}_{t,j}^B)^2 + \psi_j^2\sum_{t=1}^{T}(\tilde{X}_{t,j}^P)^2 - 2\psi_j\sum_{t=1}^{T}\tilde{X}_{t,j}^B\tilde{X}_{t,j}^P}{s_j} + T\log s_j\right)
\end{aligned}
$$

Assume that at the kth E-step the estimated parameter is $\theta^{(k)}$. Let

$$
\begin{aligned}
Cov(\tilde{X}_t|\tilde{Y}_t, \theta^{(k)}) &\equiv \begin{pmatrix} R_P^{(k)} & R_{PB}^{(k)} \\ R_{BP}^{(k)} & R_B^{(k)} \end{pmatrix} \\
\tag{3.13} &= \Sigma_X(\theta^{(k)}) - \Sigma_X(\theta^{(k)})\overline{A}'\left(\overline{A}\Sigma_X(\theta^{(k)})\overline{A}'\right)^{-1}\overline{A}\Sigma_X(\theta^{(k)})
\end{aligned}
$$

and

$$
\tag{3.14} E[\tilde{X}_t|\tilde{Y}_t, \theta^{(k)}] \equiv \begin{pmatrix} m_t^{(k)} \\ b_t^{(k)} \end{pmatrix} = \Sigma_X(\theta^{(k)})\overline{A}'\left(\overline{A}\Sigma_X(\theta^{(k)})\overline{A}'\right)^{-1}\tilde{Y}_t
$$

Now

$$
\tag{3.15} E[\sum_{t=1}^{T}\tilde{X}_t^P(\tilde{X}_t^P)'|\tilde{Y}_1, \cdots, \tilde{Y}_T, \theta^{(k)}] = TR_P^{(k)} + \sum_{t=1}^{T}m_t^{(k)}(m_t^{(t)})'
$$

Define

$$
\tag{3.16} a_{B,j}^{(k)} \equiv E[\sum_{t=1}^{T}(\tilde{X}_{t,j}^B)^2|\tilde{Y}_1, \cdots, \tilde{Y}_T, \theta^{(k)}] = T[R_B^{(k)}]_{jj} + \sum_{t=1}^{T}(b_{t,j}^{(k)})^2
$$

(3.17) $a_{P,j}^{(k)} \equiv E[\sum_{t=1}^{T}(\tilde{X}_{t,j}^{P})^2|\tilde{Y}_1,\cdots,\tilde{Y}_T,\theta^{(k)}] = T[R_P^{(k)}]_{jj} + \sum_{t=1}^{T}(m_{t,j}^{(k)})^2$

and

(3.18) $a_{PB,j}^{(k)} \equiv E[\sum_{t=1}^{T}\tilde{X}_{t,j}^{P}\tilde{X}_{t,j}^{B}|\tilde{Y}_1,\cdots,\tilde{Y}_T,\theta^{(k)}] = T[R_{PB}^{(k)}]_{jj} + \sum_{t=1}^{T}m_{t,j}^{(k)}b_{t,j}^{(k)}$

Using the above, we get the expectation step.

E-Step

$$
\begin{aligned}
Q(\theta,\theta^{(k)}) &\equiv E[l_{\tilde{X}}(\theta)|\tilde{Y}_1,\cdots,\tilde{Y}_T,\theta^{(k)}] \\
&= -\frac{1}{2}tr\Sigma_P^{-1}(TR_P^{(k)} + \sum_{t=1}^{T}m_t^{(k)}(m_t^{(t)})') - \frac{T}{2}\log|\Sigma_P| \\
&\quad -\frac{1}{2}\sum_j \left(\frac{a_{B,j}^{(k)} + \psi_j^2 a_{P,j}^{(k)} - 2\psi_j a_{PB,j}^{(k)}}{s_j} + T\log s_j \right)
\end{aligned}
$$
(3.19)

The M-step involves maximization of $Q(\theta,\theta^{(k)})$ over θ and is straightforward from the following observations. The first and second term in the last expression just involve Σ_P. The maximum likelihood estimate of Σ_P subject to the diagonal constraint is given simply by replacing the off-diagonal elements in the unconstrained MLE with 0.

Let $\mathcal{B}(\cdot)$ be the function which replaces the off-diagonal elements of a matrix with zeros. The kth stage M step then gives the following parameter estimates.

M-Step

(3.20) $$\Sigma_P^{(k+1)} = \mathcal{B}\left(R_P^{(k)} + \frac{1}{T}\sum_{t=1}^{T}m_t^{(k)}(m_t^{(k)})' \right)$$

(3.21) $$\psi_j^{(k+1)} = \psi_0^{(k+1)} = \frac{\sum_j a_{PB,j}^{(k)}/s_j^{(k)}}{\sum_j a_{P,j}^{(k)}/s_j^{(k)}}$$

and

(3.22) $$s_j^{(k+1)} = \frac{a_{B,j}^{(k)} - 2\psi_j^{(k+1)}a_{PB,j}^{(k)} + (\psi_j^{(k+1)})^2 a_{P,j}^{(k)}}{T}$$

Computational Complexity

The computational complexity of each EM step can be obtained as follows. Assume that the number of flows, J, is of order L^2, where A is $L \times J$. The matrix inversions in (3.13-3.14) involve $2L \times 2L$ matrices and hence have complexity $O(L^3)$. Computing $\overline{A}\Sigma_X$ involves multiplying a $2L \times 2J$ matrix with a $2J \times 2J$ matrix and would have complexity $O(LJ^2)$ if done naively. However, if sparsity of Σ_X is exploited then the complexity reduces to $O(LJ) = O(L^3)$. On the other hand, computing $\overline{A}\Sigma_X\overline{A}'$ from $\overline{A}\Sigma_X$ involves multiplying a $2L \times 2J$ matrix with a $2J \times 2L$ matrix, neither of which is necessarily sparse. The complexity of this operation is $O(L^2J) = O(L^4)$. Note that we never need to multiply two $L \times L$ matrices. Thus, the overall complexity of each iteration is $O(L^4)$. Note that while (3.14) is expressed in terms of individual $\tilde{Y}_t$, we only need the following "sufficient" statistic

$$\sum_{t=1}^{T} \tilde{Y}_t \tilde{Y}_t'$$

for evaluation of (3.15-3.18). This would involve a one-time cost of $O(L^2T)$.

3.5   Performance Assessment

The data set and simulation setups used in our numerical study are described next.

Data were obtained from a complete packet header trace of a high capacity link [8]. We split the data into bidirectional flows between sub-networks using the first 8 bytes of the IP-address to identify the corresponding sub-network. We aggregate flow volumes to bin size of 5 minutes. The total duration under consideration is 12.5 hours. Thus, we have data on packet and byte volumes of 55 flow pairs (110 flows) in each of 150 time intervals. The mean byte volume of each of these 110 flows is plotted versus the mean packet volume in Figure 3.3.
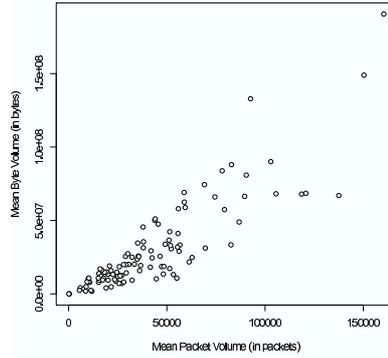
Figure 3.3: Mean byte volume versus mean packet volume for flows in Tokyo network trace data.

To generate data from the compound model we simulate the packet volumes as independently Gamma distributed with means and variances equal to the corresponding parameters in the Tokyo trace data set described above. For each time-interval and flow, given the packet volume, the byte volume is generated as normally distributed with mean and variance proportional to the packet volume. The proportionality constants are the mean packet size and variance in packet size. Mean packet size is estimated from the Tokyo-trace data set over all flows. Variance of the packet size distribution is calculated from the mean by assuming that the packet size distribution is supported entirely on 40 and 1500 bytes.

To simulate from the Independent Sub-Flow model, we generate two sub-flow (packet) volumes for each flow in each time-interval. The first sub-flow corresponds to a packet size of 40 bytes and we use Gamma distributions with common scale parameter across all flows and randomly generated shape parameter. Similarly the second sub-flow corresponds to a packet size of 1500 bytes and we use Gamma distributions with common scale parameter across all flows and randomly generated shape parameters. Finally the packet and byte volumes of the flows are generated as appropriate linear combinations of the sub-flow volumes.

Finally, we also look at data generated from the Independent Sub-Flow model
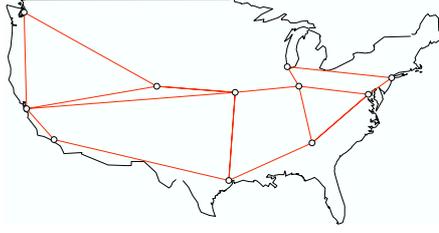
Figure 3.4: Abilene Topology used for Numerical Study

above, with the additional constraint that the scale parameter of all sub-flows is identical. In this case the means of packet volumes are proportional to variances of packet volumes over all flows. Thus, estimation based on classical tomography relation [6] would be consistent. We refer to this as the classical data generation method.

The Abilene network topology (Figure 3.4) is used in our experimental setup. It consists of 11 nodes and $16 \times 2 = 32$ directed edges between pairs of nodes (bidirectional links). Flows exist between all pairs of nodes resulting in a total of $11 \times 10 = 110$ flows. We assume that these flows are routed through minimum distance paths. Further we assume that cumulative flow volumes (SNMP data) are available from all the edges.

The key findings from the numerical study are discussed next. In the case of simulated data, 200 replications of each scenario were run to obtain the mean and standard deviation of estimates.

1. In Figures 3.5, 3.6 and 3.7 the results of estimating the mean packet volumes are shown using the Compound (top row), Independent Sub-flows (middle row) and classical tomography (bottom row), when the data generation mechanism corresponds to the Compound, the Independent Sub-flows and the classical to-mography model, respectively. It can be seen that when the model is correctly specified, the resulting estimates exhibit no discernible bias. Further, when data

are generated from the Independent Sub-flows model, classical tomography performs well (see Figure 3.6), while when the data are generated from the classical tomography mechanism, both the Compound model and the Independent Sub-flows model estimate the means well (see Figure 3.7). Finally, as expected, with increasing T the variance of the estimates reduces.

2. In Figures 3.8 and 3.9 the estimates obtained at the end of the first stage of the "hybrid" procedure (i.e. $\hat{\theta}$) are shown, when the generative model is specified as Compound and Independent Sub-flows, respectively. It can be seen that the results exhibit no discernible bias in the case of correct specification (left panels in Figure 3.8 and right panels in 3.9). On the other hand, estimates from the Independent Sub-flows model exhibit a strong systematic bias for data obtained from the Compound model (right panels in Figure 3.8), while those obtained from the Independent Sub-flows model do not when estimated by the Compound model (left panels in Figure 3.9). Figure 3.10 shows $\hat{\theta}$ when data is generated from classical model and estimation is performed with the Compound and Independent Sub-flows models. The Independent Sub-flows model performs well in general, while the Compound model estimates adequately only the s_P parameters. Finally, the variance of the estimates decreases as the sample size T increases (results not shown).

3. Table 3.1 shows the median (over flows) of relative mean squared error for various scenarios of data generation and estimation. The median is used in order to avoid the results from getting overwhelmed by lighter flows, which have large relative MSE. Relative MSE for a parameter is defined as follows. Let θ_0 be the true value of the parameter and let $\hat{\theta}_i$ be the estimate from the

Table 3.1: Median Relative MSE for various estimation and generative models

Estimation		Generative Model					
		$T = 150$			$T = 500$		
Model	Parameter	Compound	ISF	Classical	Compound	ISF	Classical
Compound	μ	0.241	0.253	0.234	0.093	0.144	0.089
	s_P	0.565	0.329	0.32	0.234	0.109	0.106
	s_{PB}	0.565	0.377	0.453	0.234	0.131	0.217
	s_B	0.565	0.307	0.354	0.234	0.107	0.293
ISF	μ	0.297	0.186	0.165	0.245	0.07	0.059
	s_P	0.47	0.307	0.235	0.353	0.107	0.07
	s_{PB}	0.488	0.425	0.38	0.411	0.134	0.165
	s_B	0.61	0.429	0.408	0.3	0.136	0.175
Classical	μ	0.7	0.293	0.276	0.589	0.105	0.091

ith replication out of a total of r replications. Then, the relative MSE is equal to $\sum_{i=1}^{r}(\hat{\theta}_i/\theta_0 - 1)^2$.

4. Figures 3.11 and 3.12 display the estimated versus true values for the Tokyo trace data for the three estimation techniques. For (s_P, s_{PB}, s_B) the Independent sub-flows model clearly does better than the Compound model. For the final estimate, $\hat{\mu}$, both the Compound model and the Independent Sub-flows model suffer from a single outlier, while the classical tomography estimates have a lot of estimates equal to 0.

5. The outlier in Figures 3.12 (a) and (b) corresponds to the flow in Figure 3.1 (c). This is clearly an exceptional flow. We substitute this flow by the flow in Figure 3.14 which is constructed through averaging the packet and byte volumes over all other flows in each time interval. Figure 3.13 shows the estimated versus true μ for the three estimation methods with this replacement. In this case the performance of the estimates based on Compound model and Independent Sub-flows model clearly outperform those based on classical tomography.
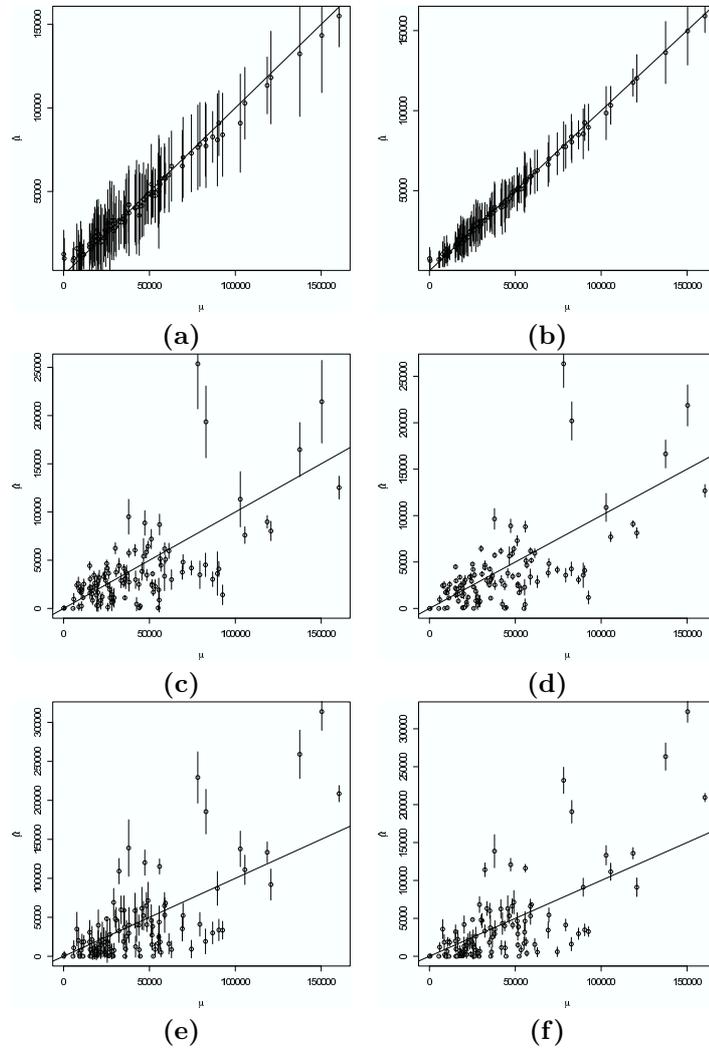
Figure 3.5: Estimated (with +/- s.d error bars) versus the true parameters for data simulated from Compound Model and estimation under Compound (Top), Independent Sub-Flow (Middle) and Classical Tomography (Bottom) model with $T = 150$ (left) and $T = 500$ (right).
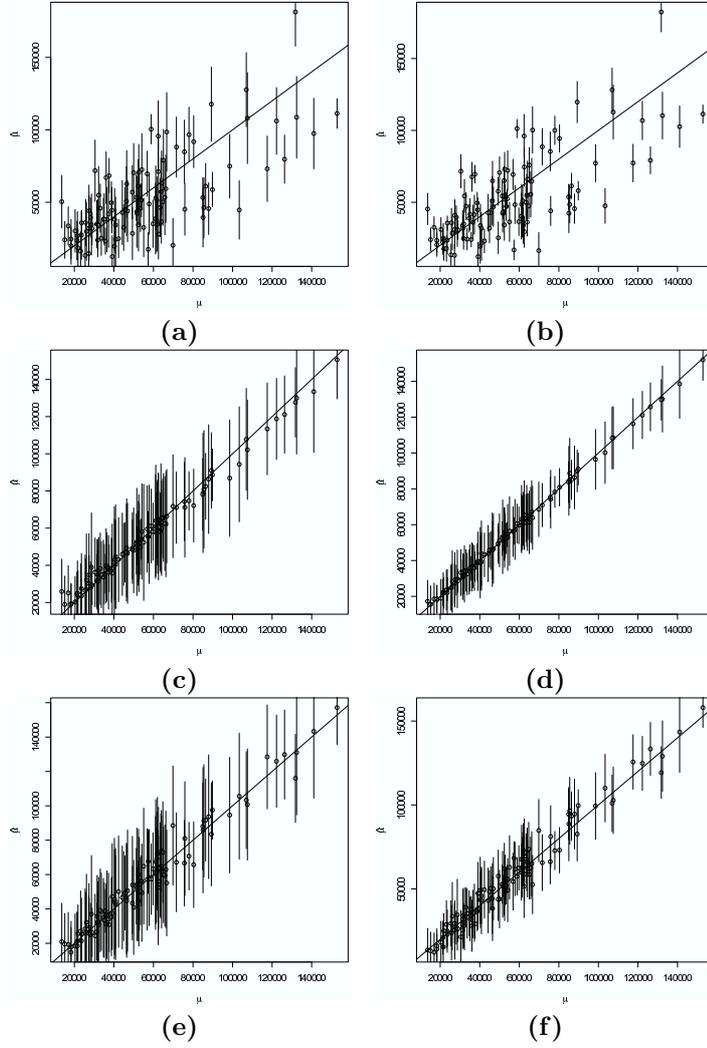
Figure 3.6: Estimated (with +/- s.d error bars) versus the true parameters for data simulated from Independent Sub-Flow Model and estimation under Compound (Top), Independent Sub-Flow (Middle) and Classical Tomography (Bottom) model with $T = 150$ (left) and $T = 500$ (right).
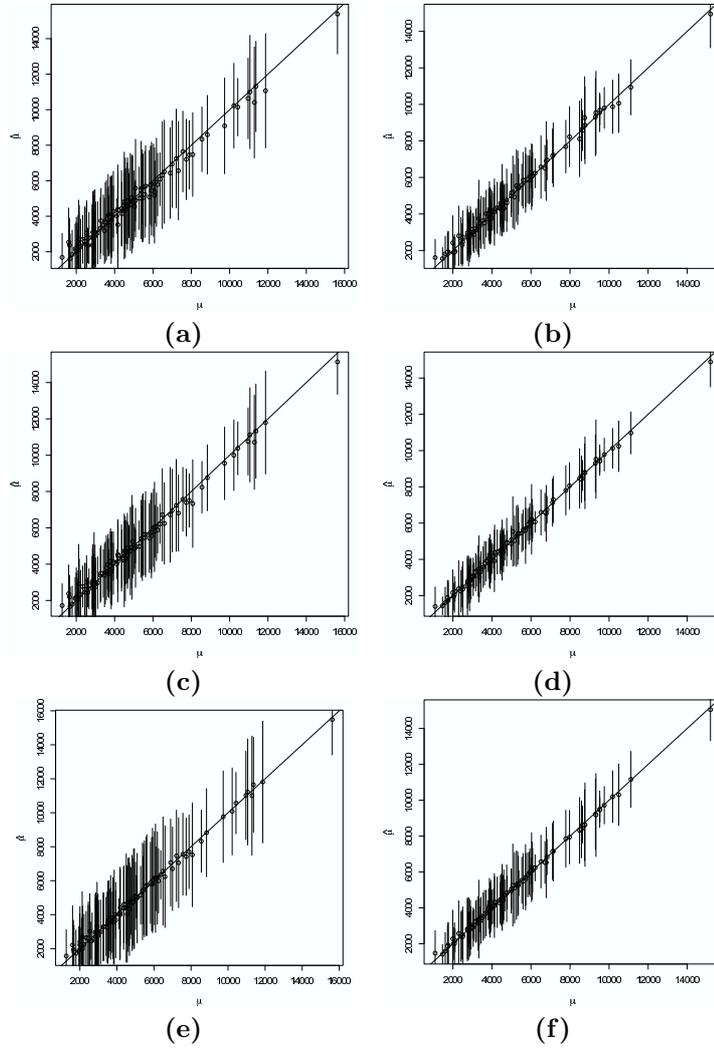
Figure 3.7: Estimated (with +/- s.d error bars) versus the true parameters for data simulated from Classical data generation Model and estimation under Compound (Top), Independent Sub-Flow (Middle) and Classical Tomography (Bottom) model with $T = 150$ (left) and $T = 500$ (right).
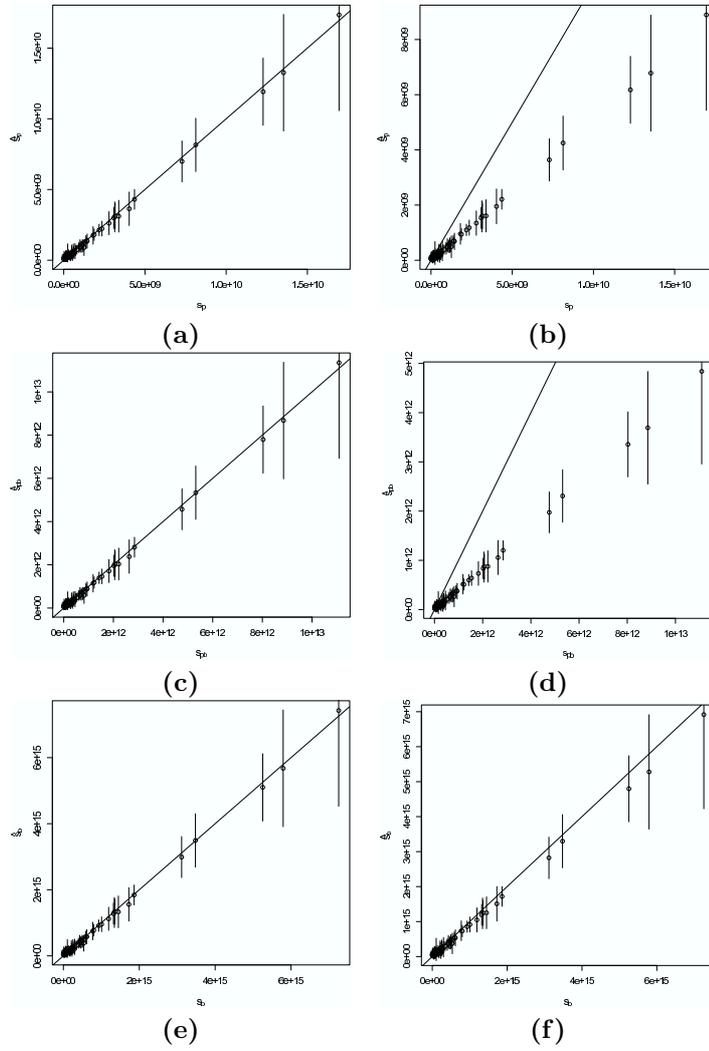
Figure 3.8: Estimated (with +/- s.d error bars) versus the true parameters for data simulated from Compound Model and estimation under Compound model (left) and Independent Sub-Flow model(right) for $T = 150$.
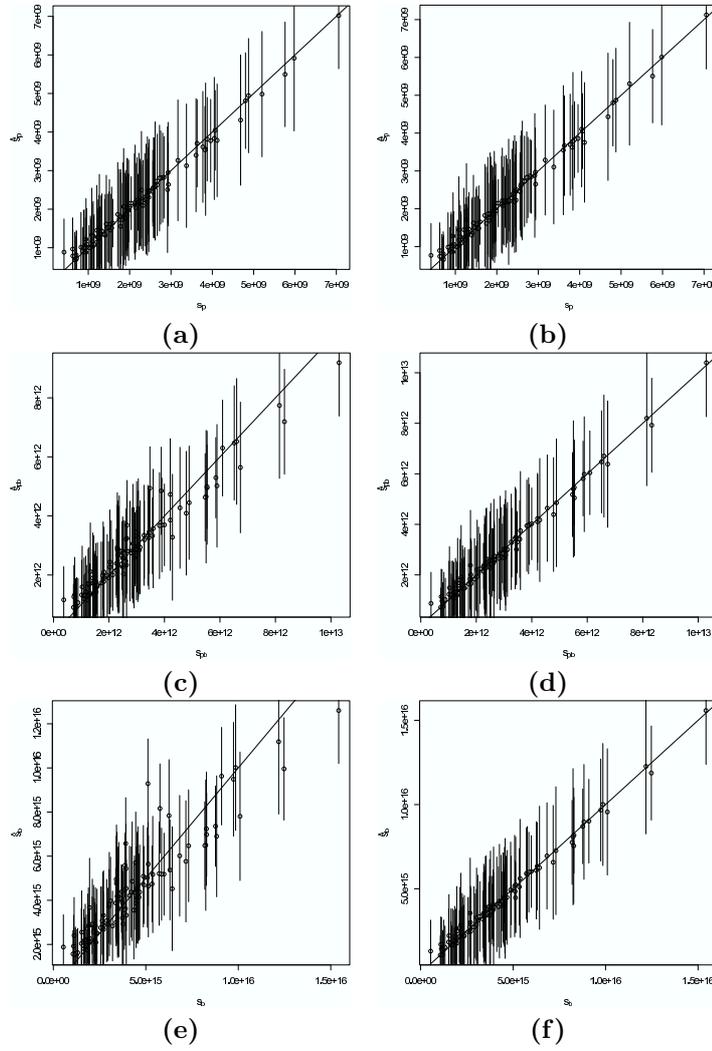
Figure 3.9: Estimated (with +/- s.d error bars) versus the true parameters for data simulated from Independent Sub-Flow model and estimation under Compound Model (left) and Independent Sub-Flow model (right) with $T = 150$.
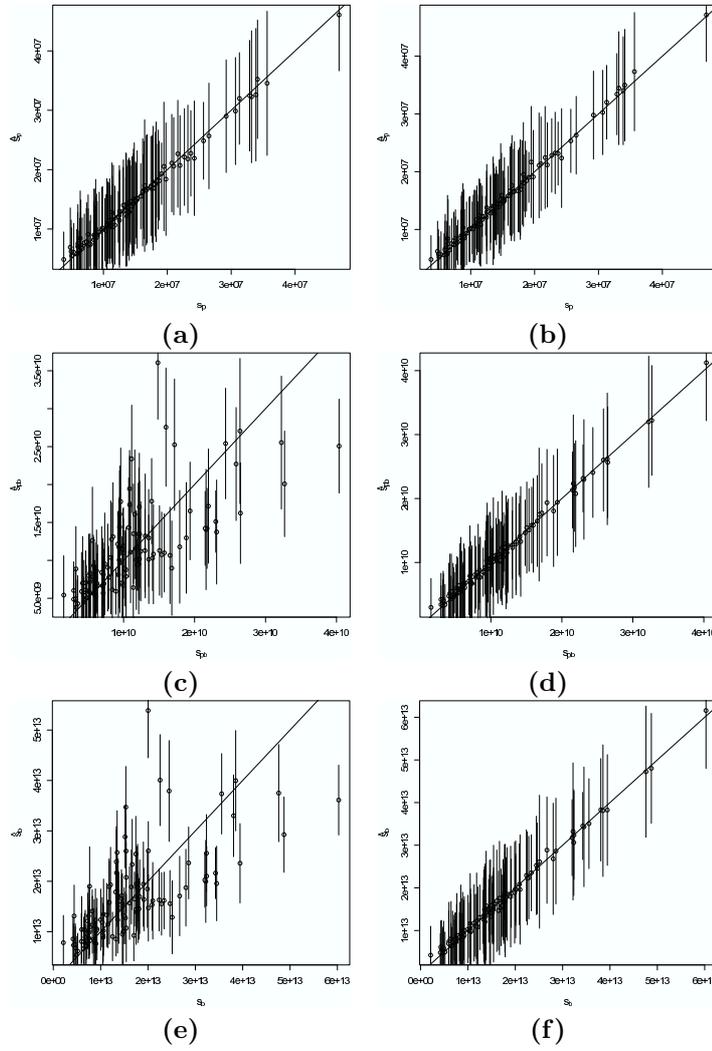
Figure 3.10: Estimated (with +/- s.d error bars) versus the true parameters for data simulated from Classical data generation model and estimation under Compound Model (left) and Independent Sub Flow Model (right) with $T = 150$.
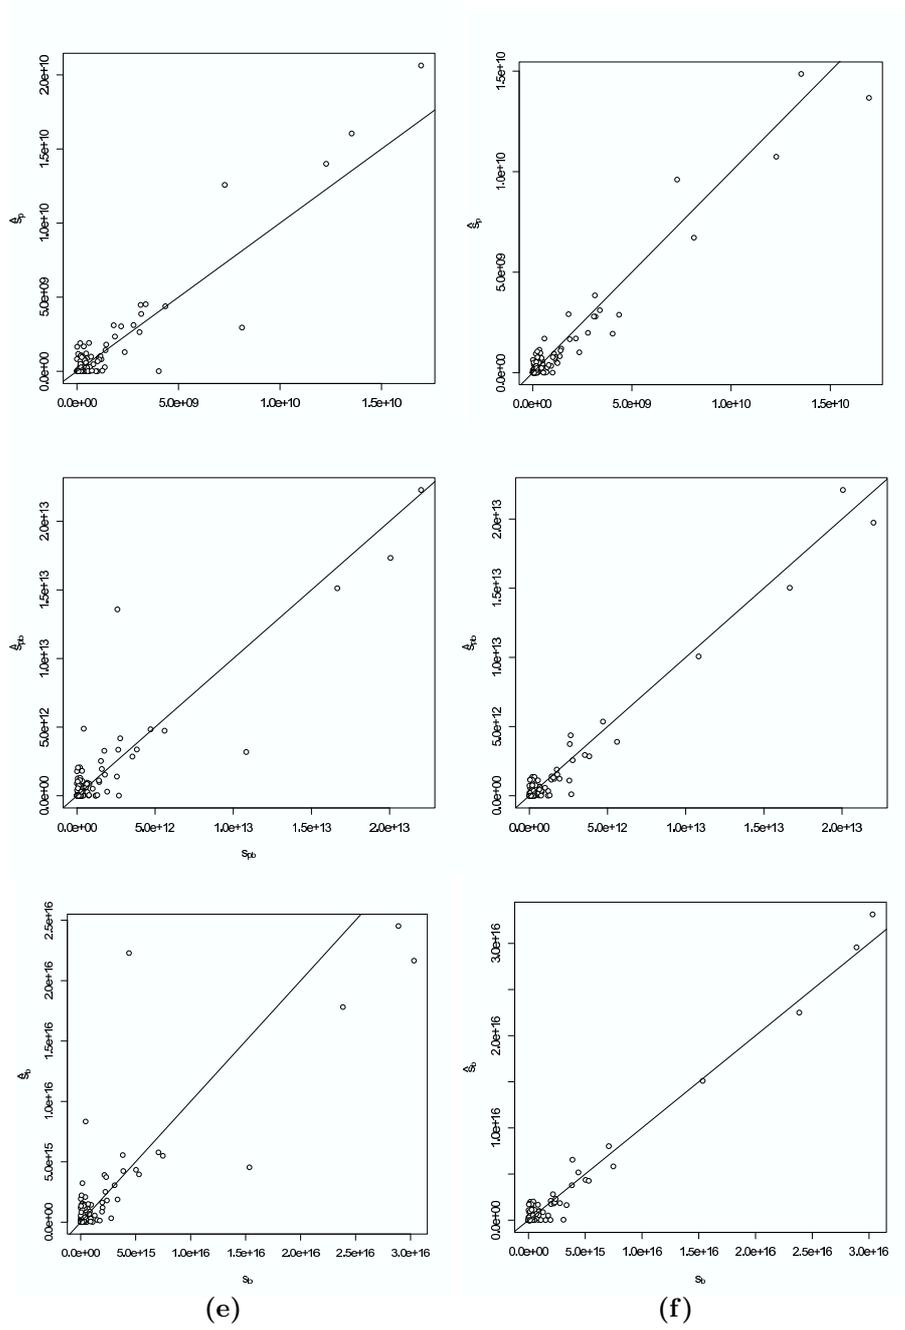
Figure 3.11: Estimated versus the true parameters for Tokyo Data ($T = 150$) assuming compound model (left) and Independent Sub-Flow model (right).
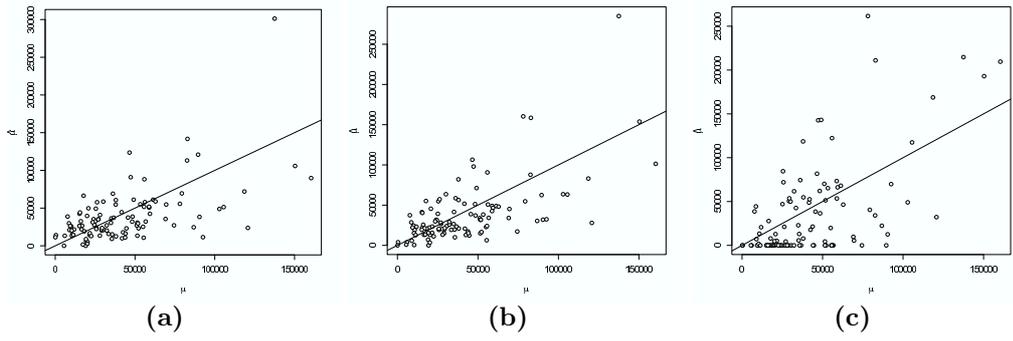
Figure 3.12: Estimated versus the true means for Tokyo Data ($T = 150$) assuming compound model (a), Independent Sub-Flow model (b) and Classical Tomography (c).
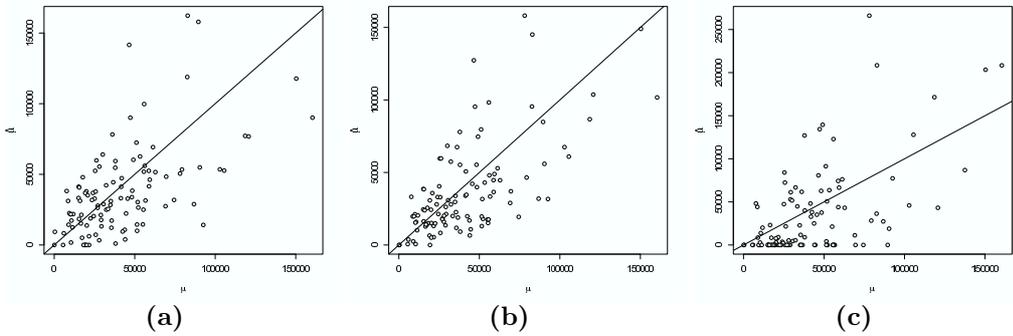

Figure 3.13: Estimated versus the true means for Tokyo Data ($T = 150$) after replacing the outlier flow, assuming compound model (a), Independent Sub-Flow model (b) and Classical Tomography (c).
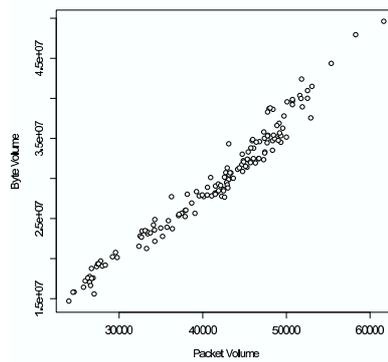

Figure 3.14: Substitute flow volumes for the outlier flow

3.6 Packet Size Tomography

The packet size distribution of a flow is a useful quantity for network monitoring purposes and is indicative of the traffic composition [44]. Joint modeling of packet and byte volumes allow us to estimate parameters of the packet size distribution from cumulative measurements, as well. This is most easily accomplished through the compound model and that is the focus on in this section.

We start by removing the constraint of common packet size means; i.e. $\psi_j = \psi_0$. The objective here is to estimate ψ, the vector of mean packet sizes of all flows. Recall that if Σ_P is constrained to be diagonal as described in Section 3.3.1, then Σ_X is identifiable from Y_t observations. This in turn means that $\theta = (\Sigma_P, \psi, s)$ is identifiable. Mean packet volumes, μ, are not identifiable and are in fact "confounded" with v. Thus, we use the parametrization (3.2). With this parametrization, $l_Y(\theta)$ the "covariance only pseudo-likelihood" (3.6), is well-behaved.

As before the pseudo likelihood estimator, $\hat{\theta}$ maximizes $l_Y(\theta)$. An EM algorithm very similar to that used for the hybrid estimator and given in the Appendix can be used for the optimization. The only difference is that equation (3.21) is replaced by

$$(3.23) \qquad \psi_j^{(k+1)} = \frac{a_{PB,j}^{(k)}}{a_{P,j}^{(k)}}$$

Remark: The computational complexity of each EM step is the same as that for the Hybrid Estimator of Section 3.4

3.6.1 Numerical Study

First, we consider the performance of our estimates under simulated data. The data are generated as described in Section 3.5 for the Compound model with the exception that the mean and variance of packet size distribution is calculated separately for each flow and data generated correspondingly. A sample size of $T = 500$

is considered.

Figure 3.15 shows the estimated versus the true values of s_p, s_{pb}, s_b and ψ. Note that the "natural" parameters of the covariance matrix, i.e. s_p, s_{pb}, s_b are well estimated. However, certain ψ_j have large MSE of estimation. The reason for this is as follows. Estimating ψ_j is similar to estimating the regression coefficient with s_{pj} being the variance of the corresponding covariate. As in any regression problem, if the covariate variances span a big range of values, the coefficients corresponding to small values of s_{pj} are not well estimated. This issue is demonstrated more clearly in Figure 3.16. The plot on the left panel shows the MSE from the above simulation versus s_p (both on a log scale), while the plot on the right panel shows the asymptotic variance (as described in the following) versus s_p (again on log scales). The asymptotic variance is calculated from the Fisher Information matrix corresponding to the covariance only likelihood (3.6) when ψ alone is unknown, evaluated at the true value of θ. Both figures show that a large variance for estimates of ψ_j is observed for small values of s_{pj}. The differences between the two plots are expected due to departures from normality in the data. In reality, the interest is primarily in estimating properties of heavy flows which usually correspond to large values of packet volume variance, s_{pj}. Since, s_p is itself well estimated, reliable estimates of ψ_j can be provided for the most interesting flows.

Figure 3.17 shows the estimated versus true values of mean packet size, ψ_j, for the Tokyo trace data for heavy flows only. Here heavy flows are defined to be the top 40% flows in terms of *estimated* packet volume variance, $\hat{s}_p$. Naturally, the data would have some departures from the compound model that would impact the performance of estimates. It is likely that more highly aggregated data would follow the compound model better and would lead to better estimation.
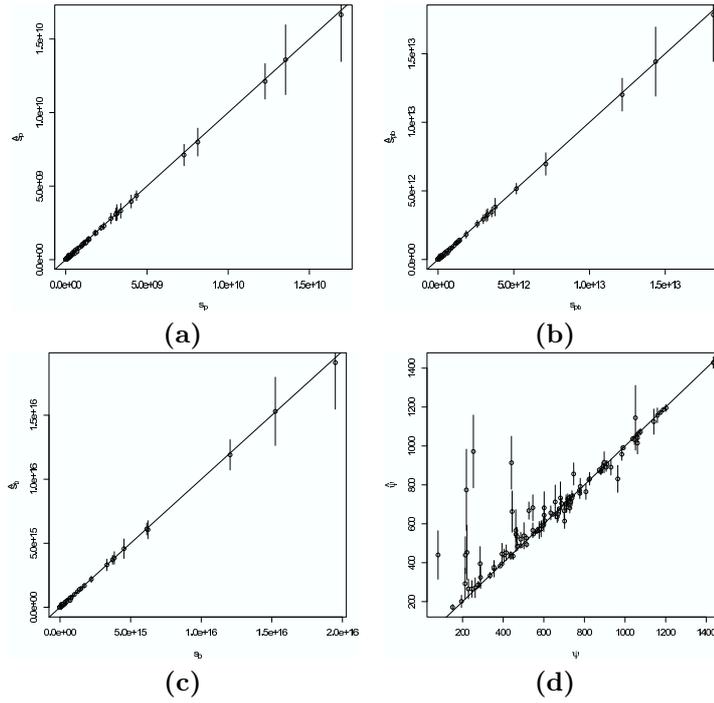
Figure 3.15: Estimated (with +/- s.d error bars) versus the true parameters for data simulated and estimated under Compound Model with $T = 500$.
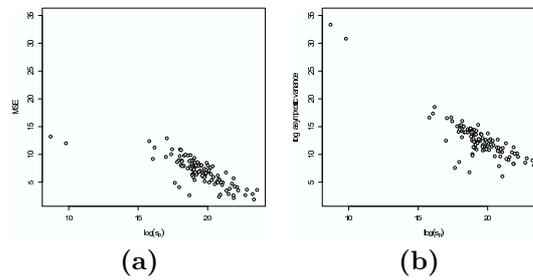

Figure 3.16: Dependence of (a) MSE (simulation) and (b) asymptotic variance (normal approximation) on packet volumes variance
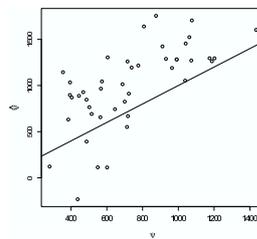

Figure 3.17: Estimated versus true ψ for heavy flows in Tokyo trace data

3.7   Discussion

The use of packet and byte information for traffic volume measurement along with structural modeling of their joint distribution opens up several options for more detailed network tomography. We have proposed two models, the Compound model and the Independent Sub-flows model for this task. Further, we made specific network-wide regularizing assumptions that lead to identifiability of parameters of interest. These choices and their performance clearly depend on the data at hand and are also closely tied to the estimation strategy. Estimation, in turn poses significant challenges. As demonstrated by the simulation studies, mis-specification of the distribution family is not necessarily the biggest challenge. The heterogeneity observed in real computer network flows and of course departures from the regularizing assumptions are significant factors. Finally the Independent Sub-flows model and the Compound model provide a framework for defining, investigating the identifiability of and estimating several interesting characteristics of the joint distribution of packet and byte volumes of a flow. In particular the Independent Sub-flows model can incorporate a larger number of sub-flows (see Corollary II.16). Indeed, it is easy to see that variances of up to 3 sub-flows are identifiable from the covariance of packet and byte volumes of flows Σ_X. Using carefully chosen parametric families and the information from higher cumulants it would be possible to estimate an even larger number of sub-flows, however, the practical viability of such an approach could be limited due to the various challenges posed by data from real networks.

CHAPTER IV

Optimal Design for Sampled Data

4.1 Introduction

In this chapter we consider the problem of tracking flow volumes in a computer network using sampled data. Consider a wide area computer network such as the one depicted in Figure 4.2. As before a *flow* is defined as all traffic with common origin and destination nodes. Flow volumes have been observed to exhibit complicated structure as seen in Figure 4.3. Real time tracking (as opposed to off-line estimation of distributional properties) of flow volumes plays an important role in network management tasks, such as identifying failures together with their causes and impact, detecting malicious activity and configuring routing protocols [2, 41]. Packets of network traffic can be observed (and sampled) at router interfaces, henceforth called *observation points*. However, during the measurement process sampling is employed due to high flow volumes and resource constraints at routers. It is increasingly common for such measurement infrastructure to be deployed in computer networks [14]. Each packet from the aggregate flow at an observation point is sampled independently with a certain probability (sampling rate) [12]. Typical sampling rates range between .001-.01. Obviously low sampling rates result in large sampling noise. For every packet sampled, its header information is recorded which allows one
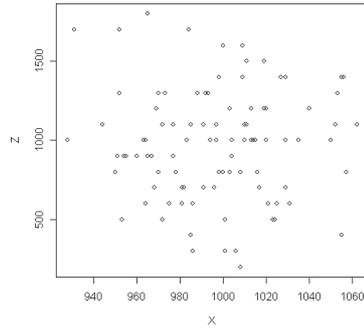
Figure 4.1: Sampling noise in estimated value Z for sampling rate $\xi = .01$

to reconstruct objects of interest, such as volumes of flows with a particular source and destination traversing the network. An important issue is how to *select* (design) the sampling rates across the network subject to resource constraints, in order to collect the maximum amount of information on the underlying source-destination flows.

Availability of sampled data on *individual* flows should clearly improve our ability to track flow volumes. However, such data can be fairly noisy at low sampling rates. As an illustrative example, suppose that a flow with volume X in a certain time interval is sampled at a rate ξ. If the number of sampled packets is N, then the usual [13] estimate of flow volume is $Z \equiv N/\xi$. Now if X is distributed as a Poisson random variable and the conditional distribution of N given X is assumed to be binomial with parameters X and ξ, then it can easily be shown that correlation between X and Z is equal to $\sqrt{\xi}$. This can be quite low for realistic sampling rates in networks. Figure 4.1 shows a scatter plot of 100 independent and identically distributed (i.i.d) samples of (X, Z) pairs. This example strongly suggests that flow volumes could be tracked more accurately by combining sampled data from across the network and (more crucially) across measurement intervals.

One way of achieving lower estimation error with the same sampling rate is

through filtering; i.e combining the present measurement with past measurements to track the time-series of flow volumes. In designing a sampling scheme for this situation one needs to take into account measurement noise and process noise.

While modeling the dynamics of flow volumes is a challenging task in itself [36], we use a simple random walk model for this purpose. This is a robust enough model to be useful in a large range of applications and leads to scalable filters. We consider the problem of minimizing the (running) estimation error through optimal design of measurement scheme in the filtering context. In this chapter, we take an optimal experiment design approach to the above problem and demonstrate its application to computer network monitoring using sampled data.

The related research on optimal design has focused on one of the following scenarios. There is a large body of work on optimal input design for dynamical systems [45, 20]. There the focus is on parameter estimation (system identification) rather than filtering as in this chapter. Another related area is sequential design for nonlinear systems [18, 19], where the optimal design depends on values of unknown parameters. While there are some commonalities, the design problem in a filtering context is unique in that the design at any time affects not just the current estimation error but also future estimation errors. The problem of optimal sensor placement in control system literature looks at an equivalent problem [1]. However, the formulation is not in terms of information matrices and the special case of random walks has not been analyzed to our knowledge.

The remainder of the chapter is organized as follows: in Section 4.2, we formulate and investigate the idealized problem of optimal design in the context of filtering for multiple random walks. In Section 4.3, we study its application to tracking flow volumes using sampled data. In Section 4.4 we look at some generalizations of both
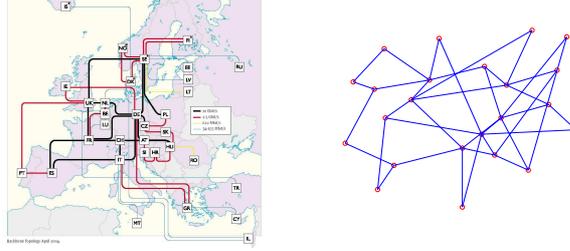
Figure 4.2: Geant Network (a) Geographic view (www.geant.net) and (b) Logical Topology
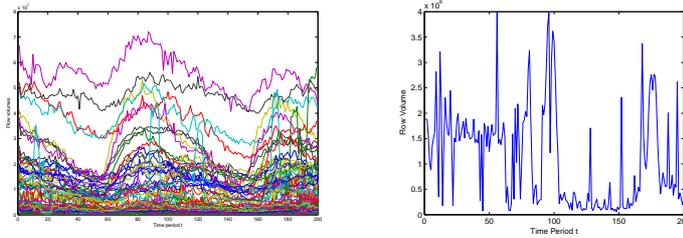
Figure 4.3: Flow volumes: (a) All flows and (b) One of the lighter flows.

the flow volume model and measurement data model. We end with discussion of a possible generalization of the steady state optimal design problem and comments in Section 4.5.

4.2   Optimal Design for Multiple Random Walks

Let us first recall the idea of E-optimality from classical experiment design literature for a simple setting. Assume we have *independent* observations

$$(4.1) \qquad\qquad y_i \sim N(x_i, 1/m_i),$$

for $i = 1, 2, \cdots, n_r$. The natural estimate for x_i is $\hat{x}_i = y_i$ for all i. It is standard to assume that the inverse variance of observation noise is roughly proportional to design variables. The inverse variance, m_i can be thought of as the information collected on parameter x_i. Specifically, we assume that the relation between an $n_r \times 1$ information vector m and an $n_o \times 1$ vector of design variables ξ is

$$(4.2) \qquad\qquad m = J\xi$$
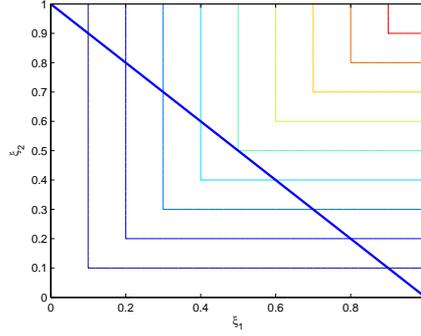
Figure 4.4: Contours of objective function for E-optimal design

For example, suppose there is a library of measurements $z_1, \cdots, z_{n_o}$, each of which is independently distributed as $z_i \sim N(x_{[i]}, \sigma_i^2 I)$, where $x_{[i]}$ is a subset of elements of x. Let ξ_i be equal to (or proportional to) the number of independent measurement of type i (replications of z_i) collected during the experiment. Then, the weighted least squares estimate y of x can be shown to have distribution given by (4.1) and (4.2). The matrix J depends on the the membership of subsets $x_{[i]}$ and variances σ_i^2 (assumed known), for $i = 1, \cdots, n_o$.

We assume that the design variables are constrained to be positive and in addition satisfy n_v linear inequality constraints. These can be written as $R\xi \le b$ where R is an $n_v \times n_o$ matrix and b is $n_v \times 1$ vector. We think of this type of constraint as a budgetary one, that specifies upper limits on weighted sums of the design variables.

Now the E-optimal design problem is given by:

$$\arg \max_{R\xi \le b} \min_i m_i$$

Note that this corresponds to minimizing the maximum mean squared error (MSE) since $1/m_i$ is the MSE in the estimate of x_i.

As an example consider the situation where $m_1 = 50\xi_1$ and $m_2 = 50\xi_2$. Further

assume the constraint

$$\xi_1 + \xi_2 \leq 1$$

Figure 4.4 shows the contours of the objective function and the boundary of the constraint. It is clear that the optimal design would be $\xi_1 = \xi_2 = .50$, which is also reasonable from the symmetry of the setup.

We extend the above criteria to steady state optimal design for random walks. Consider a collection of *independent* random walks

$$x_i(t) = x_i(t-1) + \epsilon_i(t)$$

for $i = 1, \cdots, n_r$ and $t = 1, 2, \cdots$. We assume that $Var(\epsilon_i(t)) = \sigma_i^2$ which is referred to as the innovation variance. Further, suppose we have noisy observations

$$y_i(t) = x_i(t) + \eta_i(t)$$

Let $Var(\eta_i(t)) = 1/m_i$. As before we assume the relation between observed information and design variables to be $m = J\xi$ with $n_r \times n_o$ matrix J assumed known.

The estimates of interest in this case are the ones obtained through filtering

$$\hat{x}_i(t) = E[x_i(t)|y_i(t), y_i(t-1), \cdots]$$

Let $s_i(t) = Var(x_i(t)|y_i(t), y_i(t-1), \cdots) = Var(x_i(t) - \hat{x}_i(t)|y_i(t), y_i(t-1), \cdots)$. Further, let $\tilde{m}_i = \lim_{t\to\infty} 1/s_i(t)$ when it exists. We will refer to this as the *steady state information*. When the innovation and measurement noise, $\epsilon_i(t)$ and $\eta_i(t)$ respectively, are Gaussian, the optimal filter is a Kalman filter [25]. If $s_i(t|t-1) = Var(x_i(t)|y_i(t-1), y_i(t-2), \cdots)$ then the Kalman Filter update equations give us

(4.3) $$s_i(t|t-1) = s_i(t-1) + \sigma_i^2$$
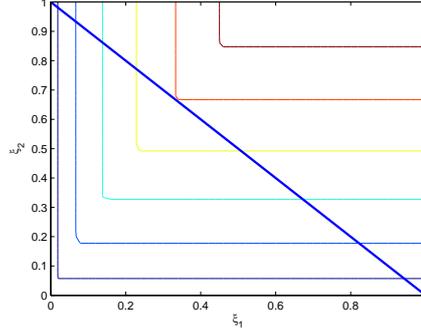
Figure 4.5: Contours of objective function for Steady State E-optimal design

and

$$(4.4) \qquad s_i(t)^{-1} = s_i(t|t-1)^{-1} + m_i$$

$$(4.5) \qquad = \left(\frac{1}{s_i(t-1)^{-1}} + \sigma_i^2\right)^{-1} + m_i$$

Thus

$$\tilde{m}_i = \left(\frac{1 + \sigma_i^2 \tilde{m}_i}{\tilde{m}_i}\right)^{-1} + m_i$$

or

$$\sigma_i^2 \tilde{m}_i^2 - \sigma_i^2 m_i \tilde{m}_i - m_i = 0$$

Hence

$$\tilde{m}_i = \frac{m_i \sigma_i^2 + \sqrt{m_i^2 \sigma_i^4 + 4 m_i \sigma_i^2}}{2\sigma_i^2}$$

We define the steady state E-optimal design problem as

$$\arg \max_{R\xi \leq b} \min_i \tilde{m}_i$$

As an example consider the situation where $m_1 = 50\xi_1$ and $m_2 = 50\xi_2$. Further, let the innovation noise be characterized by $\sigma_1 = 0.1$ and $\sigma_2 = 0.2$. As before we assume the design constraint

$$\xi_1 + \xi_2 \leq 1$$

From Figure 4.5, notice that even though there is symmetry in the measured information, the first random walk is smoother than the second one and hence less measurement resources need to be allocated to it.

4.2.1 Optimization for Steady State E-optimal Design

To solve the steady state E-optimal design problem we have to maximize θ subject to

(4.6)
$$\frac{m_i \sigma_i^2 + \sqrt{m_i^2 \sigma_i^4 + 4 m_i \sigma_i^2}}{2\sigma_i^2} \geq \theta$$

for $i = 1, \cdots, n_r$ and

$$R\xi \leq b$$

Equation (4.6) can be equivalently written as

$$\theta^2 \leq m_i \left(\theta + \frac{1}{\sigma_i^2} \right)$$

which is a hyperbolic constraint [33]. Thus, this problem can be cast as second order cone program. Such optimization programs can be solved efficiently through interior point methods [4].

4.2.2 Myopic Approach

In the following, we present a greedy alternative to steady state optimal design. As before, assume $y_i(t) = x_i(t) + \eta_i(t)$. Further, we assume that $Var(\eta_i(t)) = 1/m_i(t)$. i.e. we allow for time varying design variables $\xi(t)$ with $m(t) = J\xi(t)$. As before $s_i(t) = Var(\hat{x}_i(t)|y_i(t), y_i(t-1), \cdots)$. Define, information at time t, $\tilde{m}_i(t) = 1/s_i(t)$. Note that $\tilde{m}_i(t)$ is a function of $\xi(t), \xi(t-1), \cdots$.

The Myopic E-optimal design at time t is defined as

$$\arg \max_{R\xi(t) \leq b} \min_i \tilde{m}_i(t)$$

From 4.5 it follows that this is a linear program. Not surprisingly, the myopic optimal design is a much easier problem than steady-state optimal design even in more general settings as noted in Section 4.4. Note that since the sampling rates are allowed to vary with time, it may have an objective function larger than the steady state optimal case. However, as the objective of optimization is to maximize present information with no regard to impact on future information, such a scheme cannot be guaranteed to perform well in the long run.

4.3 Application to tracking flow volumes

The ideas developed above can be used for sampling rate design for tracking flow volumes in a computer network. As mentioned in the introduction we will use the random walk model for flow volumes. Suppose a computer network has n_r origin destination flows. Let $x_i(t)$ be the volume of ith flow in time interval t, for $i = 1, \cdots, n_r$. These flow volumes are tracked using sampled data which are noisy. Recall that flows are sampled at router interfaces which we refer to as observation points. All flows traversing an observation point (router interface) experience the same sampling rate. Each incoming edge at a node in Figure 4.2 is an interface of the corresponding router. Each router typically has multiple interfaces and each flow may traverse multiple observation points.

Suppose there are n_o observation points on the network where sampled data on flows can be collected. Further, assume that sampling rates of $\xi = (\xi_1, \cdots, \xi_{n_o})'$ are used at observation points $1, \cdots, n_o$, respectively. Any given observation point $k \in \{1, \cdots, n_o\}$ generates estimates for g_k elements of $x(t)$; i.e. the number of flows that go through that node. Thus, a total of $n_g = \sum_{k=1}^{n_o} g_k$ measurements are available, which need to be optimally combined to get the required estimates.

Assume that $k(i)$ is the observation point at which the ith measurement is collected and $l(i)$ the corresponding flow. Thus $k() : \{1, \cdots, n_g\} \rightarrow \{1, \cdots, n_o\}$ and $l() : \{1, \cdots, n_g\} \rightarrow \{1, \cdots, n_r\}$. Further let,

$$E[z_i(t)|x_{l(i)}(t)] = x_{l(i)}(t)$$

and for now assume

(4.7) $$Cov(z_i(t)|x_{l(i)}(t)) = \mu_{l(i)}/\xi_{k(i)}$$

where $\mu_i = E[x_i(t)]$. The exact sampling mechanism and approximation involved in the above relation are described in Section 4.3.2. Thus, in vector notation we get

(4.8) $$E[z(t)|x(t)] = Lx(t)$$

where L is a $n_g \times n_r$ matrix with $L_{ij} = 1$ only if $l(i) = j$ and 0 otherwise and

(4.9) $$Cov(z(t)|x(t)) = D$$

where D is a $n_g \times n_g$ diagonal matrix. Using (4.7), the inverse of D is given by $D^{-1} = \sum_k \xi_k \Psi_k$, where Ψ_k, $k = 1, \cdots, n_o$, are $n_g \times n_g$ diagonal matrices with

(4.10) $$[\Psi_k]_{ii} = \begin{cases} 1/\mu_i & \text{if } k = k(i) \\ 0 & \text{otherwise} \end{cases}$$

Let $y(t)$ be the GLS estimate of $x(t)$ under equation (4.8) and (4.9). Thus

(4.11) $$Cov(y(t)|x(t)) = (L'D^{-1}L)^{-1}$$

(4.12) $$= (\sum_k (L'\Psi_k L)\xi_k)^{-1}$$

From definition of L it follows that the columns of L are orthogonal. Thus, the matrix in (4.11) is diagonal. Further

$$Diag(Cov(y(t)|x(t))) = m = J\xi$$

where

$$[J]_{ik} = L'_{\cdot,i} \Psi_k L_{\cdot,i}$$

We will refer to the above as the linear model.

Sampling is employed in network flow measurements because measurement resources like CPU time and available storage are limited. Typically, all observation points (router interfaces) belonging to a particular router share these resources. We assume that the sampling rates are constrained to lie in a convex polygon $R\xi_t \leq b$. This includes the case where the sum of sampling rates on the interfaces of a router is bounded above by the budget for that router. We will focus on this constraint for this and next section (Section 4.4 has some other examples). In this case, the constraints are given as one linear inequality for each router.

For the available data we set up the performance evaluation as follows. We use the Geant network topology, which has $n_v = 23$ nodes (routers) and 37×2 bidirectional edges. The available data [46] correspond to flow volumes over time. Each time interval is equal to 15 minutes. The original data set spans 4 months, but we focus on the first 200 time intervals, to avoid severe non-stationarities inherent in an evolving network. Further, we focus on the top 25% of measured flows by volumes since one is typically interested in tracking heavy flows. This corresponds to $n_r = 76$ flows. We assume that these flows are routed through minimum distance paths, which is a common routing mechanism in wide area networks [37]. We assume that sampled data can be collected at each incoming edges of a router and thus we have $n_o = 37 \times 2$ observation points. We assume that the sum of sampling rates on all interfaces of a router is bounded above by .01, i.e. $b_i = .01$. Finally we estimate the σ_i^2 and μ_i,parameters associated with the flow volume processes, and assume they are available for filtering purposes and measurement design.

(a)   (b)
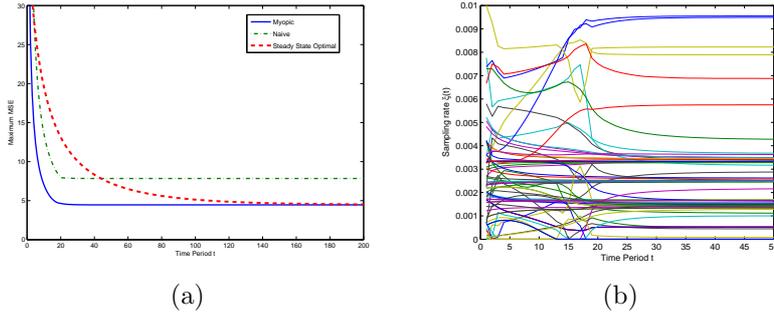
Figure 4.6: Performance of various sampling schemes (a) and sampling rates at various interfaces under myopic scheme (b)
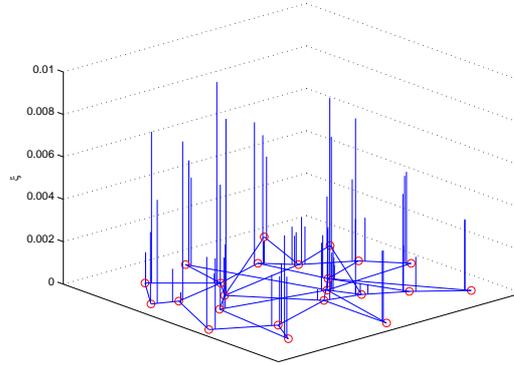
Figure 4.7: Spatial view of steady state optimal sampling rates

For the purpose of comparison we define a naive sampling scheme as follows. For any given router, equal sampling rate is allocated to every interface that carries any of the 76 flows of interest. This allocation is done so as to make the corresponding budget constraint tight. For example, suppose the ith router has 5 interfaces but only 4 of them are traversed by one of the 76 flows of interest. In this case each of the latter 4 interfaces will be allocated a sampling rate of $b_i/4$ while the remaining interface will be allocated a sampling rate of 0.

4.3.1   Linear Model: Performance of Various Sampling Schemes

Figure 4.6 (a) shows the value of maximum MSE versus time. Note that as information accumulates over time we get an improvement in performance under all three sampling mechanisms. Here performance is measured as the maximum of $s_i(t)$ over all flows, calculated using equations (4.3) and (4.5). Surprisingly, both the myopic and steady state sampling mechanisms perform equally well in steady state and we achieve a 42% improvement over the naive sampling in the steady state. Figure 4.6 (b) shows that the myopic optimal sampling rates at all observation points reach a steady state. Figure 4.7 shows the value of steady state sampling rates at various router interfaces in the network topology. Even though the myopic scheme has the flexibility of time varying sampling rates, if the sampling rates do reach a steady state its performance can clearly be no better than the steady-state optimal scheme.

4.3.2   Departures from Linear Model: Performance with Geant Data

A more detailed model for flow volumes and sampled measurements would have to include significant departures from the linear model assumed above. First, the true flow volumes clearly have more structure than independent random walks, as seen in Figure 4.3. In applying the above ideas to the Geant data, we will investigate their robustness to the independent random walk assumption.

A more serious departure is the following. Suppose that a flow with volume X in a certain time interval is sampled at a rate ξ. If the number of sampled packets is N, then the usual (approximate ML) estimate of flow volume is $Z \equiv N/\xi$. The variance of measurement noise can be shown to be $\mathrm{Var}(Z|X) \simeq X/\xi$ [11]. Thus μ_i in (4.10) is actually equal to the unknown $x_i(t)$.
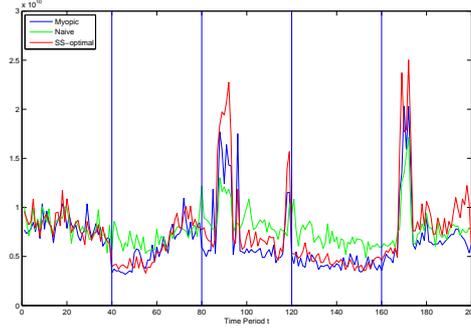
Figure 4.8: Performance of various sampling schemes using batch sequential design with flow volumes from Geant Data

The observation above implies that in application of the presented techniques to sampled network data, one would have to rely on an approximate model for measurements $z_i(t)$. We will follow an approach similar to batch sequential design [19]. Assume that the sampling rates are to be held constant for a batch of contiguous time intervals. At the beginning of each batch we use the most recent estimate $\hat{x}_i(t-1)$ in place of μ_i in (4.10) for sampling rate design. For filtering purposes, we employ a Kalman filter with $\hat{x}_i(t-1)$ in place of μ_i in equation (4.10) at each time t. We replace the budget constraint inequalities $R\xi \leq b$ with the corresponding equalities $R\xi = b$ to force full utilization of available resources. For routers that are traversed by at least one of the 76 flows of interest, we introduce additional equality constraints as follows. Design variable ξ_k for an interface k not traversed by one of the 76 flows of interest is constrained to be identically 0. Figure 4.8 shows the performance of different sampling schemes averaged over 200 realizations of sampled data. The sampled data emulates the exact sampling mechanism described above (with respective sampling rates) with the Geant data treated as the underlying (unobserved) flow volumes.

Sampling rates were adjusted only at the beginning of a 40 time period block and
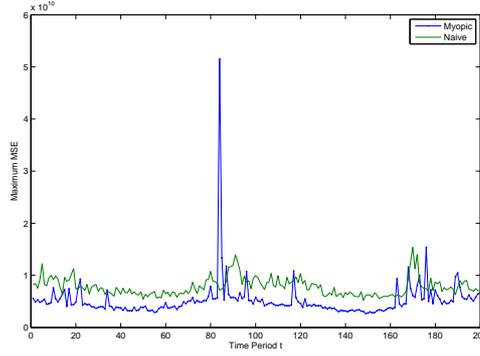
Figure 4.9: Performance of fully time varying myopic and naive sampling mechanism with flow volumes from Geant Data

were held constant over each block. In the first block the sampling rates were forced to be the same as the naive scheme irrespective of the sampling mechanism under study. Notice that for low values of the objective function (maximum mean squared error) the myopic and steady state allocation perform better than the naive allocation. On the other hand when the maximum mean squared error spikes, the naive allocation performs better indicating robustness to model departures. The median (over time periods 41 to 200) of maximum MSE for myopic, naive and steady-state optimal sampling is 5.46×10^9, 7.49×10^9 and 6.14×10^9, respectively. Thus, the myopic scheme performs better than the steady state optimal scheme, which in turn performs better than the naive scheme.

Finally, we look at the performance of myopic allocation when the above scheme is employed with a block size of just one time interval; i.e. sampling rates were adjusted at the beginning of each time period using the myopic scheme. The results are displayed in Figure 4.9. As before the current estimate of flow volumes is used in place of μ_i in equation (4.10) for both filtering and myopic sampling scheme design. The myopic sampling scheme can be seen to perform better than the naive version in most time periods. The median (over time periods 1 to 200) of maximum MSE is

4.50×10^9 and 7.44×10^9 for myopic and naive sampling respectively.

4.4 Utilizing SNMP Data: State Space Models

As before suppose that there are n_r flows in a network. However, in addition to sampled data now assume that one can obtain accurate measurements (e.g. SNMP) about the sum of the flows at the n_l links. Typically, the number of flows is significantly larger than the number of links; i.e. $n_l << n_r$.

As before let $x(t) = (x_1(t), \cdots, x_{n_r}(t))'$ be the vector of a random realization of volumes of different flows, in time interval t. Further, assume that n_l SNMP measurements $y(t) = (y_1(t), \cdots, y_{n_l}(t))'$ for the same time interval t are available. Each SNMP measurement is an estimate of the sum of all traffic traversing a particular link in a particular measurement interval. Thus, we can write

$$y(t) = Ax(t) + u(t)$$

where A is a $n_l \times n_r$ routing 0/1 matrix that describes the routes of the various network flows and $u(t)$ is independent noise (to model possible errors in SNMP data). It is assumed that $Cov(u(t)) = \sigma^2 I$. We assume that the matrix A does not change over time. However, the case of time varying routing is handled in a straightforward way.

The observation model can now be written as

$$(4.13) \qquad \begin{pmatrix} Y(t) \\ Z(t) \end{pmatrix} = \begin{pmatrix} A \\ L \end{pmatrix} X(t) + \begin{pmatrix} u(t) \\ \epsilon(t) \end{pmatrix}$$

with

$$\mathrm{Cov}(\epsilon(t)) = \Psi(t, \xi(t))^{-1}$$

where $\Psi(t, \cdot)$ is a linear function and $\xi(t)$ is the value of design variables in time

interval t. For sampled flow volume measurements $\Psi(t, \xi)^{-1} = D$, due to equation (4.9).

Generalizing the random walk model, the joint distribution of $(x(1), x(2), \cdots)$ can be modeled through a state-space approach. Such an approach has been shown [30, 41] to perform well for modeling flow volumes. Specifically, we assume the following transition equation:

$$(4.14) \qquad\qquad x(t) = Cx(t-1) + w(t)$$

Further, let $Cov(w(t)) = W$. The state transition equation may suffer from mis-specification too and we look at this issue numerically in Section 4.4.3. Note that the above equation corresponds to a vector auto-regressive model of order 1. Higher order models and model selection issues have been investigated by Zhao et. al. [52]. Flow modeling is an active area of research and one that is beyond the scope of this work. Equation (4.22) is an extremely simple and robust way of modeling flow volumes that still captures spatio-temporal dependence which is key to filtering.

As before, we will assume that there are certain budget and positivity constraints on the sampling rates $(\xi(1)', \xi(2)', \cdots)$. For the remainder of this section, it is assumed that the constraints can be simplified as $\xi(t) \in \Xi(t)$, where $\Xi(t)$ are convex sets.

Given the state space model one can use a Kalman-Filter to recursively compute estimates of $x(t)$ based on information available at time t. Such an approach using only SNMP data $(y(t)', y(t-1)', \cdots)'$ has been been demonstrated in [41]. Let $\hat{x}(t)$ be the Kalman-Filter estimate of $x(t)$ based on $(y(t)', z(t)', y(t-1)', z(t-1)', \cdots)$ and let $\Sigma(t)$ denote the covariance matrix of estimation error i.e. $\Sigma(t) = E[(x(t) - \hat{x}(t))(x(t) - \hat{x}(t))'|y(t)', z(t)', y(t-1)', z(t-1)', \cdots]$.

The conditional covariance of $x(t)$ given $(y(t-1)', z(t-1)', y(t-2)', z(t-2)', \cdots)$

can be written as [25]:

$$(4.15) \qquad \Sigma(t|t-1) = C\Sigma(t-1)C' + W$$

Further, we have the following updating equation:

$$(4.16) \qquad \Sigma(t) = \Sigma(t|t-1) - \Sigma(t|t-1)J'F(t)^{-1}J\Sigma(t|t-1)$$

where $J' = (A', L')$ and

$$F(t) = J\Sigma(t|t-1)J' + \begin{pmatrix} \sigma^2 I & 0 \\ 0 & \Psi(t, \xi(t))^{-1} \end{pmatrix}$$

4.4.1 Formulation of the Myopic Optimal Sampling Problem

It is reasonable to seek to minimize some appropriate functional of $(\Sigma(1), \Sigma(2), \cdots)$, since this would lead to maximization of information about the underlying flows. More generally, if interest is restricted to a subset of flows $K'X(t)$, selected by a matrix K, then we seek to minimize a functional of $(K'\Sigma(1)K, K'\Sigma(2)K, \cdots)$. Notice that any *joint* minimization of $(K'\Sigma(1)K, K'\Sigma(2)K, \cdots)$ would have very high complexity. We make some comments on the steady state optimal design problem in Section 4.5. Here we restrict attention to the myopic design problem, i.e. optimizing $K'\Sigma(t)K$ at every time interval $t = 1, 2, \cdots$.

We will refer to the following problem as the (myopic) optimal design problem for state space models. At each time interval t, the design objective would be to minimize $f(K'\Sigma(t)K)$ for some functional f. If $\mathbb{S}$ is the set of symmetric semi-definite matrices, then $f : \mathbb{S} \to \mathbb{R}$ is in general matrix isotonic [38]. The dimension of matrices in $\mathbb{S}$ is usually clear from the context and we will use the same notation $f()$ for all cases. Common choices of f are the determinant (also known as the D-optimality criterion in the design of experiments literature), the maximum eigenvalue (E-optimality) and

the trace (A-optimality). Our exposition will focus on the E- and A- optimality criteria.

Remark: Although there is nothing special about these criteria conceptually, and in fact network engineers may find other criteria like relative mean squared error (MSE) attractive, we focus on these for the following reasons. First, these criteria result in optimization problems that are simple convex programs and thus intuitive and computationally easier to handle. Secondly, the purpose of this work is to demonstrate that it is possible to get better performance with respect to specific design objectives though a targeted allocation of sampling resources. Generally any objective function would involve a scalarization of the covariance matrix of estimation error. Many reasonable scalarizations are well approximated by above criteria or simple modifications thereof. We discuss the connection between the optimality criteria and quantities like MSE in Section 4.4.2.

Using the matrix inversion lemma $((A + BDB')^{-1} = A^{-1} - A^{-1}B(B'A^{-1}B + D^{-1})^{-1}B'A^{-1})$, we can write

$$(4.17) \qquad \Sigma(t)^{-1} = \Sigma(t|t-1)^{-1} + \sigma^{-2}A'A + L'\Psi(t, \xi(t))L$$

The above can be interpreted as an information update equation [25]. Here $\Sigma(t)^{-1}$, the information available at time t, is represented as the sum of $\Sigma(t|t-1)^{-1}$, the information from time $t-1$, $\sigma^{-2}A'A$, information from SNMP data and $L'\Psi(t, \xi(t))L$ information from sampled data.

Now,

$$(4.18) \qquad K'\Sigma(t)K = K'(\Sigma(t|t-1)^{-1} + \sigma^{-2}A'A + L'\Psi(t, \xi(t))L)^{-1}K$$

Note that both $\Sigma(t)$ and $K'\Sigma(t)K$ are of the general form $G'M(\xi(t))^{-1}G$ where G is a known matrix and $M(\cdot)$ is a linear function. This will allow us to write the

resulting optimizations as canonical convex programs.

A value of $\sigma = 0$, implies perfect knowledge in a specific subspace; i.e. $Ax(t) = y(t)$. We proceed by reparameterizing the variable $x(t)$ so that the equality $y(t) = Ax(t)$ can be solved for a subset of these new variables and the conditional covariance of the remaining variables is positive definite. Let $\tilde{x}(t) = Q'x(t)$, where Q is the unitary matrix from the eigenvalue decomposition

$$A'A = Q\Lambda Q'$$

and Λ being the diagonal eigenvalue matrix with first n_l diagonal values positive and the rest zero. The covariance of $\tilde{x}(t)$ given $(y(t)', z(t)', y(t-1)', z(t-1)', \cdots)$ can be written using the following proposition.

Proposition IV.1. *Let X and ϵ be independent normal random vectors with covariances Σ and Ψ^{-1} respectively. Further let $Y = AX$ for A with full row rank, $Z = LX + \epsilon$ and $A'A = Q\Lambda Q'$ be the eigenvalue decomposition of $A'A$. Let $Q = (Q_1, Q_2)$ such that (w.l.o.g) $AQ_1 \equiv P$, a full rank square matrix and $AQ_2 = 0$. Define*

$$\tilde{X} \equiv \begin{pmatrix} \tilde{X}_1 \\ \tilde{X}_2 \end{pmatrix} \equiv \begin{pmatrix} Q_1'X \\ Q_2'X \end{pmatrix}.$$

Then, the covariance of $\tilde{X}_1$ given Y is 0 and the covariance of $\tilde{X}_2$ given Y and Z is $(Q_2'\Sigma^{-1}Q_2 + Q_2'L'\Psi LQ_2)^{-1}$.

Proof. Note that $AX = P\tilde{X}_1$ and thus $\tilde{X}_1 = P^{-1}Y$. Hence, clearly the covariance of $\tilde{X}_1$ given Y is 0. Further, the covariance of $\tilde{X}_2$ given $(Y', Z')'$ is the same as $\Sigma_{\tilde{X}_2|(\tilde{X}_1', \tilde{Z}')'}$, the covariance of $\tilde{X}_2$ given $(\tilde{X}_1', Z')'$. Let $\Sigma_{\tilde{X}_2|\tilde{X}_1}$, $\Sigma_{\tilde{X}_2, Z|\tilde{X}_1}$ and $\Sigma_{Z|\tilde{X}_1}$ be the covariance of $\tilde{X}_2$, cross-covariance of $\tilde{X}_2$ and Z and the covariance of Z

respectively given $\tilde{X}_1$. Then

$$\Sigma_{\tilde{X}_2|(\tilde{X}_1',\tilde{Z}')'} = \Sigma_{\tilde{X}_2|\tilde{X}_1} - \Sigma_{\tilde{X}_2,Z|\tilde{X}_1}\Sigma_{Z|\tilde{X}_1}^{-1}\Sigma_{\tilde{X}_2,Z|\tilde{X}_1}'$$

$$= \Sigma_{\tilde{X}_2|\tilde{X}_1} -$$

$$\Sigma_{\tilde{X}_2|\tilde{X}_1}Q_2'L'(LQ_2\Sigma_{\tilde{X}_2|\tilde{X}_1}Q_2'L' + \Psi^{-1})^{-1}LQ_2\Sigma_{\tilde{X}_2|\tilde{X}_1}$$

$$(4.19) \qquad\qquad = (\Sigma_{\tilde{X}_2|\tilde{X}_1}^{-1} + Q_2'L'\Psi LQ_2)^{-1}$$

The last equality follows from the matrix inversion lemma. Further

$$\Sigma_{\tilde{X}_2|\tilde{X}_1} = Q_2'\Sigma Q_2 - Q_2'\Sigma Q_1(Q_1'\Sigma Q_1)^{-1}Q_1'\Sigma Q_2$$

Therefore,

$$\begin{aligned}\Sigma_{\tilde{X}_2|\tilde{X}_1}^{-1} &= [(Q'\Sigma Q)^{-1}]_{22} \\ &= [Q'\Sigma^{-1}Q]_{22} = Q_2'\Sigma^{-1}Q_2\end{aligned}$$

which establishes the desired result.

□

Hence the covariance of $\tilde{x}(t)$ given

$(y(t)', z(t)', y(t-1)', z(t-1)', \cdots)$ is given by

$$\tilde{\Sigma}(t) = \begin{pmatrix} 0 & 0 \\ 0 & (Q_2'\Sigma(t|t-1)^{-1}Q_2 + Q_2'L'\Psi(t,\xi(t))LQ_2)^{-1} \end{pmatrix}$$

Finally,

$$\begin{aligned}\Sigma(t) = Q\tilde{\Sigma}(t)Q' &= Q_2(Q_2'\Sigma(t|t-1)^{-1}Q_2 \\ (4.20) &\quad + Q_2'L'\Psi(t,\xi(t))LQ_2)^{-1}Q_2'\end{aligned}$$

which shows that both $\Sigma(t)$ and $K'\Sigma(t)K$ are of the form $G'M(\xi(t))^{-1}G$.

4.4.2 Optimality Criteria and Convex Programs

For the classical optimality criteria mentioned earlier, the optimization problems can be written as canonical convex programs [29] that can be solved using any of the commonly available software packages. The usual class of algorithms used for solving convex optimization problems, Interior Point Methods, have polynomial time complexity and have been successfully employed for various large scale optimization problems. Since we focus on one measurement interval at a time, the general optimization problem is to minimize $f(G'M(\xi)^{-1}G)$ subject to $\xi \in \Xi$.

1. *E-optimality.* E-optimality corresponds to minimizing the maximum eigenvalue of a covariance matrix. Equivalently, this minimizes the worst possible variance of any linear combination of the error vector. Roughly, this would lead to a small value for the largest of the MSE.

   In this case we get the following program:

   $$\underset{\theta \in \mathbb{R}, \xi \in \Xi}{\text{maximize}}\, \theta$$

   subject to

   (4.21) $$G'M(\xi)^{-1}G \preceq \theta^{-1}I$$

   Recall that if $A \succ 0$ and $C \succ 0$ then

   $$\begin{pmatrix} A & B \\ B' & C \end{pmatrix} \succeq 0$$

   if and only if $C \succeq B'A^{-1}B$ or equivalently if and only if $A \succeq BC^{-1}B'$. This follows from taking the Schur complement of C and A respectively.

   Thus, the constraint (4.21) is equivalent to

   $$M(\xi) \succeq \theta GG'$$

The above program is a Semi-Definite Program (SDP) if Ξ is a convex polygon [4].

2. *A-optimality.* A-optimality corresponds to minimizing the trace of a covariance matrix and hence minimizing the average MSE. In this case we get the following program:

$$\underset{t\in\mathbb{R}^n, \xi\in\Xi}{\text{minimize}} \sum_{i=1}^{n} t_i$$

subject to

$$t_i \geq e_i' G' M(\xi)^{-1} G e_i$$

for $i = 1, \cdots, n$, where n is the number of columns in G and e_i is the ith unit-vector of $\mathbb{R}^n$. Using Schur complement we can write the given constraints as:

$$\begin{pmatrix} M(\xi) & G e_i \\ e_i' G' & t_i \end{pmatrix} \succeq 0$$

Thus we get a SDP if Ξ is a convex polygon.

3. *D-optimality.* D-optimality corresponds to minimizing the determinant of a covariance matrix and hence minimizing the volume of the associated confidence ellipsoid. In this case we get the following program:

$$\underset{\xi\in\Xi}{\text{minimize}} \log \det G' M(\xi)^{-1} G$$

Note that if G does not have full column rank then the minimization is trivially unbounded. Although D-optimality has a natural interpretation for polynomial regression models we will focus on A- and E- optimality due to their simplicity and interpret-ability in terms of error variances.
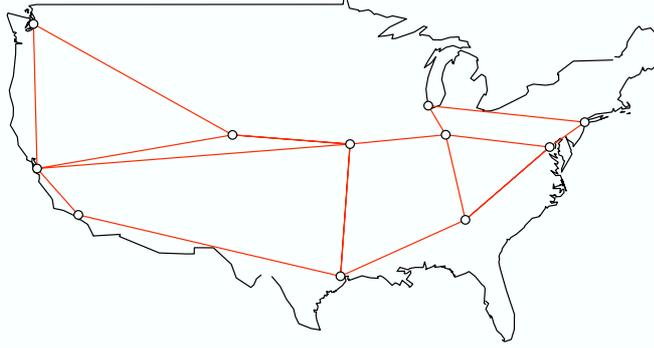
Figure 4.10: Abilene Topology used for Performance Evaluation Purposes

4.4.3 Performance Evaluation

In this Section, the performance of the proposed design scheme for the sampling rates is evaluated. Specifically, the following issues are addressed: in the first part, the cumulative effect of the greedy design optimization steps is examined, when the true model of flow volumes is given by the state-transition equation (4.14) and errors in the observation equation (4.13) have constant variance. In the second part, the effects of model misspecification in the state transition equation and approximations in the supposed variances of errors of the observation equation are investigated.

In the following, we describe the overall setup for the numerical study.

1. *Topology, Routing and Observation Points.* The Abilene network topology (Figure 4.10) is used in our experimental setup. It consists of 11 nodes and $16 \times 2 = 32$ directed edges between pairs of nodes (bidirectional links). Flows exist between all pairs of nodes resulting in a total of $11 \times 10 = 110$ flows. We assume that these flows are routed through minimum distance paths. Further we assume that SNMP data are available from all the edges. Similarly, sampled data are collected at each edge. All the incoming edges at a node are considered as the interfaces of the corresponding router.

2. *Budget Constraints.* The budget constraint represents an effective way of lim-

iting the number of samples collected at a particular router. For this numerical study, we use a budget constraint of the form $R\xi_t \leq b$. Here R_{ij} is equal to the square of the number of flows traversing edge j if edge j is incident on node i and 0 otherwise. Thus, the weighted sum of sampling rates on the incoming edges of node i is bounded above by b_i. It can be seen that the cost of sampling on an edge goes up as the square of number of flows traversing that edge. This is a very reasonable assumption, because not only does sampling on a heavily used link results in a large number of samples and thus a significant demand on resources, but also on such a link less resources are available for sampling to begin with. The elements of b were identically chosen to be some b_0 to ensure realistic sampling rates. The above router level constraint implies the network level constraint $\mathbf{1}'R\xi_t \leq \mathbf{1}'b$. While most of the following assumes router level constraints we will also investigate the result of imposing only the network level budget constraint.

Notice that the above setup has the following symmetry property. Define a short flow to be a flow that traverses exactly one edge or equivalently exactly one router interface. Further, each interface is traversed by exactly one short flow. By definition, any given interface is the only point in the network from where sampled data is available for the short flow traversing it.

The above symmetry has important consequences for optimal design for the scenario described in the following. An intuitive understanding of its effects can be described as follows. A low sampling rate on any interface would result in an information deficit on the corresponding short flow, which in turn will result in a large measurement error variance due to the inverse relation between variance and sampling rate. From equation (4.17), the total information at time t is the sum of

information obtained from the previous time period $t-1$, from SNMP measurements and from sampled data. Suppose that the sum of information from time $t-1$ and from SNMP measurements do not ameliorate the problem of information deficit on short flows. Further, suppose that there is roughly equal information from this sum on each short flow. In this case, the objective function of A-optimality (i.e. average estimation error variance) is heavily influenced by the large measurement variances of short flows. Further, the objective function of E-optimality is determined to an even larger extent by the largest measurement error variance. Thus, to the extent possible given the budget constraints, all the short flows must be sampled roughly equally under the optimal allocation. This in turn implies that all interfaces should be allocated roughly equal sampling rates, again to the extent possible under the budget constraints. This was indeed found to be true as described in the following.

On the other hand, when interest focuses on a specific subsystem of flows $K'x(t)$, then it is not a priori clear what an optimal or near optimal solution should be. For most of following, we assume that the interest is restricted to long flows; i.e. flows that traverse four or more edges in the employed topology. A matrix K was used to subset $x(t)$ accordingly.

Cumulative Effect of Myopic Optimality

Recall that at each time interval t, the measurement scheme was designed to better utilize the side information available from past measurements $\Sigma(t|t-1)^{-1}$ and from SNMP measurements $\sigma^{-2}A'A$. The gains compared to a homogeneous allocation may be small in a single step but these gains propagate across time in the form of prior information according to equation (4.15). The purpose of this section is to study how the gains from using an optimal design accumulate over time.

When the assumed observation and state transition models (4.134.14) are the true

model and the true C, W, σ and $\Psi(t, \cdot)$ matrices are known, then $\Sigma(t)$ calculated using (4.15,4.16) is the true covariance of estimation error. Hence, we can calculate the covariance of estimation error for any sequence of allocations $(\xi(1)', \xi(2)', \cdots)$.

For this investigation we assume $\sigma = 0$, $C = I$, $W = I$ and $\Psi(t, \xi) = \Psi(\xi) = \sum_k \xi_k \Psi_k$. The important calibration issue here is to choose Ψ_k to maintain realistic order of magnitude of measurement error $\Psi(\xi)^{-1}$ with respect to the innovations covariance W. Empirical data has shown [30] that the innovations in flow volumes have a variance of roughly the same magnitude as the flow volume itself. Therefore, it is expected that $[W]_{ii}$ is of the same order of magnitude as X_i. Using this observation and equation (4.10) we choose

$$[\Psi_k]_{ii} = \begin{cases} 1 & \text{if } k = k(i) \\ 0 & \text{otherwise} \end{cases}$$

For the sake of comparison, a set of naive allocations is defined as follows. If we are interested in the estimation error for all flows and there is a router level budget constraint, it is given by ξ for which each interface on a router has equal sampling rate and $R\xi = b$. Note that interfaces on different routers can have different sampling rates. This ξ is also used as the common initial allocation $\xi(0)$ for all allocation schemes considered in this section. On the other hand, if there is only a network level budget constraint and we are interested in the estimation error for all flows, the naive allocation correspond to ξ for which all interfaces in the network have equal sampling rates and $\mathbf{1}'R\xi = \mathbf{1}'b$. Finally, when we are interested in the estimation error in $K'x(t)$ (the subset of long flows) the naive allocation is given by ξ for which any interface that is not traversed by any long flow has 0 sampling rate, all other interfaces have sampling rates determined as above for router level and network level budget constraints, respectively.
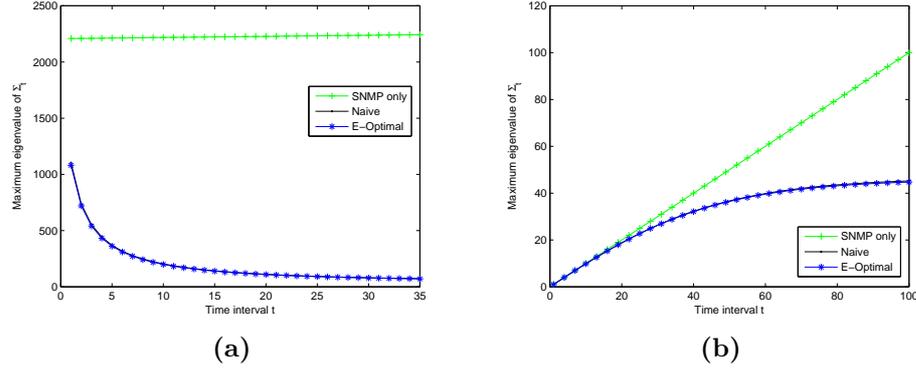
Figure 4.11: Performance of E-optimal sampling for (a)noisy and (b) noise-free initialization

We consider two sets of initial conditions. The first one, referred to as noise-free initialization, assumes perfect knowledge at time 0 ($\Sigma(0) = 0$). The second, referred to as noisy initialization, assumes

$$\Sigma(0) = Q_2(Q_2'L'\Psi(\xi(0))LQ_2)^{-1}Q_2'$$

From equation (4.20), the latter corresponds to starting with only one measurement interval worth of information. All results are for router level budget constraints, unless otherwise stated.

First consider the case when we are interested in estimation errors of all flows. Figures 4.11 (a) and (b) show the value of the objective function for E-optimality, i.e. the largest eigenvalue of $\Sigma(t)$, over time for 3 Kalman filters. These filters respectively use SNMP data only, SNMP and naively sampled data and SNMP and optimally sampled data. Clearly inclusion of sampled data gives substantial improvement over SNMP alone. Moreover, as information accumulates over time we get an improvement in performance. This can be viewed as the gain from using a temporal model to combine samples from across time intervals. For the E-optimality criterion, the performance of the naive and optimal allocations is fairly similar. This is not surprising given the symmetry in our setup, discussed above. Figure 4.12 presents
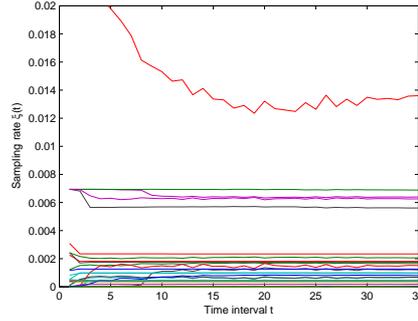
Figure 4.12: Evolution of E-optimal sampling rates

the evolution of the allocations $\xi(t)$ over time for noisy initialization. Clearly the sampling rates achieve a steady state for most of the interfaces after very few time intervals (3-5) and certainly before $t \approx 15$.

The above qualitative assessments hold for both the E- and A- optimality criteria. As an example, Figure 4.13 shows the value of the objective function for A-optimality, i.e. the trace of $\Sigma(t)$ over time starting from noisy initialization based on the three Kalman filters employed. In this case, the third filter uses A-optimally sampled data. Notice that the Kalman filter based on optimally sampled data has slightly better transient behavior than the one based on naively sampled data. The difference from the E-optimal case (Figure 4.11(a)) is due to the fact that A-optimality, being determined by the sum of variances, is less sensitive to high variance or low information on particular flows. Hence, the A-optimal solution is driven by short flows to a lesser degree than the E-optimal solution, despite the symmetry property mentioned earlier.

Next, we focus on the case when the interest is on $K'x(t)$, the subset of long flows. Figures 4.14 (a), (b), (c) and (d) show the relevant performance metric over time for E- and A- optimal sampling strategies for the two sets of initial conditions. In this case, we see a clear advantage in using optimal sampling over naive sampling. In
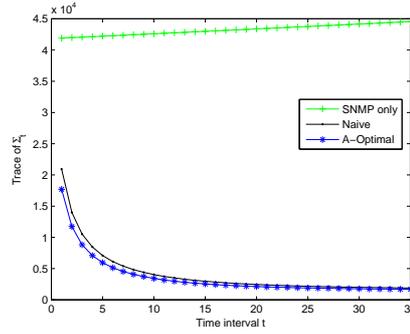
Figure 4.13: Performance of A-optimal sampling for noisy initialization

respective performance metrics the improvement over naive sampling is 29.7% and 26.0% for E-optimal design and 18.0% and 15.7% for A-optimal design.

Next, we look more closely at the steady state sampling rates. Figures 4.15(a) and (b) show the evolution and spatial profile of the final E-optimal sampling rates under noisy initialization. In Figure 4.15 (c) the ξ_k values versus k for the above allocation and the naive allocation are shown. Clearly the optimal allocation has a sparser support. Intuitively all interfaces that have a high cost but do not provide enough information on the *long flows* have zero sampling rate. Figure 4.15 (d) shows for each flow the cumulative sampling rate over all observation points that the flow traverses for naive and optimal allocation, respectively. Notice that the minimum allocation to a long flow is larger under optimal sampling (.0046) than for naive sampling (.0033). It clearly illustrates the benefits obtained from being able to design the sampling rates in an optimal manner, which can be responsive to the objective set forth by network engineers.

Finally, we look at some variations on the above setup. Figure 4.16 shows the performance of A-optimal sampling from a noise-free initialization under a *network-level* constraint alone. Here the improvement over naive sampling is 31.8%. Clearly there is improvement in performance over the corresponding performance under router-
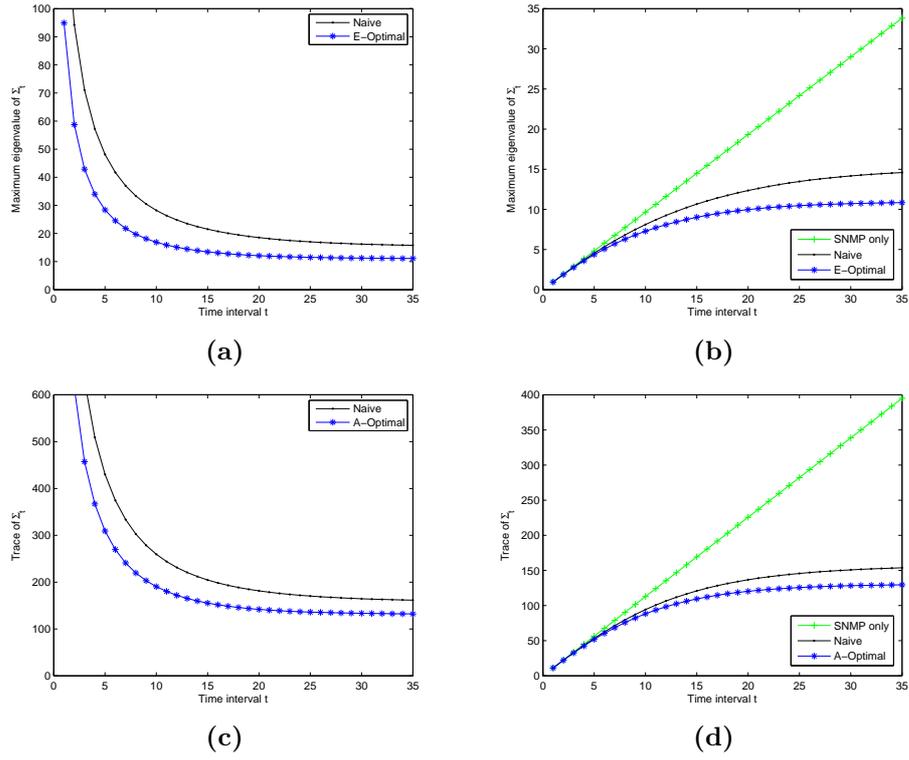
Figure 4.14: Performance of E-optimal (a)(b) and A- optimal sampling for noisy (a)(c) and noise-free (b)(d)initialization
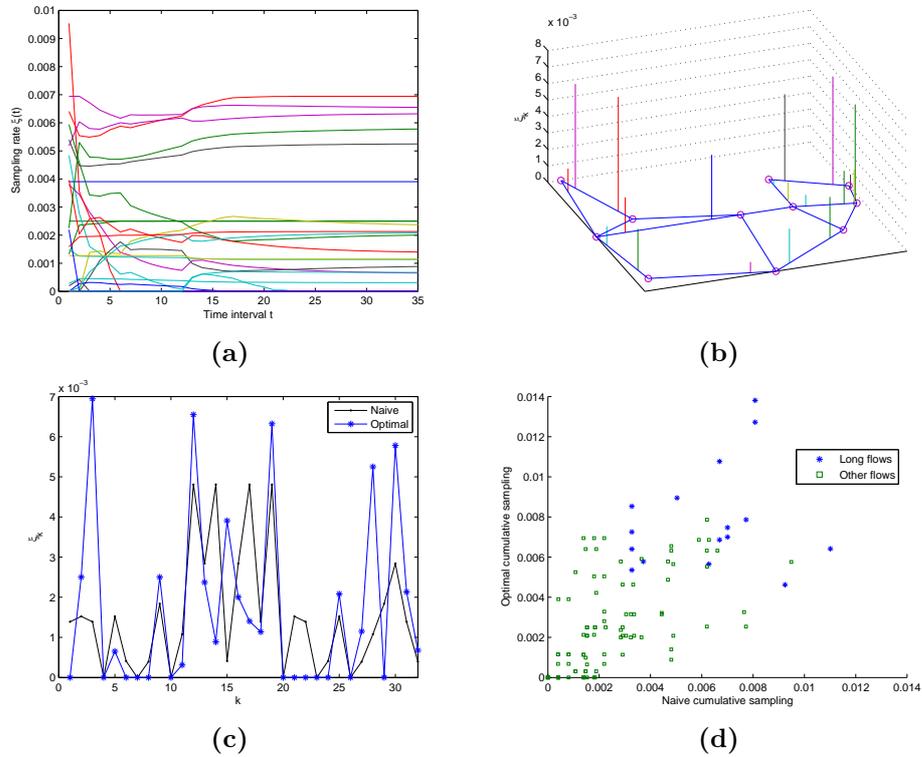

Figure 4.15: Spatio-temporal behavior of E-optimal ξ
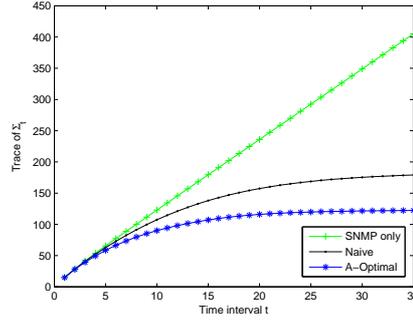
Figure 4.16: Performance of A-optimal sampling under network level constraint
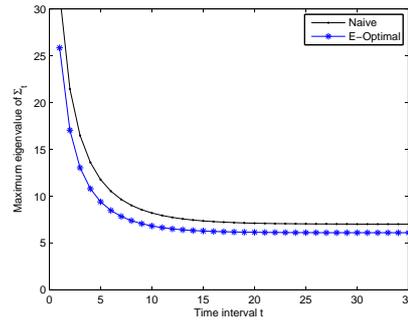
Figure 4.17: Performance of E-optimal sampling under linear cost.

level constraints (Figure 4.14(d)). Another variation is to replace the matrix R by its element-wise square-root. Thus, the cost of sampling on an interface goes up linearly with the number of flows traversing that interface. Notice that from a geometric point of view we have not changed the nature of the constraints. Figure 4.17 shows the performance of E-optimal sampling under such a constraint. Notice that compared to our usual cost (Figure 4.14(a)) the gain over the corresponding naive allocation is smaller. This is perhaps not surprising since under linear cost the gains from sampling on an interface are roughly equal to the cost and hence any shift along the surface $\{\xi : R\xi = b\}$ would have small effect on the objective function.

Effects of Misspecification

The state transition equation (4.14) and observation equation (4.13) necessarily suffer from model misspecification and parameter inaccuracies. Note that if $x(t)$ has

a constant covariance, then (4.14) can be true only if the spectral radius of C is less than 1 [34]. However, in this case, the steady state mean of $x(t)$ would be 0. This model (4.14) works well in practice [41] possibly because the misspecification in the conditional distribution of $x(t)$ given $x(t-1)$ is small. The misspecification in the observation equation (4.13) may potentially be more serious. Firstly, the sampled data $z(t)$ are not distributed normally given $x(t)$ to begin with. Thus, Kalman filtering is not optimal. Second, it is clear from the observations in Section 4.3.2 that the true value of the covariance of the measurement error, $\Psi(t, \cdot)$, can not be known with high precision and one would use an approximation for both filtering and design optimization. We study the effects of such misspecification in the following.

We start by describing the model used for data generation. Assume,

$$x_i(t) = \text{ROUND}(aV_i(t) + \cdots + aV_i(t+n))$$

where $V_i(t)$ are independent Poisson random variables with mean λ_i. Here $\text{ROUND}(\cdot)$ is the round off function. Variables a and λ_i are chosen to ensure $E[x_i(t)] = E[(x_i(t) - x_i(t-1))^2] = \mu_i$, where μ_i is the required mean. Thus, we again ensure that the variance of innovations for a flow are of the same magnitude as the volume of the flow, as seen in empirical studies [30]. The above can be solved for $a = (n+1)/2$ and $\lambda_i = 2\mu_i/(n+1)^2$ Further, the above model implies that the corresponding fitted model for the state transition equation (4.14) would have $C = I$ and $W = \text{Diag}(\mu)$. These are the values we use for both design optimization and filtering. For most of the results presented here it assumed that $n = 50$, while the μ_i's were chosen from a uniform distribution with support between 9000 and 11000.

The sampled measurements were generated as follows. The SNMP measurements $y(t)$ were generated from a normal distribution with mean $x(t)$ and covariance $\sigma^2 I$ with $\sigma^2 = 100$. The i-th sampled measurement, $z_i(t)$ was chosen to be $1/\xi_k(t)$ times

a binomial random variable with parameters $x_j(t)$ (the corresponding flow volume) and $\xi_k(t)$ (the corresponding sampling rate), where $j = l(i)$ and $k = k(i)$.

Finally, for the purpose of design optimization and filtering it is assumed that $\Psi(t, \xi) = \sum_k \xi_k \Psi_k$, where Ψ_k are diagonal matrices with

$$(4.22) \qquad [\Psi_k]_{ii} = \begin{cases} 1/\hat{x}_{l(i)}(t-1) & \text{if } k = k(i) \\ 0 & \text{otherwise} \end{cases}$$

where $\hat{x}(t-1)$ is the Kalman filter estimate of $x(t-1)$. This choice works well if $\hat{x}(t-1)$ is a good estimate of $x(t)$. However, if $\hat{x}(t-1)$ is based on a few intervals worth of data, as in the first few time intervals after a noisy initialization, it may be extremely noisy and this can potentially lead to instabilities. In our simulations rare instances of explosive growth in estimation error were observed (these instances were extremely infrequent, occurring at most thrice in 200×2 realizations of sampled data, corresponding to optimal and naive sampling of the 200 realizations of $x(t)$. We have dropped such cases from our results and indicated accordingly).

In this setup, two sets of initial conditions were considered as well. The noise-free initialization assumes $\hat{x}(0) = x(0)$ and $\Sigma(0) = 0$, while the noisy one assumes

$$\hat{x}(0) = (J'J)^{-1} J' \begin{pmatrix} Y(0) \\ Z(0) \end{pmatrix}$$

where $J' = (A', L')$ and

$$\Sigma(0) = (J'J)^{-1} J' \begin{pmatrix} \sigma^2 I & 0 \\ 0 & \Psi^{-1}(0, \xi(0)) \end{pmatrix} J(J'J)^{-1}$$

We define $\hat{x}(-1) \equiv \hat{X}(0)$ for the purpose of evaluating $\Psi(0, \cdot)$. Also, recall that $\xi(0)$ is the naive allocation, which was used to generate $Z(0)$.

For a given optimization setup we generate μ as described and simulated up to 200 realizations of $x(t)$. We calculate $\hat{x}(t)$ for different versions of the Kalman filter
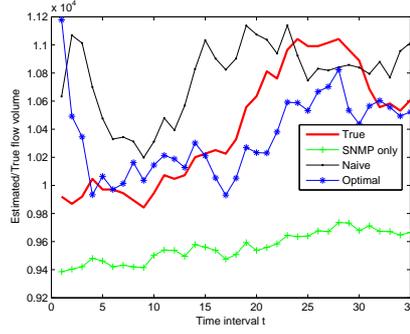
Figure 4.18: Sample trajectory of true and estimated flow volume

operating on different data depending on the respective sampling rates. Figure 4.18 depicts one realization of the trajectories of the true flow volume and estimated flow volumes for a particular flow. Thus, we calculate the MSE for all flows, averaging over all realizations. Given the discussion in Section 4.4.2, we study the performance of the maximum MSE among all flows for E-optimality and the mean of MSEs of all flows for A-optimality.

Figure 4.19 depicts the performance of optimal sampling for different initializations and different target sets of flows (for (a) and (c) two and one simulation, respectively, out of 200 were dropped due to numerical instability). Comparing Figures 4.19(a)(b) to Figures 4.11 (a)(b), and Figures 4.19(c)(d) to Figures 4.14 (a)(b) we observe that the relative performance of different Kalman filters under misspecification is very similar to that of the corresponding set of Kalman Filters when there is no misspecification. Specifically, we again observe that when interest is restricted to the subset of long flows there is optimal sampling performs considerably better than naive sampling.

Finally, we look at some variations on the above settings. Recall that $\Sigma = \mathrm{Diag}(\mu)$ and thus one needs knowledge of μ for both filtering and tracking purposes. We study the effect of incorrectly specified μ on the performance of the Kalman filter. Figure
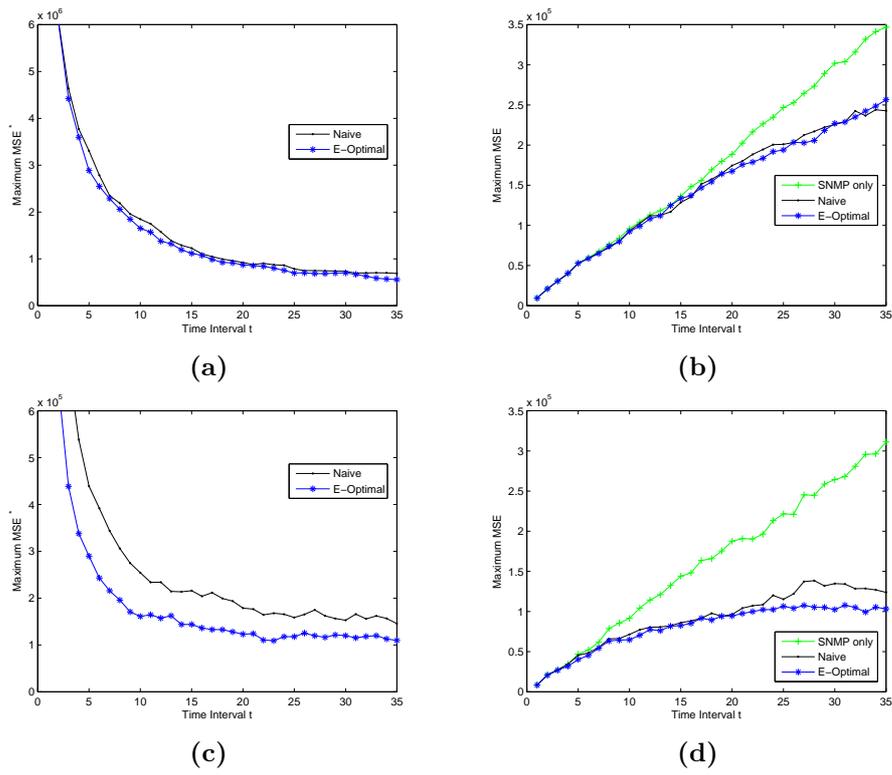
Figure 4.19: Performance of E-optimal for noisy (a)(c) and noise-free (b)(d) realizations. Interest is restricted to long flows in (c)(d).
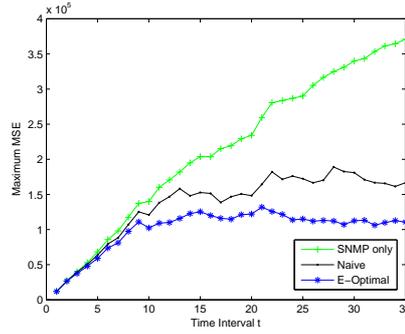
Figure 4.20: Performance under misspecification of μ
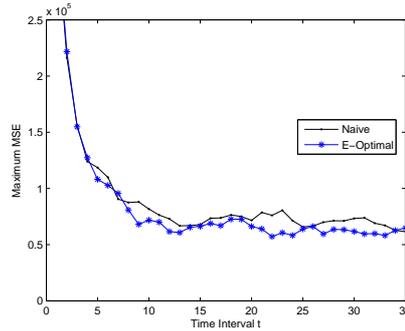
Figure 4.21: Performance of E-optimality under linear cost

4.20 shows the performance of respective Kalman filters when the presumed μ is still chosen as described earlier, but the true μ used for data generation has two arbitrary entries changed to 20000. Notice that for the above misspecification the relative improvement from optimal sampling over naive sampling is considerably larger than the corresponding improvement otherwise (Figure 4.19(d)). Figure 4.21 shows the performance of E-optimal sampling when R is replaced by its element-wise square-root. Given our observation from Figure 4.17 it is not surprising that the relative gain from optimal sampling over naive sampling is reduced.

4.5 Discussion and Future Work

The previous section extends the myopic design problem to state space models. The *specification* of steady state optimal design problem can be easily generalized to

such systems. Recall that these systems are described by a pair of equations [25]. The State Transition Equation can be written as:

$$x(t) = Cx(t-1) + w(t)$$

where $\text{Cov}(w(t)) = W$. Then Observation Equation can be written as:

$$y(t) = Lx(t) + \epsilon(t)$$

Assume $\text{Cov}(\epsilon(t)) = \Psi(\xi)^{-1}$ where $\Psi(\cdot)$ is a linear function and ξ is the value of design variables.

For the above dynamical system a Kalman Filter can be used to iteratively compute $E[x(t)|y(t), y(t-1), \cdots]$. Let the steady state estimation error covariance be $\Sigma = \tilde{M}^{-1}$. Then $\tilde{M}$ satisfies the Algebraic Riccatti Equation:

$$(4.23) \qquad \tilde{M} = (C\tilde{M}^{-1}C' + W)^{-1} + L'\Psi(\xi)L$$

Such equations have no analytic solution in general.

The Steady State Optimal Design problem can now be defined as:

$$\arg \max_{R\xi \leq b} f(\tilde{M}^{-1})$$

where $f(\cdot)$ is an appropriate scalarization of the covariance matrix [17]. An interesting open problem is to solve the above optimization efficiently in the absence of an analytic solution to (4.23). The sensor placement problem in control system literature [1] is equivalent though not identical. The newton-type algorithm proposed in [1] for this problem requires the solution of the Algebraic Riccati Equation at each time step. It would be desirable to develop more efficient algorithms.

In summary, we have shown that steady state E-optimal design for random walks is a second order cone program. We have shown numerically that the performance

of the Kalman filter can be significantly improved by optimal experiment design. We investigated numerically the extension of myopic design to instances of state space models. The linear state space model is of general interest and one would like to investigate the steady state optimal design problem described above. From a practical point of view it would be useful to extend these ideas to non-linear filtering.

CHAPTER V

Concluding Remarks

The common theme running through most of the present work is utilizing the dependence structure in flow volumes distribution for estimation purposes. It is quite likely that these ideas would be useful in the context of other networks like road networks or supply chain networks. The results on identifiability and steady state optimal design seem to be of general statistical interest. The work on dual modality tomography highlights some of the challenges in, and possible solutions for, estimation of flow volumes in realistic computer networks. Future work has been indicated at the end of chapters as appropriate. We conclude here with some general remarks in various related directions.

Recall that most of our identifiability results concern identifiability up to mean. When interest is in estimating flow volume means, as is often the case, some kind of mean-variance relationship typically needs to be assumed for regularization. A possible exception could be assuming a compound model as noted in Section 2.6. Estimation for compound models is challenging [5] even when there is no aggregation of the kind under consideration. However, this is certainly an interesting direction for future exploration. On the other hand, often times the interest may be in estimating variances of flow volumes, e.g. to use as input for anomaly detection

or filtering. Sometimes second moments like temporal auto-correlations or inter-modality covariances are of independent interest. Even in these cases identifiability of a certain model may not be enough to recommend the use of such a model for straightforward estimation when sample size is small. There have been several exciting developments in the field of regularized covariance estimation [3] and it would be interesting to use some of those ideas for estimating covariances of flow volumes under identifiable models.

Of course, if sampled data collection infrastructure is widely deployed in a network then questions of identifiability of flow volumes distribution from cumulative measurements are no longer crucial. In this case sampling rate design and schemes for merging sampled and SNMP-like data are more important from a practical point of view. Finally, an area of enduring promise and some controversy is probabilistic modeling of individual connections in a computer network to explain aggregate traffic characteristics [22].

BIBLIOGRAPHY

BIBLIOGRAPHY

[1] Ami Arbel. Sensor placement in optimal filtering and smoothing problems. *IEEE Transactions on Automatic Control*, 27(1):94–98, 1982.

[2] P. Barford, J. Kline, D. Plonka, and A. Ron. A signal analysis of network traffic anomalies. In *IMW '02: Proceedings of the 2nd ACM SIGCOMM Workshop on Internet measurment*, pages 71–82, New York, NY, USA, 2002. ACM Press.

[3] Peter J. Bickel and Elizaveta Levina. Covariance regularization by thresholding. *The Annals of Statistics*, 36(6):2577–2604, 2008.

[4] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

[5] Boris Buchmann and Rudolf Grübel. Decompounding: An estimation problem for poisson random sums. *The Annals of Statistics*, 31(4):1054–1074, 2003.

[6] J. Cao, D. Davis, S. Wiel, and B. Yu. Time-varying network tomography : Router link data. *Journal of the American Statistical Association*, 95:1063–75, 2000.

[7] Aiyou Chen, Jin Cao, and Tian Bu. Network tomography: Identifiability and fourier domain estimation. *INFOCOM 2007. 26th IEEE International Conference on Computer Communications. IEEE*, pages 1875–1883, May 2007.

[8] Kenjiro Cho. WIDE-TRANSIT 150 Megabit Ethernet Trace 2007-01-09 (Anonymized) (collection). http://mawi.wide.ad.jp/mawi/samplepoint-F/20070109/.

[9] B. Choi and S. Bhattacharrya. On the accuracy and overhead of Cisco sampled netflow. In *ACM Sigmetrics Workshop on Large-Scale Network Inference (LSNI)*, Banff, Canada, June 2005.

[10] Yuan Shih Chow and Henry Teicher. *Probability Theory*. Springer-Verlag New York, Inc., 1997.

[11] N. Duffield, C. Lund, and M. Thorup. Properties and prediction of flow statistics from sampled packet streams. In *IMW '02: Proceedings of the 2nd ACM SIGCOMM Workshop on Internet measurment*, pages 159–171, New York, NY, USA, 2002. ACM Press.

[12] Nick Duffield, Carsten Lund, and Mikkel Thorup. Flow sampling under hard resource constraints. In *SIGMETRICS '04/Performance '04: Proceedings of the joint international conference on Measurement and modeling of computer systems*, pages 85–96, New York, NY, USA, 2004. ACM.

[13] Nick Duffield, Carsten Lund, and Mikkel Thorup. Optimal combination of sampled network measurements. In *IMC'05: Proceedings of the Internet Measurement Conference 2005*, pages 8–8, Berkeley, CA, USA, 2005. USENIX Association.

[14] Nick G. Duffield. Sampling for passive internet measurement: A review. *Statistical Science*, 19(3):472–498, 2004.

[15] Nick G. Duffield, Carsten Lund, and Mikkel Thorup. Learn more, sample less: control of volume and variance in network measurement. *IEEE Transactions on Information Theory*, 51(5):1756–1775, 2005.

[16] V. Erramilli, M. Crovella, and N. Taft. An independent-connection model for traffic matrices. In *IMC '06: Proceedings of the 6th ACM SIGCOMM on Internet measurement*, pages 251–256, New York, NY, USA, 2006. ACM Press.

[17] V. V. Fedorov and P. Hackl. *Model-Oriented Design of Experiments*. Springer-Verlag, 1997.

[18] I. Ford, D. M. Titterington, and C. P. Kitsos. Recent advances in nonlinear experimental design. *Technometrics*, 31(1):49–60, 1989.

[19] Roland Gautier and Luc Pronzato. Sequential design and active control. In *New Developments and Applications in Experimental Design*, pages 138–151. Institute of Mathematical Statistics, 1998.

[20] G. C. Goodwin. *Dynamic System Identification. Experiment Design and Data Analysis*. Academic Press Inc, 1977.

[21] C. Gourieroux, A. Monfort, and A. Trognon. Pseudo maximum likelihood methods: Theory. *Econometrica*, 52(3):681–700, 1984.

[22] Matthias Grossglauser and Jean-Chrysostome Bolot. On the relevance of long-range dependence in network-traffic. *IEEE/ACM Transactions on Networking*, 7(5):629–640, 1999.

[23] Sprint Academic Research Group. IP Data Analysis. `http://research.sprintlabs.com/packstat/packetoverview.php`.

[24] C. D. Hardin. On the spectral representation of symmetric stable processes. *Journal of Multivariate Analysis*, 12(3):385–401, 1982.

[25] A. C. Harvey. *Forecasting, structural time series models and the Kalman filter*. Cambridge University Press, 1990.

[26] A. M. Kagan, Yu. V. Linnik, and C. Radhakrishna Rao. *Characterization Problems in Mathematical Statistics*. John Wiley & Sons, Inc., 1973.

[27] A. Lakhina, M. Crovella, and C. Diot. Diagnosing network-wide traffic anomalies. In *SIGCOMM '04: Proceedings of the 2004 conference on Applications, technologies, architectures, and protocols for computer communications*, pages 219–230, New York, NY, USA, 2004. ACM Press.

[28] E. Lawrence, G. Michailidis, V.N. Nair, and B. Xi. Network tomography: a review and recent developments. In *Frontiers in Statistics*, pages 345–364. Fan and H. Koul (eds), Imperial College Press, London, 2006.

[29] Jon Lee. Semidefinite programming in experimental design. In *H. Wolkowicz, R. Saigal and L. Vandenberghe, editors, "Handbook of Semidefinite Programming", International Series in Operations Research and Management Science, Vol. 27*. Kluwer, 2000.

[30] G. Liang, N. Taft, and B. Yu. A fast lightweight approach to origin-destination IP traffic estimation using partial measurements. *IEEE Transactions on Information Theory*, 52(6):2634–2648, 2006.

[31] G. Liang and B. Yu. Pseudo likelihood estimation in network tomography. In *Proceedings of the 22nd Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM 2003)*, San Francisco, CA, 2003.

[32] Y. Liu, D. Towsley, T. Ye, and J. C. Bolot. An information-theoretic approach to network monitoring and measurement. In *IMC '05, Internet Measurement Conference*, New Orleans, LA, USA, October 2005.

[33] Miguel Sousa Lobo, Lieven Vandenberghe, Stephen Boyd, and Herv Lebret. Applications of second-order cone programming. *Linear Algebra and its Applications*, 284(1-3):193 – 228, 1998. International Linear Algebra Society (ILAS) Symposium on Fast Algorithms for Control, Signals and Image Processing.

[34] H. Lutkepohl. *Introduction to multiple time series analysis*. Springer-Verlag, New York, 1993.

[35] A. Medina, N. Taft, K. Salamatian, S. Bhattacharyya, and C. Diot. Traffic matrix estimation: existing techniques and new directions. In *Proceedings of the 2002 conference on Applications, technologies, architectures, and protocols for computer communications*. ACM Press, 2002.

[36] Kihong Park and Walter Willinger. *Self-Similar Network Traffic and Performance Evaluation*. John Wiley and Sons, Inc., 2000.

[37] L.L. Peterson and B.C. Davie. *Computer Networks: A Systems Approach*. Morgan Kaufmann, San Francisco, 2003.

[38] F. Pukelsheim. *Optimal Design of Experiments*. John Wiley & Sons, Inc., 1993.

[39] David Rincón, Matthew Roughan, and Walter Willinger. Towards a meaningful mra of traffic matrices. In *IMC '08: Proceedings of the 8th ACM SIGCOMM conference on Internet measurement*, pages 331–336, New York, NY, USA, 2008. ACM.

[40] Harsh Singhal and George Michailidis. Identifiability of flow distributions from link measurements with applications to computer networks. *Inverse Problems*, 23(5):1821–1849, 2007.

[41] A. Soule, A. Lakhina, N. Taft, K. Papagiannaki, K. Salamatian, A. Nucci, M. Crovella, and C. Diot. Traffic matrices: Balancing measurements, inference and modeling. In *Proceedings of the joint international conference on Measurement and modeling of computer systems*. ACM Press, 2005.

[42] A. Soule, A. Nucci, R. Cruz, E. Leonardi, and N. Taft. How to identify and estimate the largest traffic matrix elements in a dynamic environment. In *SIGMETRICS '04/Performance '04: Proceedings of the joint international conference on Measurement and modeling of computer systems*, pages 73–84, New York, NY, USA, 2004. ACM Press.

[43] M. Taniguchi and Y. Kakizawa. *Asymptotic Theory of Statistical Inference for Time Series*. Springer-Verlag New York, Inc., 2000.

[44] K. Thompson, G.J. Miller, and R. Wilder. Wide-area internet traffic patterns and characteristics. *Network, IEEE*, 11(6):10–23, Nov/Dec 1997.

[45] D. M. Titterington. Aspects of optimal design in dynamic systems. *Technometrics*, 22(3):287–299, 1980.

[46] S. Uhlig, B. Quoitin, S. Balon, and J. Lepropre. Providing public intradomain traffic matrices to the research community. *ACM SIGCOMM Computer Communication Review*, 36(1), January 2006.

[47] Y. Vardi. Network tomography: Estimating source-destination traffic intensities from link data. *Journal of the American Statistical Association*, 91:365–377, 1996.

[48] W. Willinger, V. Paxson, and M. S. Taqqu. *Self-similarity and heavy tails: structural modeling of network traffic*. Birkhauser Boston Inc., Cambridge, MA, USA, 1998.

[49] Lili Yang and George Michailidis. Estimation of flow lengths from sampled data. In *Proceedings of IEEE GLOBECOM*, 2006.

[50] Lili Yang and George Michailidis. Sample based estimation of network traffic characteristics. In *Proceedings of IEEE INFOCOM*, 2007.

[51] Y. Zhang, M. Roughan, N. G. Duffield, and A. G. Greenberg. Fast accurate computation of large-scale IP traffic matrices from link loads. In *Proceedings of the International Conference on Measurements and Modeling of Computer Systems, SIGMETRICS 2003*, pages 206–217, 2003.

[52] Qi Zhao. *Towards Ideal Network Traffic Measurement: A Statistical Algorithmic Approach.* PhD thesis, College of Computing, Georgia Institute of Technology, Atlanta, Georgia, USA, 2007.

[53] Qi Zhao, Zihui Ge, Jia Wang, and Jun Xu. Robust traffic matrix estimation with imperfect information: making use of multiple data sources. In *SIGMETRICS '06/Performance '06: Proceedings of the joint international conference on Measurement and modeling of computer systems*, pages 133–144, New York, NY, USA, 2006. ACM.