

Identifiability of flow distributions from link measurements with applications to computer networks

Harsh Singhal and George Michailidis

Department of Statistics, The University of Michigan, MI, USA

Received 21 February 2007, in final form 13 July 2007

Published 22 August 2007

Online at stacks.iop.org/IP/23/1821

Abstract

We study the problem of identifiability of distributions of flows on a graph from aggregate measurements collected on its edges. This is a canonical example of a statistical inverse problem motivated by recent developments in computer networks. In this paper (i) we introduce a number of models for multi-modal data that capture their spatio-temporal correlation, (ii) provide sufficient conditions for the identifiability of n th order cumulants and also for a special class of heavy tailed distributions. Further, we investigate conditions on network routing for the flows that prove sufficient for identifiability of their distributions (up to mean). Finally, we extend our results to directed acyclic graphs and discuss some open problems.

(Some figures in this article are in colour only in the electronic version)

1. Introduction

An increasing variety of network data is available from modern computer networks. These data differ in their granularity, accuracy, volume and delay [13]. An important area of interest is to collect aggregate data on the network's edges in order to infer origin–destination traffic volumes. This is a canonical example of a statistical inverse problem and has applications in network capacity planning and fault diagnosis, traffic forecasting and provisioning and routing protocol configuration [11, 17]. Another application comes from road networks, where it is of interest to decompose aggregate traffic volumes collected from loop detectors to the corresponding origin–destination traffic flows [8]. The objective of this study is to obtain conditions under which network-wide traffic volumes (flows) can be *estimated* from aggregate (limited resolution) data.

A computer network 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 payload is expressed in bytes, while on road networks, the traffic is carried on vehicles. The volume of traffic measured on a link may refer to either the number of packets and/or the number of bytes in computer networks, and such data for a particular time interval—typically of the order of a couple of minutes—are available through queries using the simple network management protocol (SNMP) [15]. 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.

Modelling and estimation of flow volumes in computer networks has attracted a lot of attention recently and has implications at all time scales. At short time scales (less than 1 h) network service providers are interested in anomaly detection [1, 9], where sudden change in the flow volume distribution may indicate a malicious attack or equipment failure. At moderate time-scales (1 h to 1 week) estimates of flow volumes are useful for traffic engineering tasks such as load balancing and routing protocol configuration [17]. At larger time-scales it is of interest to monitor the changing nature of network traffic. For example, it is fairly well documented that the bulk of network traffic is moving away from connection-oriented HTTP to connection-less peer-to-peer traffic [6].

The aggregate nature of link load (SNMP) data leads to an inverse problem. The estimability and usability of any model of flow volumes requires that it should be uniquely identifiable. If two distinct sets of values for the free parameters of the model lead to the same distribution of observable data then the model is unidentifiable from the data. In this paper, we introduce models that can capture the spatio-temporal dependence observed in flow volumes and can be shown to be identifiable (up to mean) under reasonable conditions.

1.1. Literature review

The problem of estimating flow volumes from aggregate link traffic measurements was introduced by Vardi [19], where the term *network tomography* was used to describe a particular class of statistical inverse problems. In [19], the focus was on estimating mean flow volumes under the assumption that they are Poisson distributed with parameter λ_j for flow j . Estimability (identifiability) of λ_j was proved using the parametric form of the density of a Poisson random variable. The proof proceeds through writing the exact non-zero probability associated with observing a certain vector of link measurements. Due to the Poisson assumption, enough such non-zero probability events can be constructed to show identifiability of all λ_j . To obtain a solution for the λ_j , maximum likelihood estimators based on the normal approximation of the Poisson distribution, as well as based on the method of moments were also proposed. A mean–variance relationship that generates a full rank system of linear equations was also used in [3], where flow volumes were modelled as being normally distributed with flow variances proportional to their means. The proportionality assumption leads to identifiability of means through identifiability of variances; however, it has been found that such models do not estimate accurately enough the distribution of X in large high-speed computer networks [14].

Another class of models imposes other types of constraints for obtaining identifiable (estimable) solutions. For example, gravity models [20] assume that flow f_j between nodes n_k and n_ℓ is proportional to the total amount of traffic departing node n_k and the total amount of traffic entering node n_ℓ . This model assumes complete independence between source and destination nodes that also tends to be violated in backbone networks [11]. The above assumption introduces enough constraints to regularize the problem for a unique solution. A Kalman filter based approach suggested in [16] provides best linear estimates of flow volumes assuming a specific temporal dependence structure with known parameters. Recently, a

sufficient condition for identifiability of second- and higher order cumulants (see section 4) of the distributions of flow volumes was established in [4], under strong assumptions regarding independence of the flows. Further, an estimator based on the characteristic function of the aggregate data was proposed. The ideas developed in some of the above papers have been employed in [17] and [12] to develop practical traffic volume estimators for continuous monitoring of real networks.

1.2. Problem formulation and basic notation

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 $f_j, j = 1, \dots, J$, 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 $P = (n_1, \dots, n_{L_P+1}), (n_i, n_{i+1}) \in E$, for $i = 1, \dots, L_P$. We say $e_i = (n_i, n_{i+1}) \in P, i = 1, \dots, L_P$, 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 $\{\mathcal{P}(j), w_j(P)\}$ determines 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 j th flow that traverses the i th 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 ($O(n)$). 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 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

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




Figure 1. Example topology.




Figure 2. Aggregate volume measurements.

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 $X_{(i,j)}$. Here the number of links $L = 10$ and the number of flows $J = 30$.

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

$$Y_1 = X_1 + X_2, \quad Y_2 = X_2 + X_3.$$

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:

$$v_y \equiv \begin{pmatrix} \text{Var}(Y_1) \\ \text{Var}(Y_2) \\ \text{Cov}(Y_1, Y_2) \end{pmatrix} = \begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & 1 \\ 0 & 1 & 0 \end{pmatrix} \begin{pmatrix} \text{Var}(X_1) \\ \text{Var}(X_2) \\ \text{Var}(X_3) \end{pmatrix} \equiv Bv_x.$$

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 paper, 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

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

Thus, we define the matrix function $B(R, \tilde{R})$, with R and \tilde{R} both $n \times m$ as follows. Let function $g(i_1, i_2) : \{1, \dots, n\}^2 \rightarrow \{1, \dots, n^2\}$ be an ordering of the pairs $(i_1, i_2) \in \{1, \dots, n\}^2$. Then

define $B(R, \tilde{R})_{g(i_1, i_2), j} = R_{i_1, j} \tilde{R}_{i_2, j}$. For example, for $R = [r_{ij}]_{2 \times 3}$, $\tilde{R} = [\tilde{r}_{ij}]_{2 \times 3}$ we get

$$B(R, \tilde{R}) = \begin{pmatrix} r_{11}\tilde{r}_{11} & r_{12}\tilde{r}_{12} & r_{13}\tilde{r}_{13} \\ r_{11}\tilde{r}_{21} & r_{12}\tilde{r}_{22} & r_{13}\tilde{r}_{23} \\ r_{21}\tilde{r}_{11} & r_{22}\tilde{r}_{12} & r_{23}\tilde{r}_{13} \\ r_{21}\tilde{r}_{21} & r_{22}\tilde{r}_{22} & r_{23}\tilde{r}_{23} \end{pmatrix}.$$

Matrices of the form $B(A, A)$ play a crucial role in subsequent developments. Note that rows of $B(A, A)$ are element-wise products of rows in A and each row in $B(A, A)$ indicates common flows between a pair of links. It can therefore be seen that ‘identifiability’ of variances of the X_i is related to the matrix $B(A, A)$ being full rank when the X_i 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:

$$Y(t) = AX(t), \quad t = 1, \dots \tag{1}$$

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, (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 modelled 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 paper, we present certain conditions on 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. Note 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. $\mathcal{L}(X) \in \mathcal{M}$. Then, identifiability is formally defined as follows.

Definition 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 $Y_2 = AX_2$, $\mathcal{L}(X_1), \mathcal{L}(X_2) \in \mathcal{M}$, $Y_1 \stackrel{d}{=} Y_2$ implies that $X_1 \stackrel{d}{=} X_2 + c$ for some constant $c \in \mathbb{R}^J$.*

1.3. Main contributions

In this paper, we present a broad framework that models dependence between flow volumes that are present in computer networks. This framework is based on a *latent variable* model which represents a random vector as a product of two terms, one that captures the dependence structure and another that captures other parameters. This framework also allows for various types of measurements on each flow and thus accommodates different modalities and temporal correlations. As a special case of this framework, we investigate the independent connections model for data networks, which takes into account the most important dependence between flow volumes and yet is simple enough that identifiability (up to mean) can be guaranteed for reasonably large classes of networks. Next, conditions for identifiability are derived for the general model, by separately handling variances and the remaining parameters of the distribution to get around the inherent nonlinearity of a latent variable model. Some

elements of these ideas have been previously explored for a single type of measurement with no spatial or temporal dependence structure in [3, 4, 17]. It is worth noting that using a latent variable model to account for dependence, makes the conditions for identifiability significantly more involved. Further, identifiability is established for stable (heavy tailed) distributed flow volumes, by leveraging the theory of minimal representations to establish a connection between identifiability in the stable and the Gaussian case, since in the former second moments (variances) do not exist.

The sufficient condition for identifiability under the proposed general model involves a variant of the B matrix being of full rank. We subsequently show that identifiability is guaranteed for a reasonably large class of networks for the independent connections model, by stating explicit assumptions on the network routing and edge weights. This result is of significant practical importance for designing networks with desired properties. The proof borrows ideas from [17], but is significantly more involved due to the presence of a dependence structure amongst flows. Finally, stronger results regarding identifiability can be obtained for networks represented by directed acyclic graphs and/or when two-dimensional measurements can be gathered.

The remainder of the paper is organized as follows: In section 2, we present the general distributional modelling framework under which we derive our results. In section 3, we specialize the general framework to the independent connections model which is useful for real data networks. In section 4, we derive conditions for identifiability of various distributional parameters under the general model. In section 5, we show that under reasonable assumptions on the routing scheme, the required conditions for identifiability are satisfied for the independent connections model, while in section 6 we look at certain special cases of interest for computer networks. Finally, in section 7 we present a discussion on some of the results and point to open problems.

2. A general modelling framework

In the network tomography literature, flow volumes are usually modelled as being independent across time and space, while packet and byte volumes have never been considered simultaneously. We provide next a framework to capture the most interesting dependences in flow volumes. Assume that there are measurements of type $1, \dots, K$ on each flow. These different measurement types may correspond to different ‘modalities’, such as packet volume or byte volume, different time lags or both. Let $X^{(k)}$ denote the flow measurements of type k . Each $X^{(k)}$ is a vector of length J , where J is the number of flows. Thus, $X_j^{(k)}$ denotes the k th measurement on the j th flow, $1 \leq k \leq K$ and $1 \leq j \leq J$. For type k observations we have the following model

$$Y^{(k)} = AX^{(k)} \quad (2)$$

for $k = 1, \dots, K$. Stacking together all types of measurements, we get

$$Y \equiv \begin{pmatrix} Y^{(1)} \\ Y^{(2)} \\ \vdots \\ Y^{(K)} \end{pmatrix} = \begin{pmatrix} A & 0 & \dots & 0 \\ 0 & A & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & A \end{pmatrix} \begin{pmatrix} X^{(1)} \\ X^{(2)} \\ \vdots \\ X^{(K)} \end{pmatrix} \equiv \bar{A}X, \quad (3)$$

where Y is a vector of length LK , \bar{A} is a $LK \times JK$ matrix and X is a vector of length JK . Let $\Sigma_X \equiv \text{Cov}(X)$ be the $JK \times JK$ covariance matrix of the flows. Define $\mathcal{Q} \equiv \{1, \dots, JK\}^2 = \{(q_1, q_2) : q_1, q_2 \in \{1, \dots, JK\}\}$, the set of pairs of indices in X .

Now, $q \in \mathcal{Q}$ can be used to index the covariance matrix, $[\Sigma_X]_q \equiv [\Sigma_X]_{q_1, q_2}$. The dependence between elements of X can be specified in one of the following ways:

- (1) *Covariance model.* The covariance model will be specified through $\mathcal{Q}_M \subset \mathcal{Q}$ such that $(q, q) \in \mathcal{Q}_M$ for all $q \in \{1, \dots, JK\}$ and if $(q_1, q_2) \in \mathcal{Q}_M$ then $(q_2, q_1) \in \mathcal{Q}_M$. Note that \mathcal{Q}_M gives the set of admissible non-zero covariance matrices, i.e. if $q \in \mathcal{Q} - \mathcal{Q}_M$, then $[\Sigma_X]_q \equiv 0$. Restricting the number of non-zero covariances reduces the number of parameters which is important for identifiability purposes.
- (2) *Latent variable model.* Let $Z \in \mathbb{R}^{JK}$ be a random vector of latent variables of length JK , where its elements are assumed to be independent with identical (say unit) variances. Further, assume that

$$X = CZ \equiv \begin{pmatrix} \tilde{C}_{11} & 0 & \cdots & 0 \\ \tilde{C}_{21} & \tilde{C}_{22} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ \tilde{C}_{K1} & \tilde{C}_{K2} & \cdots & \tilde{C}_{KK} \end{pmatrix} \begin{pmatrix} \tilde{Z}_1 \\ \tilde{Z}_2 \\ \vdots \\ \tilde{Z}_K \end{pmatrix}, \quad (4)$$

where C (and thus $\tilde{C}_{kk}, k = 1, \dots, K$) are lower triangular matrices. The coefficient matrix C belongs to a set \mathcal{C} that is expected to capture the spatial dependence and defines the latent variable model space. Note that even though we refer to \mathcal{C} as our model, the distribution of Z is an important component of this modelling framework. We would concern ourselves with identifiability of parameters (other than variances) of the distribution of Z . Clearly identifiability of C and distributional parameters of Z would depend on whether \mathcal{C} is sufficiently small.

Examples of the models \mathcal{Q}_M and \mathcal{C} are given in section 3.

Note that Σ_X is a function of C and model definitions \mathcal{Q}_M and \mathcal{C} should be compatible in order to be used simultaneously. The latent variable model is more general, as it accommodates more distributional parameters for X than just variances and covariances. Now, let $\mathcal{S}(\mathcal{Q}_M)$ be the set of symmetric positive definite matrices Σ for which $\Sigma_q = 0$ for $q \notin \mathcal{Q}_M$. It can be immediately seen that $\mathcal{S}(\mathcal{Q}_M)$ is a convex cone in the n_M -dimensional space, where $n_M = |\{q \in \mathcal{Q}_M : q_1 \leq q_2\}|$. Due to the fact that the set of positive semidefinite matrices has a non-empty interior [2] and since $\mathcal{S}(\mathcal{Q}_M)$ is a projection of that set, $\mathcal{S}(\mathcal{Q}_M)$ has a non-empty interior in the n_M -dimensional space. In general, there may not be a one-to-one correspondence between elements of $\mathcal{S}(\mathcal{Q}_M)$ and \mathcal{C} . However, for the following specification of \mathcal{Q}_M and \mathcal{C} this correspondence follows immediately.

3. Independent connections model for data networks

We start by illustrating some important features of network traffic volume using a publicly available data set that motivates specific features considered in the proposed model (data are obtained from <http://www-dirt.cs.unc.edu/ts/>). The data are essentially a 4-variate time series where the four 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 s interval. We limit ourselves to the first 700 observations for illustration. The four time series are plotted in figure 3.

The temporal dependence is visible in the time series and can be more clearly seen through their auto-correlation functions shown in figure 4, which for a time series $x(t)$, at a given lag ℓ , corresponds to the observed correlation between $x(t)$ and $x(t - \ell)$ over all values of t . For each of the four time series considered, the auto-correlation functions are significantly




Figure 3. Byte (top) and packet (bottom) volume time series from the Abilene network.

greater than zero and decay with increasing lag. The simplest possible model for such time series is an auto-regressive model. Auto-regressive models were fitted to each of the four time series and the appropriate order of each model was chosen by the Akaike information criterion (AIC). The orders of the models were four for a forward byte volume, eight for a reverse byte volume and five for both forward and reverse packet volumes, respectively. The residuals from these models have the following correlation matrix (FB = forward byte, FP = forward packet, RB = reverse byte, RP = 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, it can be clearly seen the strong dependence between packet and byte volumes and between forward and reverse flows. Finally, the quantiles of the residuals obtained from the auto-regressive models were compared to those of a standard normal distribution. Figure 5 clearly shows that the quantiles of the observed error distribution are more extreme than those of a normal distribution indicating the presence of heavier tails.

The spatial correlation between network-wide flow volumes was subsequently examined. The data set used for this analysis (see [18]) gives byte volumes of all flows in a backbone




Figure 4. Auto-correlation functions for forward bytes (top left), forward packets (top right), reverse bytes (bottom left) and reverse packets (bottom right).

European educational network for each 15 min interval over a period of four 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 in the time series and to 76 flows that had no missing values and also comprised the top quarter of flows in terms of average traffic. The time series for each flow was smoothed by a spline model in order to deal with non-stationarities such as the well-known diurnal patterns [5, 9]. The residuals from the above step were used to fit an auto-regressive model, with AIC-based order selection, to account for temporal dependences. The pair-wise correlations obtained from the residual time series from the above analysis correspond to the spatial correlations. These pair-wise correlations can be divided into two sets: the forward–reverse ones and all the others. Figure 6 plots the observed densities of these two sets of correlation and it can clearly be seen that the forward–reverse correlations are stronger than the remaining ones. Although there is a bimodality 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, together with similar empirical findings from previous studies (see comments at the end of the section), we outline next some useful models for




Figure 5. Quantile–quantile plots for forward bytes (top left), forward packets (top right), reverse bytes (bottom left) and reverse packets (bottom right).

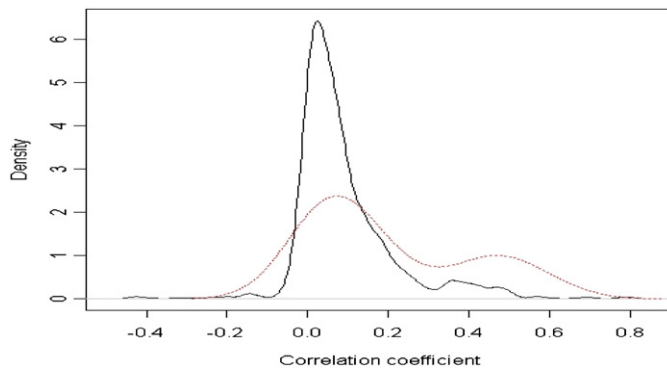


Figure 6. Densities of observed correlations: forward–reverse(dashed) and the remaining (solid).

computer networks. 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

$$Y^{(p)} = A_F X_F^{(p)} + A_R X_R^{(p)}. \tag{5}$$




Figure 7. Example topology.

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 7 comprised of four nodes and four 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 (5) 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 [15], 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 [5]. One of these flows is labelled 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

$$\Sigma_X = \begin{pmatrix} \text{Diag}(\delta_{FF}) & \text{Diag}(\delta_{FR}) \\ \text{Diag}(\delta_{FR}) & \text{Diag}(\delta_{RR}) \end{pmatrix}, \quad (6)$$

where each of δ_{FF} , δ_{FR} , δ_{RR} is a vector of length $J/2$ and component-wise they correspond to the variances of $X_F^{(p)}$, covariances of $X_R^{(p)}$ and $X_R^{(p)}$ and variances of $X_R^{(p)}$, respectively. Thus, Σ_X is a matrix of dimension $J \times J$. Using the framework developed in the previous section, this model can be represented as the following covariance model:

$$\mathcal{Q}_M = \{q \in \mathcal{Q} : q_1 = q_2 \bmod J/2\}. \quad (7)$$

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

$$X_{F_j}^{(p)} = c_{1j}Z_{1j}, \quad X_{R_j}^{(p)} = c_{2j}Z_{1j} + c_{3j}Z_{2j}.$$

with Z_{ij} independent normal with (possibly) different means and unit variances for all i, j . Note that two independent latent variables, Z_{1j} and Z_{2j} , are associated with flow pair j . The reverse flow in the 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

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

Hence, the latent variable model space is given by

$$\mathcal{C} = \{C : C_{i,i} > 0, C_{i,j} = 0 \text{ if } i < j \text{ or } i \neq j \text{ mod } J/2\}. \tag{8}$$

For \mathcal{Q}_M and \mathcal{C} as in (7) and (8), if $C \in \mathcal{C}$ then $\Sigma_X = \Sigma_X(C) \in \mathcal{S}(\mathcal{Q}_M)$ and (specifically) is positive definite. Conversely, if the covariance matrix $\Sigma_X \in \mathcal{S}(\mathcal{Q}_M)$ and (specifically) is positive definite, then there is a unique $C \in \mathcal{C}$ corresponding to it. This simply follows from the Cholesky decomposition of Σ_X , since Σ_X is assumed to be a symmetric positive definite matrix. Thus, taking C to be the lower triangular matrix from the factorization of Σ_X , we get $\text{Cov}(CZ) = CC\text{Cov}(Z)C' = CC' = \Sigma_X$. It should be noted that if X (or equivalently Z) is modelled as multi-variate normally distributed, then identifiability of Σ_X (or C) is equivalent to identifiability of the distribution of X up to mean.

The above model corresponds to having exactly one type of measurement ($K = 1$). Models (7) and (8) are reasonable for larger values of K as well. Different types of measurements on each flow ($K > 1$) 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 the section. Now consider another model, with dependence within flows and between packet counts and byte counts of the same flow:

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

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

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

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 $2J \times 2J$. If X is assumed to be multivariate normal the above model has a one-to-one correspondence with the following latent variable model:

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

$$\begin{aligned}
 X_{Fj}^{(b)} &= c_{4j}Z_{1j} + c_{5j}Z_{2j} + c_{6j}Z_{3j}, \\
 X_{Rj}^{(b)} &= c_{7j}Z_{1j} + c_{8j}Z_{2j} + c_{9j}Z_{3j} + c_{10j}Z_{4j},
 \end{aligned}$$

with Z_j independent normal with (possibly) different means and variances for all j . Here, four independent latent variables Z_{1j}, Z_{2j}, Z_{3j} and Z_{4j} are associated with the flow pair j . Dependence between different measurements of a flow pair is induced due to shared latent variables. Specifically Z_{1j} induces dependence between $X_{Fj}^{(p)}$, the forward packet volume and $X_{Rj}^{(p)}$, the reverse packet volume of flow pair j . Similarly Z_{1j} and Z_{2j} induce dependence between the forward and reverse volumes of flow pair j . Note that models (7) (8) also apply.

(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 modelled as follows:

$$\begin{aligned}
 Y^{(p)}(t) &= AX^{(p)}(t), \\
 Y^{(b)}(t) &= AX^{(b)}(t), \\
 X^{(p)}(t) &= \Phi_{p,1}X^{(p)}(t-1) + \dots + \Phi_{p,m}X^{(p)}(t-m) + \epsilon^{(p)}(t), \\
 X^{(b)}(t) &= \Phi_{b,1}X^{(b)}(t-1) + \dots + \Phi_{b,m}X^{(b)}(t-m) + \epsilon^{(b)}(t),
 \end{aligned}$$

where the various Φ_{\cdot} matrices contain the lag coefficients and $\epsilon^{(p)}(t), t = 1, \dots$, are i.i.d. mean 0 random vectors and so are $\epsilon^{(b)}(t), t = 1, \dots$. 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 autoregressive models, it is easy to verify the following:

$$\Sigma_{X,pp} = \Phi_p \Sigma_{X,pp} \Phi_p' + \Sigma_{pp}, \tag{9}$$

$$\Sigma_{X,pb} = \Phi_p \Sigma_{X,pp} \Phi_b' + \Sigma_{pb}, \tag{10}$$

$$\Sigma_{X,bp} = \Phi_b \Sigma_{X,bp} \Phi_p' + \Sigma_{bp}, \tag{11}$$

$$\Sigma_{X,bb} = \Phi_b \Sigma_{X,bb} \Phi_b' + \Sigma_{bb}, \tag{12}$$

$$\begin{aligned}
 \text{Cov}(X^p(t), X^{(p)}(t-l)) &= \Phi_p^l \Sigma_{X,pp} \equiv \Sigma_{X,pp}^l, \\
 \text{Cov}(X^{(p)}(t), X^{(b)}(t-l)) &= \Phi_p^l \Sigma_{X,pb} \equiv \Sigma_{X,pb}^l, \\
 \text{Cov}(X^{(b)}(t), X^{(p)}(t-l)) &= \Phi_b^l \Sigma_{X,bp} \equiv \Sigma_{X,bp}^l, \\
 \text{Cov}(X^{(b)}(t), X^{(b)}(t-l)) &= \Phi_b^l \Sigma_{X,bb} \equiv \Sigma_{X,bb}^l,
 \end{aligned}$$

where $\Sigma_{pp}, \Sigma_{pb}, \Sigma_{bp}$ 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 $\Sigma_{pp}, \Sigma_{pb}, \Sigma_{bp}$ and Σ_{bb} are block diagonal matrices of the form (6), 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 (9)–(12) imply that $\Sigma_{Xpp}, \Sigma_{Xpb}, \Sigma_{Xbp}$ and Σ_{Xbb} have the same block diagonal form given in (6). 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 (6) and the model (7) applies.

The independent connections model accommodates multi-modal measurements with spatio-temporal dependence together with non-parametric distributions. It is significantly more general than previous proposals in the literature that have usually focused on either the independent case [3, 4] or temporal dependence alone under parametric distribution models [16, 12]. One notable approach to modelling spatial dependence is the gravity model [20]

and its extensions [5]. These models assume that the distribution of total traffic originating at a node to various destinations happens in accordance with a specific (though empirically justified) pattern. Further, some stability in distribution patterns and availability of access link data are assumed for identifiability. We have avoided these assumptions while still modelling the most important spatial dependence and other interesting features of flow traffic data.

4. Identifiability results for the general model

In this section, we establish conditions for identifiability of the distribution of the flows $X_j, j = 1, \dots, J$, under the general dependence model introduced above. As discussed earlier, since the routing matrix A is not full rank, the means of the flows are in general not identifiable. Specifically, we establish the identifiability of second- and higher order cumulants provided they exist. Second-order cumulants of a random vector correspond to covariances, third order to skewness coefficients, etc. They can be obtained from the characteristic function, which for a random vector $T = (T_1, \dots, T_K)$ is given by

$$\psi_T(\theta) \equiv E[e^{i(\theta_1 T_1 + \dots + \theta_K T_K)}],$$

for $\theta = (\theta_1, \dots, \theta_K)$. The cumulant generating function for the same T is $\mathcal{K}_T(u) \equiv \log \psi_T(-iu)$. The cumulants of T are given by

$$\kappa_{r_1, \dots, r_K} = \frac{\partial^{r_1}}{\partial u_1^{r_1}} \cdots \frac{\partial^{r_K}}{\partial u_K^{r_K}} \mathcal{K}_T(u)|_{u=0}.$$

For a random vector T , for which the characteristic function is infinitely differentiable at 0, the distribution is determined completely by all the cumulants or equivalently by all partial derivatives of $\log \psi_T()$ evaluated at 0. The identifiability of order n cumulants of X under model \mathcal{M} is defined as follows.

Definition 2. *Order n cumulants (or other distributional parameters) of X are identifiable under model \mathcal{M} from observations $Y = AX$, if for $Y_1 = AX_1$ and $Y_2 = AX_2, \mathcal{L}(X_1), \mathcal{L}(X_2) \in \mathcal{M}, Y_1 \stackrel{d}{=} Y_2$, implies that the order n cumulants (or the corresponding distributional parameters) of X_1 and X_2 are the same.*

Our strategy would be to first identify second-order cumulants of X_j and then use this result to identify the n th order cumulants of Z_j and hence of X_j , assuming these quantities exist. Finally, we investigate the identifiability of X_j for the case of a non-analytic characteristic function, i.e. when they have a symmetric α stable (S- α -S) distribution.

4.1. Identifiability of second-order cumulants

Given the covariance model \mathcal{Q}_M , it is fairly straightforward to get conditions for identifiability of second-order cumulants. Define

$$\mathcal{Q}_M(i, j) = \{(q_1, q_2) : ((i - 1)J + q_1, (j - 1)J + q_2) \in \mathcal{Q}_M\},$$

for $1 \leq i, j \leq K$. Elements of $\mathcal{Q}_M(i, j)$ indicate the admissible non-zero elements of $\Sigma_{X^{(i)}, X^{(j)}}$ under model \mathcal{Q}_M . Now for $r = (r_1, r_2)$ we have that

$$[\Sigma_{Y^{(i)}, Y^{(j)}}]_r = \sum_{q \in \mathcal{Q}_M(i, j)} A_{r_1, q_1} [\Sigma_{X^{(i)}, X^{(j)}}]_q A_{r_2, q_2}.$$

If $i \neq j$,

$$[\Sigma_{Y^{(i)}, Y^{(j)}}]_r = \sum_{\substack{q_1=q_2 \\ q \in \mathcal{Q}_M(i, j)}} A_{r_1, q_1} [\Sigma_{X^{(i)}, X^{(j)}}]_q A_{r_2, q_2} + \sum_{\substack{q_1 \neq q_2 \\ q \in \mathcal{Q}_M(i, j)}} A_{r_1, q_1} [\Sigma_{X^{(i)}, X^{(j)}}]_q A_{r_2, q_2}.$$

Let $g(r)$ be an ordering of $r \in \{1, \dots, L\}^2$ and $h(q)$ an ordering of $q \in \mathcal{Q}_M(i, j)$. Define matrices $\bar{B}(i, j, A)$, $i \neq j$ of order $L^2 \times |\mathcal{Q}_M(i, j)|$, as

$$[\bar{B}(i, j, A)]_{g(r), h(q)} = A_{r_1, q_1} A_{r_2, q_2}. \tag{13}$$

If $i = j$,

$$\begin{aligned} [\Sigma_{Y^{(i)}, Y^{(i)}}]_r &= \sum_{\substack{q_1=q_2 \\ q \in \mathcal{Q}_M(i, j)}} A_{r_1, q_1} [\Sigma_{X^{(i)}, X^{(i)}}]_q A_{r_2, q_2} \\ &+ \sum_{\substack{q_1 < q_2 \\ q \in \mathcal{Q}_M(i, j)}} (A_{r_1, q_1} A_{r_2, q_2} + A_{r_1, q_2} A_{r_2, q_1}) [\Sigma_{X^{(i)}, X^{(i)}}]_q. \end{aligned}$$

Now let $h'(q)$ be an ordering of $q \in \mathcal{Q}' \equiv \{(q_1, q_2) : q_1 \leq q_2, (q_1, q_2) \in \mathcal{Q}_M(i, j)\}$. Analogously define matrices $\bar{B}(i, i, A)$ of order $L^2 \times |\mathcal{Q}'|$, as

$$[\bar{B}(i, i, A)]_{g(r), h'(q)} = \begin{cases} A_{r_1, q_1} A_{r_2, q_2} & q_1 = q_2 \\ A_{r_1, q_1} A_{r_2, q_2} + A_{r_1, q_2} A_{r_2, q_1} & q_1 \neq q_2. \end{cases} \tag{14}$$

Note from equations (13) and (14) that each row of $\bar{B}(i, j, A)$ corresponds to a pair of rows in A . Further, each column corresponds to a unique non-zero element in Σ_X . Thus, for $i \neq j$, the (l, m) th element of $\bar{B}(i, j, A)$ is obtained as the product of elements in A corresponding to the pair of rows determined by l (through g) and the pair of columns determined by m (through h). The $\bar{B}(i, j, A)$ matrices play a similar role to the $B(A, A)$ matrix introduced in section 1, where flow volumes were assumed to be independent.

Proposition 1. *The covariance matrix of the flow volumes Σ_X is identifiable for model \mathcal{Q}_M , if the matrices $\bar{B}(i, j, A)$ are full rank for all $i, j \in \{1, \dots, K\}$. Conversely, if $\bar{B}(i, j, A)$ is not a full rank matrix for some i, j , then there exist multivariate normally distributed random vectors, X_1 and X_2 , with covariances $\Sigma_1, \Sigma_2 \in \mathcal{S}(\mathcal{Q}_M)$, respectively, for which $\Sigma_1 \neq \Sigma_2$, but $\bar{A}X_1 \equiv Y_1 \stackrel{d}{=} Y_2 \equiv \bar{A}X_2$.*

Proof. Let $\text{vec}_1, \text{vec}_2$ and vec_3 be matrix vectorizing operators such that $[\text{vec}_1(\Sigma)]_{g(r)} = \Sigma_r$, $[\text{vec}_2(\Sigma)]_{h(q)} = \Sigma_q$ and $[\text{vec}_3(\Sigma)]_{h'(q)} = \Sigma_q$, respectively. In other words, $\text{vec}_1, \text{vec}_2$ and vec_3 vectorize an entire $L \times L$ matrix, elements of a $J \times J$ matrix corresponding to $\mathcal{Q}_M(i, j)$ and elements of a $J \times J$ matrix corresponding to \mathcal{Q}' , respectively. These vectorization operations are consistent with orderings $g(r), h(q)$ and $h'(q)$. Thus, for $i \neq j$ we get

$$\text{vec}_1(\Sigma_{Y^{(i)}, Y^{(i)}}) = \bar{B}(i, j, A) \text{vec}_2(\Sigma_{X^{(i)}, X^{(i)}})$$

and

$$\text{vec}_1(\Sigma_{Y^{(i)}, Y^{(i)}}) = \bar{B}(i, i, A) \text{vec}_3(\Sigma_{X^{(i)}, X^{(i)}}).$$

The first assertion now follows directly.

For the converse, recall that $\mathcal{S}(\mathcal{Q}_M)$ has a non-empty interior. Further, for X being a multivariate normal random vector with mean 0, the distribution of $Y = \bar{A}X$ is completely determined by its covariance. Therefore, by choosing $\Sigma_1, \Sigma_2 \in \mathcal{S}(\mathcal{Q}_M)$ such that (the appropriate vectorization of) $\Sigma_1 - \Sigma_2$ is in the null space of $\bar{B}(i, j, A)$, we can get the required result. Specifically, if v is a vector in the null space of $\bar{B}(i, j, A)$, and $s(q) \equiv (q_1 - (i - 1)J, q_2 - (j - 1)J)$, then choose Σ_1 and Σ_2 such that

$$[\Sigma_1 - \Sigma_2]_q = \begin{cases} v_{l(q)} & s(q) \in \mathcal{Q}_M(i, j) \\ 0 & \text{otherwise,} \end{cases}$$

where $l(q) \equiv h(s(q))$ if $i \neq j$ or $l(q) \equiv h'(s(q))$ if $i = j$. □

Recall that under the latent variable model $\Sigma_X = \Sigma_X(C)$ is a function of the latent variable coefficient matrix. We then obtain the following result.

Corollary. *If $\Sigma_X(C) : \mathcal{C} \rightarrow \mathcal{S}(\mathcal{Q}_M)$ is a one-to-one map, then C is identifiable if the matrices $\bar{B}(i, j, A)$ are full rank for all $i, j \in \{1, \dots, K\}$.*

4.2. Identifiability of higher order cumulants

For the latent variable model, the above corollary implies that once the coefficient matrix C has been identified, we can proceed to identify higher order cumulants, provided they exist for latent variables Z .

Recall that $Y^{(k)} = AX^{(k)}$ for $1 \leq k \leq K$. Further

$$X^{(k)} = \sum_{l=1}^k \tilde{C}_{kl} \tilde{Z}_l$$

and therefore

$$Y^{(k)} = A \sum_{l=1}^k \tilde{C}_{kl} \tilde{Z}_l.$$

Notation: The notation $C_{i,\cdot}$ and $C_{\cdot,j}$ will be used to refer to the i th row and j th column respectively of a matrix C .

Let e_i be a row vector of length L , with 1 at the i th position and 0 otherwise. Write

$$e_{r_1} Y^{(k)} = (A_{r_1,\cdot}) \sum_{l=1}^k \tilde{C}_{kl} \tilde{Z}_l = \sum_{l=1}^k \sum_{j=1}^J ((A_{r_1,\cdot})[\tilde{C}_{kl}]_{\cdot,j}) [\tilde{Z}_l]_j.$$

Denote the logarithm of the characteristic function of the latent variables by

$$\phi_{Z,j,l}(t) \equiv \log E[e^{t[Z_j]_l}].$$

Therefore,

$$\begin{aligned} \phi_{r_1,r_2,k}(t, s) &\equiv \log E[e^{tY_{r_1}^{(k)} + isY_{r_2}^{(k)}}] \\ &= \sum_{l=1}^k \sum_{j=1}^J \phi_{Z,j,l}(t(A_{r_1,\cdot})[\tilde{C}_{kl}]_{\cdot,j} + s(A_{r_2,\cdot})[\tilde{C}_{kl}]_{\cdot,j}). \end{aligned}$$

Let $\phi_{Z,j,l}^{(n)}(0)$ be the n th derivative of $\phi_{Z,j,l}(t)$ at $t = 0$. Differentiating at $(0, 0)$ gives

$$\begin{aligned} \left. \frac{\partial^n \phi_{r_1,r_2,k}(t, s)}{\partial t^m \partial s^{n-m}} \right|_{(t,s)=(0,0)} &= \sum_{l=1}^k \sum_{j=1}^J ((A_{r_1,\cdot})[\tilde{C}_{kl}]_{\cdot,j})^m ((A_{r_2,\cdot})[\tilde{C}_{kl}]_{\cdot,j})^{n-m} \phi_{Z,j,l}^{(n)}(0) \\ &= \sum_{l=1}^{k-1} \sum_{j=1}^J ((A_{r_1,\cdot})[\tilde{C}_{kl}]_{\cdot,j})^m ((A_{r_2,\cdot})[\tilde{C}_{kl}]_{\cdot,j})^{n-m} \phi_{Z,j,l}^{(n)}(0) \\ &\quad + \sum_{j=1}^J ((A_{r_1,\cdot})[\tilde{C}_{kk}]_{\cdot,j})^m ((A_{r_2,\cdot})[\tilde{C}_{kk}]_{\cdot,j})^{n-m} \phi_{Z,j,k}^{(n)}(0). \end{aligned} \tag{15}$$

Define the matrices $\mathcal{B}(n, m, \tilde{C})$ as follows:

$$[\mathcal{B}(n, m, \tilde{C})]_{g(r),j} = ((A_{r_1,\cdot})\tilde{C}_{\cdot,j})^m ((A_{r_2,\cdot})\tilde{C}_{\cdot,j})^{n-m} \tag{16}$$

and sets \mathcal{C}_k of matrices such that for

$$C = \begin{pmatrix} \tilde{C}_{11} & 0 & \cdots & 0 \\ \tilde{C}_{21} & \tilde{C}_{22} & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ \tilde{C}_{K1} & \tilde{C}_{K2} & \cdots & \tilde{C}_{KK} \end{pmatrix},$$

$\mathcal{C}_k = \{\tilde{C}_{kk} : C \in \mathcal{C}\}$, $1 \leq k \leq K$. We can then establish the following proposition.

Proposition 2. *The n th order cumulants ($n \geq 2$) of Z are identifiable if the matrices $\bar{B}(i, j, A)$ are full rank for all $i, j \in \{1, \dots, K\}$ and the matrices $B(n, m, \tilde{C})$ are full rank for all $\tilde{C} \in \mathcal{C}_k$ for $1 \leq k \leq K$ and some $1 \leq m < n$.*

Proof. From the corollary of proposition 1, matrix C is identifiable. Since (15) is true for each pair of indices $r = (r_1, r_2) \in \{1, \dots, L\}^2$, we arrange all such equations in vector notation to get

$$v = z + B(n, m, \tilde{C}_{kk})\phi, \tag{17}$$

where

$$[v]_{g(r)} = \left. \frac{\partial^n \phi_{r_1, r_2, k}(t, s)}{\partial t^m \partial s^{n-m}} \right|_{(t,s)=(0,0)},$$

$$[z]_{g(r)} = \sum_{l=1}^{k-1} \sum_{j=1}^J ((A_{r_1, \cdot})[\tilde{C}_{kl} \cdot, j])^m ((A_{r_2, \cdot})[\tilde{C}_{kl} \cdot, j])^{n-m} \phi_{Z, j, l}^{(n)}(0)$$

and $\phi_j = \phi_{Z, j, k}^{(n)}(0)$. We will proceed inductively. Assume that $\phi_{Z, j, l}^{(n)}(0)$ have been identified for $l = 1, \dots, k - 1$. This means that z in equation (17) is known. Since v is a parameter vector of the distribution of Y , it is also known. The result now follows easily. \square

4.3. Identifiability of flow volumes under stable distributions

The above strategy for identifying the parameters of the latent variable model fails in the presence of heavy tailed distributions for the flow volumes, since second-order cumulants do not exist. However, under the assumption that the latent variables Z are symmetric α -stable distributed (S- α -S), sufficient conditions for identifiability can be derived. Specifically, if the Z_i are i.i.d. standard normal random variables, the corollary of proposition 1 gives a sufficient condition for identifiability of the distribution of X , which turns out to be sufficient for identifiability under a S- α -S distribution.

The characteristic function of a random variable W distributed according to a S- α -S law and scale parameter σ is given by

$$\psi_W(\theta) \equiv E[e^{i\theta W}] = e^{-\sigma^\alpha |\theta|^\alpha}.$$

Let $Z_i, i = 1, \dots, n$, be i.i.d. S- α -S random variables with unit (say) scale coefficients. Then, it is easy to see that any linear combination of Z_i is also S- α -S. Let $Y = \sum_i a_i Z_i$ to get

$$\begin{aligned} \psi_Y(\theta) \equiv E[e^{i\theta Y}] &= E[e^{i\theta \sum_i a_i Z_i}] = \prod_i E[e^{i a_i \theta Z_i}] = \prod_i \psi_{Z_i}(a_i \theta) \\ &= \prod_i e^{-|a_i \theta|^\alpha} = e^{-(\sum_i |a_i|^\alpha) |\theta|^\alpha}. \end{aligned}$$

Define measure μ as the counting measure on $E = \{1, \dots, n\}$ and view a_1, \dots, a_n as a function $a : E \rightarrow \mathbb{R}, a(i) = a_i$. We then get that the scale coefficient of Y is given by

$(\int_E |a|^\alpha d\mu)^{1/\alpha}$. This is a useful way of viewing linear combinations of i.i.d. S- α -S random variables. Now for a S- α -S random vector given by $Y = MZ$ where M is an $m \times n$ matrix, we have a representation $\bar{M} \subset L^\alpha(E, \mathcal{E}, \mu)$, where $\mathcal{E} = 2^E$ and $L^\alpha(E, \mathcal{E}, \mu)$ is the space of functions $a : E \rightarrow \mathbb{R}$ with ‘norm’ $(\int_E |a|^\alpha d\mu)^{1/\alpha}$, which is a real norm for $\alpha \geq 1$, and

$$\bar{M} = \{m \in L^\alpha(E, \mathcal{E}, \mu) : m(j) = M_{i,j} \text{ for all } j \text{ and some } i\}.$$

A linear isometry from $L^\alpha(E_1, \mathcal{E}_1, \mu_1)$ to $L^\alpha(E_2, \mathcal{E}_2, \mu_2)$ is a linear map $T : L^\alpha(E_1, \mathcal{E}_1, \mu_1) \rightarrow L^\alpha(E_2, \mathcal{E}_2, \mu_2)$ such that

$$\int_{E_1} |a|^\alpha d\mu_1 = \int_{E_2} |T(a)|^\alpha d\mu_2,$$

for all $a \in L^\alpha(E_1, \mathcal{E}_1, \mu_1)$.

As an illustration consider three representations of the same bivariate S- α -S random vector Y , with Z_i , i.i.d. S- α -S with $\alpha = 1$ and $Z = (Z_1, Z_2, Z_3)'$.

$$Y \equiv \begin{pmatrix} Y_1 \\ Y_2 \end{pmatrix} \stackrel{d}{=} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 2 & 3 \end{pmatrix} Z \stackrel{d}{=} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 3 & 2 \end{pmatrix} Z \stackrel{d}{=} \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 4 \end{pmatrix} Z.$$

Now Y_1 and Y_2 are independent and thus the only remaining parameters in the distribution of (Y_1, Y_2) are the two scale coefficients which are given by the l^1 norm of the corresponding rows. Clearly, the first and second representations give the same distribution since they only differ through a permutation of columns. A key question is whether that is the only operation under which we get different representations. In this example, there is another representation which cannot be obtained by a permutation of the columns. It turns out that this is related to the third column being proportional to the second column, in each of the three representations. It has been shown that when two representations are minimal in a certain sense, they are related through a linear isometry, which for $0 < \alpha < 2$ corresponds to (generalized) permutations.

Definition 3. A matrix M is called rigid if no column in M is proportional to another column in M . In particular this implies that, no column in M is identically 0.

We will show next that a rigid matrix corresponds to a minimal representation (lemma 1). Subsequently, we establish that linear isometries for $0 < \alpha < 2$ correspond to generalized permutations (lemma 2) and that there is no linear isometry of $l^\alpha(\mathbb{R}^{n_1})$ onto $l^\alpha(\mathbb{R}^{n_2})$ for $n_1 > n_2$ (lemma 3). Finally, we show that for a set \mathcal{C} of matrices sufficiently large, non-identifiability in the stable case leads to non-identifiability in the normal case (proposition 3). The set \mathcal{C} —representing the possible latent variable models—is required to be sufficiently large in the following sense.

Definition 4. A set of $n \times n$ real matrices \mathcal{C} will be said to be closed under column scaling if $C = (c_1, \dots, c_n) \in \mathcal{C}$ implies $C \text{Diag}(\alpha) = (\alpha_1 c_1, \dots, \alpha_n c_n) \in \mathcal{C}$ for all $\alpha \in \mathbb{R}^n$.

Lemma 1. If $M_1 Z \stackrel{d}{=} M_2 Z$ for Z_i i.i.d. S- α -S, $0 < \alpha < 2$, $1 \leq i \leq n$, and M_1, M_2 are rigid with n columns then $[M_1]_{k,\cdot} = T([M_2]_{k,\cdot})$, where T is a linear isometry of $l^\alpha(\mathbb{R}^n)$ onto $l^\alpha(\mathbb{R}^n)$.

Proof. This result basically follows from theorem 5.2 in [7], which states that minimal representations are related through linear isometries. All that remains to be verified, in our setting, is that a rigid matrix M as defined above is a minimal representation as defined in [7]. Minimality in this setting is defined as follows.

Let $E = \{1, \dots, n\}$, $\mathcal{E} = 2^E$ and μ be the counting measure on E . Let $F \subset L^\alpha(E, \mathcal{E}, \mu)$ be the set of functions given by rows of M . M is a minimal representation if

$$\rho(F) \equiv \sigma\{f/g : f, g \in F\} = \mathcal{E}.$$

Since $\mathcal{E} = 2^E$, the above simply means that the ratio of functions in F should ‘shatter’ each element of E individually, i.e. no two elements in E should have the same image under all ratio functions. Now since no two columns are proportional, for pair $i, j \in E$ there exist $f, g \in F$ such that $f(i)/g(i) \neq f(j)/g(j)$. Thus, we can define set $A_{i/j} \in \rho(F)$ such that $i \in A_{i/j}$ and $j \notin A_{i/j}$. Finally

$$\{i\} = \bigcap_{j \neq i} A_{i/j}.$$

Thus $\mathcal{E} = \rho(F)$. □

Lemma 2. *A linear isometry T of $l^\alpha(\mathbb{R}^n)$ onto $l^\alpha(\mathbb{R}^n)$ for $0 < \alpha < 2$ is given by a generalized permutation, i.e. for a row vector v , $T(v) = vPD$, where P is a permutation matrix and D is a diagonal matrix with all diagonal entries equal to either 1 or -1 .*

Proof. This result basically follows from theorem 3.1 in [10], which states that linear isometries are induced by regular set isomorphisms for $\alpha \neq 2$. All that remains to be verified is that for $E = \{1, \dots, n\}$, $\mathcal{E} = 2^E$ and μ the counting measure on E , the only regular set isomorphisms are permutations. A mapping T of \mathcal{E} into itself is called a regular set isomorphism if for all set $A, A_k \in \mathcal{E}$

- (1) $T(E - A) = T(E) - T(A)$.
- (2) $T(\bigcup_{k=1}^\infty A_k) = \bigcup_{k=1}^\infty T(A_k)$ for disjoint A_k .
- (3) $\mu(T(A)) = 0$ if and only if $\mu(A) = 0$.

We will show that for $i, j \in E, i \neq j, T(\{i\}) \cap T(\{j\}) \neq \phi$ leads to a contradiction. Assume $T(\{i\}) \cap T(\{j\}) = A$. Now $T(E - \{i\}) = T(E) - T(\{i\})$. But $j \in E - \{i\}$ and thus $A \cap T(E - \{i\}) \neq \phi$ but since $A \subseteq T(\{i\})$, $A \cap (T(E) - T(\{i\})) = \phi$. Thus, we obtain a contradiction that completes the proof. □

Lemma 3. *There is no linear isometry of $l^\alpha(\mathbb{R}^{n_1})$ onto $l^\alpha(\mathbb{R}^{n_2})$ for $n_1 > n_2, \alpha \neq 2$. In other words, for E and \mathcal{E} as before and μ_1 and μ_2 being counting measures that assign non-zero measure to n_1 and n_2 elements respectively of E , there is no linear isometry of $L^\alpha(E, \mathcal{E}, \mu_1)$ onto $L^\alpha(E, \mathcal{E}, \mu_2)$, for $\alpha \neq 2$.*

Proof. From corollary 2.1 in [10], for $\alpha \neq 2, f, g \in L^\alpha(E, \mathcal{E}, \mu)$ we have that

$$\|f + g\|_\alpha^\alpha + \|f - g\|_\alpha^\alpha = 2\|f\|_\alpha^\alpha + 2\|g\|_\alpha^\alpha,$$

if and only if $f(x)g(x) = 0$ a.e. (μ). Thus, if $T : L^\alpha(E, \mathcal{E}, \mu_1) \rightarrow L^\alpha(E, \mathcal{E}, \mu_2)$ is a linear isometry, $f, g \in L^\alpha(E, \mathcal{E}, \mu_1)$ and $f(x)g(x) = 0$ a.e. (μ_1).

$$\begin{aligned} \|T(f) + T(g)\|_\alpha^\alpha + \|T(f) - T(g)\|_\alpha^\alpha &= \|T(f + g)\|_\alpha^\alpha + \|T(f - g)\|_\alpha^\alpha \\ &= \|f + g\|_\alpha^\alpha + \|f - g\|_\alpha^\alpha \\ &= 2\|f\|_\alpha^\alpha + 2\|g\|_\alpha^\alpha \\ &= 2\|T(f)\|_\alpha^\alpha + 2\|T(g)\|_\alpha^\alpha. \end{aligned}$$

This implies that $T(f)(x)T(g)(x) = 0$ a.e. (μ_2). Hence, any linear isometry T , for $\alpha \neq 2$, maps functions with disjoint supports to functions with disjoint supports. Since μ_1 has $n_1 > n_2$ distinct support points with positive measure the result follows. □

Proposition 3. *Assume the set \mathcal{C} of matrices is closed under column scaling. If there are $C_1, C_2 \in \mathcal{C}$ such that $C_1 \neq C_2$ and $AC_1Z \stackrel{d}{=} AC_2Z$ for Z_i i.i.d. S - α - $S, 0 < \alpha < 2$, then there exist $C_3, C_4 \in \mathcal{C}, C_3 \neq C_4$ such that $AC_3\tilde{Z} \stackrel{d}{=} AC_4\tilde{Z}$ for \tilde{Z}_i i.i.d. standard normal.*

Proof. We examine the following four possible cases:

Case 1. If AC_1 and AC_2 are both rigid, then $AC_1 = AC_2PD$, with P being a permutation matrix and D a diagonal matrix with all diagonal entries 1 or -1 . But $Z \stackrel{d}{=} PDZ$ and therefore $C_3 = C_1$ and $C_4 = C_2$ gives the required result.

Case 2. On the other hand, if all columns of AC_1 and AC_2 are zero clearly $C_3 = C_1$ and $C_4 = C_2$ gives the desired result.

Case 3. If AC_1 and AC_2 are rigid after removing the zero columns, proceed as follows. Define counting measures μ_1 and μ_2 that assign no measure to the indices corresponding to the zero columns in AC_1 and AC_2 , respectively. Now argue as in the rigid case under measures μ_1 for rows of AC_1 and μ_2 for rows of AC_2 . Note that the number of elements with non-zero measure under μ_1 and μ_2 respectively, has to be the same from the previous lemma and theorem 5.2 in [7].

Case 4. Now, whenever AC_1 and AC_2 are not both rigid, without loss of generality assume that $AC_1 = (v_1, \dots, v_n)$ with $v_1 = \alpha v_2, v_2 \neq 0, \alpha \neq 0$. Choose $C_3 \in \mathcal{C}$ such that $AC_3 = (\tilde{v}_1, \dots, \tilde{v}_n)$ where $\tilde{v}_i = v_i$ for $i > 2$ and $\tilde{v}_1 = (\sqrt{1+\epsilon})v_1 = (\sqrt{1+\epsilon})\alpha v_2$ and $\tilde{v}_2 = (\sqrt{1-\epsilon\alpha^2})v_2$ for $0 < \epsilon < 1/\alpha^2$. Let e_i be the i th column of a $(n-1) \times (n-1)$ identity matrix. Thus

$$AC_1 = (v_2, \dots, v_n)(\alpha e_1, e_1, e_2, \dots, e_{n-1}) \equiv (v_2, \dots, v_n)E_1,$$

$$AC_3 = (v_2, \dots, v_n)((\sqrt{1+\epsilon})\alpha e_1, (\sqrt{1-\epsilon\alpha^2})e_1, e_2, \dots, e_{n-1}) \equiv (v_2, \dots, v_n)E_3.$$

Now, $E_1E_1' = \alpha^2 e_1 e_1' + I = E_3E_3'$. So, $AC_1(AC_1)' = AC_3(AC_3)'$. Hence, $C_4 \equiv C_1, C_3 \in \mathcal{C}$ and $AC_4\tilde{Z} \stackrel{d}{=} AC_3\tilde{Z}$ for \tilde{Z}_i i.i.d. standard normal. □

5. Identifiability in the independent connections model

In this section, we derive sufficient conditions on the network routing that guarantee full rankness of the matrices appearing in propositions 1, 2 and 3.

In order to gain more insight, the matrices in propositions 1 and 3 will be simplified in the following for the independent connections model (see (7)–(8)). Recall that we defined $B(A_1, A_2)$ as the matrix in which row $g(i)$ is the component-wise product of row i_1 of A_1 and row i_2 of A_2 , where $i \equiv (i_1, i_2) \in \{1, \dots, L\}^2$. Recall that if all the X_j are independent, identifiability of variances follows from full-rankness of $B(A, A)$. Recalling the interpretation of the (l, m) th element of $\bar{B}(i, j, A)$ as a product of the pair of elements in A corresponding to rows determined by l , and columns by the m th element of $\mathcal{Q}_M(i, j)$, it is straightforward to compute the required matrices. Now define

$$\bar{B} = (B(A_F, A_F) \quad B(A_F, A_R) \quad B(A_R, A_F) \quad B(A_R, A_R)) \tag{18}$$

and

$$\bar{B}_c = (B(A_F, A_F) \quad B(A_F, A_R) + B(A_R, A_F) \quad B(A_R, A_R)). \tag{19}$$

Here \bar{B} is a $L^2 \times 4J$ matrix and \bar{B}_c is a $L^2 \times 3J$ matrix. Now, proposition 1 implies the following. For $K = 1$, the second-order cumulants, Σ_X , of the independent connections model (7), are identifiable if and only if \bar{B}_c is full rank. For $K > 1$, the second-order cumulants of the independent connections model are identifiable if and only if \bar{B} is full rank.

Once second moments are identified, the corollary to proposition 1 would imply identifiability of C in model (8). Thus, it only remains to use proposition 2 to identify

higher moments. Note that for the independent connections model, the set \mathcal{C}_k consists of matrices of the form

$$\tilde{C} = \begin{pmatrix} \text{Diag}(c_1) & 0 \\ \text{Diag}(c_2) & \text{Diag}(c_3) \end{pmatrix}.$$

Thus, $\tilde{C}_{\cdot,j}$ have either 1 or 2 non-zero entries and therefore the product of a row from A and a column from \tilde{C} appearing in (16) has either 1 or 2 terms. It can then be seen that the required matrix takes the form

$$\mathcal{B}_{n,m}(c_1, c_2, c_3) = [H_{n,m}(c_1, c_2), J_{n,m}(c_3)],$$

where

$$[H_{n,m}(c_1, c_2)]_{g(p),j} = (c_{1j}[A_F]_{p_1,j} + c_{2j}[A_R]_{p_1,j})^m (c_{1j}[A_F]_{p_2,j} + c_{2j}[A_R]_{p_2,j})^{n-m}$$

and

$$[J_{n,m}(c_3)]_{g(p),j} = c_3^n ([A_R]_{p_1,j})^m ([A_R]_{p_2,j})^{n-m}.$$

Proposition 2 implies the following: if C in model (8) is identifiable and n order cumulants of Z exist, then they can be identified if $\mathcal{B}_{n,m}(c_1, c_2, c_3)$ is a full rank for some $1 \leq m < n$ and all $c_1, c_2, c_3 \in \mathbb{R}^{J/2}$.

5.1. Minimum weight routing

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 $j = 1, \dots, J$. Define the operator $\mathcal{R}(\cdot)$ on paths such that if path $P = (m_1, m_2, \dots, m_{k-1}, m_k)$ then

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

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 4. *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 .*




Figure 8. Lemma 1, case 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 $P_1 = (f_1, f_2, \dots, l_1, l_2)$ 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 , a minimum weight path P_1 from k_1 to k_2 and a path $P_2 = \mathcal{R}(P_1)$ such that the edge (f_1, f_2) lies in P_1 and the 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 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 two possible structures:

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




Figure 9. Lemma 1, case 2.

One can now establish the following:

Proposition 4. *Under balanced minimum weight routing on a symmetric graph the matrix \bar{B} is full rank.*

Proof. Let

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

where $v_{FF}, v_{FR}, v_{RF}, v_{RR} \in \mathbb{R}^{J/2}$. We need to show that $v_{FF} = v_{FR} = v_{RF} = v_{RR} = 0$. Now let \mathcal{F} be the ordered set of node pairs with the same ordering as in vectors X_F and X_R . Also 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 . Define operators \mathcal{P}_F and \mathcal{P}_R which map a set of indices to sets of paths as follows:

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

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 4 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 \bar{B} . Consider the row of \bar{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 4. The third claim follows from the third part of lemma 4. 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 4 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 \bar{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 4. The second claim follows from the third part of lemma 4. 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. □

Corollary. *The matrices \bar{B}_c and $\mathcal{B}_{n,m}(c_1, c_2, c_3)$ are full rank for all $1 \leq m < n$ and all $c_1, c_2, c_3 \in \mathbb{R}^{J/2}$ under balanced minimum weight routing on a symmetric graph.*

These results follow easily by comparing the above matrices to \bar{B} .

6. Special cases

We examine next certain special cases of interest for computer networks.

6.1. 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* [3] proved that second moments 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 subnetworks, 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 of second cumulants depends on full rankness of $B(A, A)$. The matrix $B(A, A)$ can easily be shown to be full rank for hierarchical networks. The proof (see [3]) 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. Note that neither minimum cost routing, nor deterministic routing is required for the argument. In fact, the matrix \bar{B} (and hence \bar{B}_c and $B_{n,m}(c_1, c_2, c_3)$) can be shown to be full rank, which implies identifiability of second (and higher) cumulants under the independent connections model.

Proposition 5. *The matrix \bar{B} is full rank for hierarchical networks.*

Proof. We will prove that given $i \in \{1, \dots, 2J\}$ there exists a row r in \bar{B} such that r is the i th row of a $2J \times 2J$ identity matrix. Index i corresponds to flow pair $i' = ((i - 1) \bmod J/2) + 1$. For $i \in \{1, \dots, J/2\}$ ($i \in \{1, \dots, J/2\}$) 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 $i \in \{J/2 + 1, \dots, J\}$ ($i \in \{J + 1, \dots, 3J/2\}$) 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 \bar{B} gives the required result. \square

In the following, we use a similar idea to prove identifiability of second moments for directed acyclic graphs.

6.2. 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 $B = B(A, A)$.

Proposition 6. *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 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:

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

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.

6.3. Two-dimensional measurements

Note that the identifiability of second cumulants is driven by the fact that the covariance between two links is the sum of variances of flows common to both of them. Now, assume deterministic routing, i.e. that A is binary. In principle, if we could get direct measurements of the type \mathcal{T}_{e_1, e_2} for all pairs of edges e_1 and e_2 that give the volume of flows common to e_1 and e_2 , then we should be able to solve for individual flow volumes. To see this note that $T = \{\mathcal{T}_{ij}\}$ satisfies

$$\text{vec}(T) = B(A, A)X.$$

Thus, we have the following result.

Proposition 7. *If $B(A, A)$ is full rank and A binary, the full distribution of X is identifiable from T . In particular, means of the flow volume distribution are also estimable.*

Remark. The technology in current computer networks does not allow collection of the necessary data \mathcal{T}_{e_1, e_2} for all pairs of edges e_1 and e_2 , since that would require tagging individual packets. At present this seems feasible only for edges incident on the same node. On the other hand, loop detectors on road networks makes collection of such data feasible for a number of edge pairs. One direction of future work would be to study the case of having such measurements available for a subset of edge pairs.

7. Discussion and concluding remarks

The problem of identifiability of distributions of flows on a graph from aggregate measurements collected on its edges has been considered and resolved for a fairly general dependence structure. We presented a modelling framework that can accommodate multi-modal measurements with spatio-temporal distributions and provided empirical justification for it. The first proposition and its corollary provide verifiable conditions for identifiability of covariance and the latent variable models, respectively. The second proposition provides sufficient conditions for identifiability of higher order cumulants under the latent variable model. Proposition 3 shows that identifiability under a normal latent variable model implies identifiability under the related S- α -S model under mild conditions. Proposition 4 and its corollary provide reasonable sufficient conditions on the network structure (topology and routing scheme) so that the second and higher cumulants of the independent connection model are identifiable. Propositions 5, 6 and 7 deal with identifiability of second-order cumulants for special networks and measurement schemes. We address next some open issues.

7.1. Dependence structure

A natural question to consider is if it is possible to establish our results under a more general dependence structure between the various measurements of the forward and reverse flows of a flow pair. Specifically, we assume flow pairs to be independent but do not assume a latent variable model. Consider a general characteristic function for a flow pair j

$$\psi_j(t, s) = E[\exp(it X_{Fj} + is X_{Rj})],$$

where X_{Fj} and X_{Rj} are the forward and reverse components of the flow pair j . Similarly let $\psi_{Y_1, Y_2}(t, s)$ be the joint characteristic function of Y_1 and Y_2 and $\phi_{Y_1, Y_2}(t, s) = \log \psi_{Y_1, Y_2}(t, s)$. Then it can be shown that

$$\frac{\partial^n \phi_{Y_1, Y_2}(t, s)}{\partial t^k \partial s^{n-k}} \Big|_{(t,s)=(0,0)} = \sum_j \sum_{r=0}^k \sum_{l=0}^{n-k} {}^k C_r {}^{n-k} C_l A_{F,1j}^r A_{R,1j}^{k-r} A_{F,2j}^l A_{R,2j}^{n-k-l} \phi_j^{l+r, n-r-2k-l}(0, 0),$$

where ${}^n C_m$ is the coefficient of x^m in $(1+x)^n$ and

$$\phi_j^{n,k}(t, s) = \frac{\partial^{n+k} \log \psi_j(t, s)}{\partial t^n \partial s^k}.$$

Attempting to arrange all such equations into a linear system gets cumbersome and leads to implicitly defined matrices. On the other hand, assuming minimum cost routing should lead to considerable simplification. One direction of future work would be to investigate in greater detail such a dependence structure.

7.2. Network routing and a counter-example

In section 5, it was shown that minimum cost routing and symmetric weights was sufficient to ensure the full rankness of the B and hence (\bar{B}) matrices. The following example shows that absence of these conditions renders the result invalid.

Consider the network shown in figure 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 $B = B(A, A)$ 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, the full rankness of B (and hence \bar{B}) is not guaranteed.

7.3. Weaker conditions for identifiability of the independent connections model

The above counter-example shows that the only possibility of relaxing the conditions for proposition 4 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 \bar{B} (or $B(A, A)$ 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 nodes 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 the 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 4 to the case of non-deterministic routing remains an open problem.

Acknowledgments

The authors would like to thank two anonymous referees for useful comments and suggestions that improved the presentation of the material. They would also like to thank Professor Stilian Stoev for many helpful discussions and suggestions on the material of section 4.3. This research was supported in part by NSF grant DMS-0505535.

References

- [1] Barford P, Kline J, Plonka D and Ron A 2002 A signal analysis of network traffic anomalies *IMW '02: Proc. 2nd ACM SIGCOMM Workshop on Internet Measurement* (New York: ACM Press) pp 71–82
- [2] Boyd S and Vandenberghe L 2004 *Convex Optimization* (Cambridge: Cambridge University Press)
- [3] Cao J, Davis D, Wiel S and Yu B 2000 Time-varying network tomography: router link data *J. Am. Stat. Assoc.* **95** 1063–75
- [4] Chen A and Cao J 2006 Method of one-dimensional projections for network tomography *Statistical Inverse Problems (IMS Lecture Note Series)* ed R Liu, W Strawderman and C H Zhang
- [5] Erramilli V, Crovella M and Taft N 2006 An independent-connection model for traffic matrices *IMC '06: Proc. 6th ACM SIGCOMM on Internet Measurement* (New York: ACM Press) pp 251–6
- [6] Hamada T, Chujo K, Chujo T and Yang X 2004 Peer-to-peer traffic in metro networks: analysis, modeling, and policies *Network Operations and Management Symposium* pp 425–38
- [7] Hardin CD 1982 On the spectral representation of symmetric stable processes *J. Multivariate Anal.* **12** 385–401
- [8] Kwon J and Varaiya P 2005 Real-time estimation of origin-destination matrices with partial trajectories from electronic toll collection tag data *Transp. Res. Record* **1923** 119–26
- [9] Lakhina A, Crovella M and Diot C 2004 Diagnosing network-wide traffic anomalies *SIGCOMM '04: Proc. 2004 Conf. on Applications, Technologies, Architectures, and Protocols for Computer Communications* (New York: ACM Press) pp 219–30
- [10] Lamperti J 1958 On the isometries of certain function spaces *Pacific J. Math.* **8** 459–66
- [11] Lawrence E, Michailidis G, Nair V N and Xi B 2006 Network tomography: a review and recent developments *Frontiers in Statistics* ed Fan and H Koul (London: Imperial College Press) pp 345–64
- [12] Liang G, Taft N and Yu B 2006 A fast lightweight approach to origin-destination IP traffic estimation using partial measurements *IEEE Trans. Inf. Theory* **52** 2634–48
- [13] Liu Y, Towsley D, Ye T and Bolot J C 2005 An information-theoretic approach to network monitoring and measurement *IMC '05, Internet Measurement Conf. (New Orleans, LA, USA)*
- [14] Medina A, Taft N, Salamatian K, Bhattacharyya S and Diot C 2002 Traffic matrix estimation: existing techniques and new directions *Proc. 2002 Conf. on Applications, Technologies, Architectures, and Protocols for Computer Communications* (New York: ACM Press)
- [15] Peterson L L and Davie B C 2003 *Computer Networks: A Systems Approach* (San Francisco: Morgan Kaufmann)
- [16] Soule A, Lakhina A, Taft N, Papagiannaki K, Salamatian K, Nucci A, Crovella M and Diot C 2005 Traffic matrices: balancing measurements, inference and modeling *Proc. Joint Int. Conf. on Measurement and Modeling of Computer Systems* (New York: ACM Press)
- [17] Soule A, Nucci A, Cruz R, Leonardi E and Taft N 2004 How to identify and estimate the largest traffic matrix elements in a dynamic environment *SIGMETRICS '04/Performance '04: Proc. Joint Int. Conf. on Measurement and Modeling of Computer Systems* (New York: ACM Press) pp 73–84
- [18] Uhlig S, Quoitin B, Balon S and Lepropre J 2006 Providing public intradomain traffic matrices to the research community *ACM SIGCOMM Comput. Commun. Rev.* **36**
- [19] Vardi Y 1996 Network tomography: Estimating source-destination traffic intensities from link data *J. Am. Stat. Assoc.* **91** 365–77
- [20] Zhang Y, Roughan M, Duffield N G and Greenberg A G 2003 Fast accurate computation of large-scale IP traffic matrices from link loads *Proc. Int. Conf. on Measurements and Modeling of Computer Systems, SIGMETRICS 2003* pp 206–17