First Order Characterization of Internet Traffic Matrices

Matthew Roughan
School of Mathematical Sciences, University of Adelaide
Adelaide, 5005, Australia
matthew.roughan@adelaide.edu.au

Introduction

A traffic matrix (giving traffic volumes from ingress to egress in a network) is the basic input to many network engineering tasks. However, an Internet traffic matrix is hard to measure directly, and so much work (for some examples see [1, 2, 3, 4, 5, 6]) has been conducted in recent years on how one can estimate a traffic matrix from indirect measurements of link traffic. Such measurements may be easily collected using SNMP (the Simple Network Management Protocol), but only reveal link traffic volumes, not the source or destination of the traffic, and one must solve a highly underconstrained problem in order to obtain the desired traffic matrix. Solving such problems requires side-information, which may be brought in if we know *a priori* a reasonable model for the traffic in question. A variety of such models have been proposed, and we examine critically one such, the gravity model. Importantly, such a simple model is unlikely to be perfectly accurate, but can still be useful in determining a viable traffic matrix. The simplicity of this model also makes it useful in simulations of networks where one wishes to generate many instances of random networks, with random traffic matrices.

Background and Related Work

An IP network can be abstractly thought of as a graph, whose nodes are routers, and whose edges are links between these. A Traffic Matrix (TM) describes the volumes of traffic traversing a network from the point at which it enters the network, to the exit point. Such a matrix is useful in capacity planning, traffic engineering, network reliability analysis, and many other network engineering tasks. It is possible to measure such a matrix using measurement technologies such as flow level traffic collection [7], but typically these are hard to implement across a large network [5]. On the other hand SNMP data is easy to collect, and almost ubiquitous. However, SNMP data only provides link load measurements, not TM measurements [5].The link measurements $\mathbf{y}$ are related to the TM, which is written as a column vector $\mathbf{x}$, by the relationship $\mathbf{y} = A\mathbf{x}$ where A is called the routing matrix. The routing matrix is simply defined so that A_{ij} equals the fraction of traffic from route j that uses link i in the network, and is zero for most elements (and thus forms a sparse matrix). The resulting problem of inferring the TM from link measurements is a classic underconstrained, linear-inverse problem.

There is extensive experience with ill-posed linear inverse problems from fields as diverse as seismology, astronomy, and medical imaging [8, 9, 10, 11, 12], all leading to the conclusion that some sort of side information must be brought in. Examples of such include, Vardi [1] and Tebaldi and West [2] who assume a Poisson traffic model; Cao et al. [3] assume a Gaussian traffic model; Zhang et al. [5] assume an underlying gravity model; and Medina et al. [4] assume a logit-choice model. Finally, a more general approach to the problem of linear inverse problems is to use *regularization*, as suggested in [6]. Zhang *et al* showed that an intuitive regularization approach, based on Minimization of Mutual Information (MMI), resulted in fast, accurate and robust estimate of the traffic matrix, which interestingly can be closely approximated by using a gravity model as a starting point. There are also alternative approaches that rely on gathering additional information, for instance by changing network routing [13], but this class of methods is only practical in special cases where control over network routing is possible.

Each method of estimation is sensitive to the accuracy of this prior: for instance, [4] showed that the methods in [1, 2, 3] were sensitive to their prior assumptions (which took the form of a relationship between the mean and variance of the traffic matrix elements when these are considered as a time series). Some effort has gone into validating the time series type models of traffic matrices, e,g, see [4, 14, 15], but these approaches assume independence of traffic matrix elements. The gravity model brings in spatial correlations between traffic matrix elements. The assumption was tested in [5, 6], on a large set of traffic data from a large Internet Service Provider (ISP) in North America (AT&T), where it was not found to be accurate enough for TM inference, but it was accurate enough to use as a starting point. All of these tests were aimed at assessing the model as a starting point estimation algorithms, not characterization of the traffic matrices for other applications such as simulation of traffic matrices, which we shall make the focus of this paper.

Further effort on modeling the relationships between TM elements has been performed in [16], and successfully exploited for anomaly detection in [17, 18]. These papers focused on a Principle Components Analysis (PCA) of the traffic matrices, as times series. The PCA exploits the correlations between traffic matrix elements to separate the periodic components of the traffic (as seen in [14]), from random fluctuations, and anomalous events. It is not obvious how the structures described within [16] would lead to a simple model for use in simulations of network traffic matrices. On the other hand, the gravity model is so simple that it has already been used in a number of simulation contexts as a model for network traffic, e.g. see [19].

Gravity models

Gravity models, taking their name from Newton's law of gravitation, are commonly used by social scientists to model the movement of people, goods or information between geographic areas [20, 21]. In Newton's law of gravitation the force is proportional to the product of the masses of the two objects divided by the distance squared. Similarly, in gravity models for interactions between cities, the relative strength of the interaction might be modeled as proportional to the product of the cities' populations. A general formulation of a gravity model is given by $X_{ij} = \frac{R_i \cdot A_j}{f_{ij}}$, where X_{ij} is the matrix element representing the force from i to j; R_i represents the *repulsive* factors that are associated with "leaving" from i; A_j represents the *attractive* factors that are associated with "going" to j; and f_{ij} is a friction factor from i to j.

In network applications, gravity models have been used to model mobility in wireless networks [22], and the volume of telephone calls in a network [23]. In the context of Internet TMs, we can naturally interpret X_{ij} as the traffic volume that enters the network at location i and exits at location j, the repulsion factor R_i as the traffic volume entering the network at location i, and the attractivity factor A_j as the traffic volume exiting at location j. The friction matrix (f_{ij}) encodes the locality information specific to different source-destination pairs, however, as locality is not as large a factor in Internet traffic as in the transport of physical goods, we shall assume a common constant for the friction factors. The resulting gravity model simply states that the traffic exchanged between locations is proportional to the volumes entering and exiting at those locations.

In this very simple gravity model, we aim to estimate the amount of traffic between nodes. Denote the nodes by n_1, n_2,.... We estimate the volume of traffic $T(n_i, n_j)$ that enters the network at node n_i and exits at node n_j. Let $T^{\text{in}}(n_i)$ and $T^{\text{out}}(n_j)$ denote the total traffic that enters the network via node n_i, and exits the network via node n_j, respectively. The gravity model can then be computed by either of

$$(1) \qquad T(n_i, n_j) = T \frac{T^{\text{in}}(n_i)}{\sum_k T^{\text{in}}(n_k)} \frac{T^{\text{out}}(n_j)}{\sum_k T^{\text{out}}(n_k)} = T p^{\text{in}}(n_i) p^{\text{out}}(n_j)$$

where T is the total traffic across the network, and $p^{\text{in}}(n_i)$ and $p^{\text{out}}(n_j)$ denote the probabilities of traffic entering and exiting the network at nodes i and j respectively. Under the (reasonable) assumption that the network is neither a source nor sink of traffic in itself, so all traffic crosses the network, then $T = \sum_k T^{\text{in}}(n_k) = \sum_k T^{\text{out}}(n_k)$ and we can also write

$$(2) \qquad p(n_i, n_j) = p^{\text{in}}(n_i) p^{\text{out}}(n_j)$$

where $p(n_i, n_j)$ is the probability that a packet (or byte) enters the network at node n_i and departs at node n_j. Hence we can see that the gravity model corresponds to an assumption of independence between source and destination of the traffic. More importantly, using the above, the gravity model can be written as a matrix formed from the product of two vectors, e.g.

$$P = \mathbf{p}_{\text{in}} \mathbf{p}_{\text{out}}^T \tag{3}$$

so by characterizing these two vectors, we obtain a reasonable characterization of the matrix.

In this paper, we use data derived from the Abilene research network `http://abilene.internet2.edu/` to examine the gravity model. As this is a research network, they have made their traffic, and network topology data public, including traffic matrix measurements. We have used this data to estimate link loads in the network, and from these test various approaches to estimation and characterization of traffic matrices. The particular result of interest here is how well the gravity model estimates the traffic matrix data. The answer is, not well enough for many applications: the average accuracy of the simple gravity model above is $\pm 39\%$. As in previous tests of estimation techniques [5, 6] these results can be considerably improved by using a better initial gravity model (which incorporates natural routing asymmetries), or by regularization techniques. However, in this paper, we shall allow this to be a somewhat inaccurate starting point for 1st order modeling of traffic matrices, and consider characterization of this gravity model as our goal.

In fact such a characterization turns out to be quite easy. We start by approximating the vectors $\mathbf{p}_{\text{in}}$ and $\mathbf{p}_{\text{out}}$ with an exponentially distributed random variable. Figure 1 (a) shows one example of such a fit. Note that, although the fit is not perfect, there are only 12 samples used in drawing this distribution, with obvious limitations on the accuracy, and also the accuracy of the characterization is still better than the accuracy of the gravity model overall. Hence, we consider this to be an excellent starting point for generating a gravity model. The traffic matrix is then generated as the matrix product described above. While, in reality, it may make sense for there to be correlations between the source and destination traffic volumes (the vectors $\mathbf{p}_{\text{in}}$ and $\mathbf{p}_{\text{out}}$), we found that the approximation works extremely well in generating a traffic matrix with similar statistics to the observed traffic matrices, as shown in Figure 1 (b).

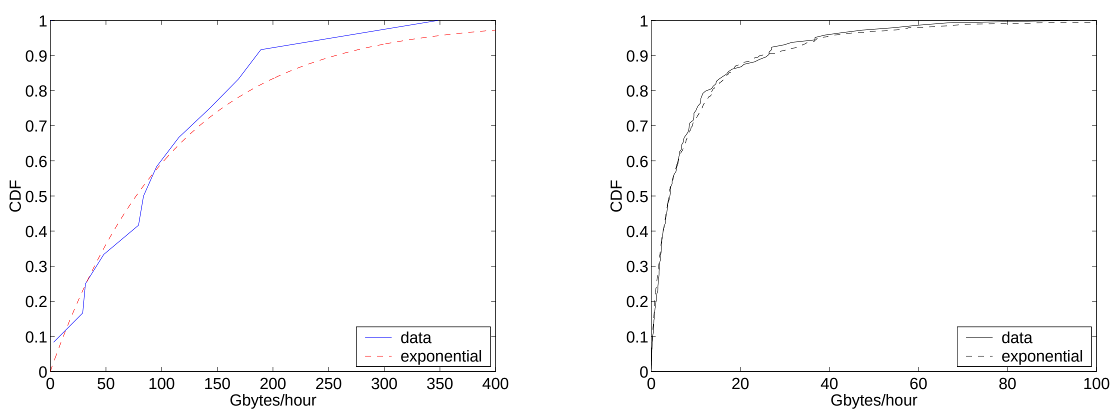

(a) Edge node traffic (b) Traffic matrix

Figure 1: ***Plots showing the Cumulative Distribution Functions (CDFs) of the traffic data at the edge of the network*** p_{in}***, and for the traffic matrix itself. The solid lines show the data, while the dashed lines are exponential fits to the data.***

Conclusion

In this paper we have presented a very simple first order model for the Abilene traffic matrix data. The model is a gravity model, and while such a model provides poor accuracy for estimation of traffic matrices, it appears to provide quite a reasonable approach for generating traffic matrices with similar statistics to the observed traffic matrices. Hence, such a model could be the basis for techniques for simulation of traffic matrices in many other applications. Future work will extend the validation of this model, and derive more refined models.

Acknowledgment

The data used in this paper generated from Abilene data by Yin Zhang of the University of Texas, to whom the author is grateful.

References

[1] Y. Vardi, "Network tomography: estimating source-destination traffic intensities from link data," *J. Am. Statist. Assoc.*, vol. 91, pp. 365–377, 1996.

[2] C. Tebaldi and M. West, "Bayesian inference on network traffic using link count data," *J. Amer. Statist. Assoc*, vol. 93, no. 442, pp. 557–576, 1998.

[3] J. Cao, D. Davis, S. V. Wiel, and B. Yu, "Time-varying network tomography," *J. Amer. Statist. Assoc*, vol. 95, no. 452, pp. 1063–1075, 2000.

[4] A. Medina, N. Taft, K. Salamatian, S. Bhattacharyya, and C. Diot, "Traffic matrix estimation: Existing techniques and new directions," in *ACM SIGCOMM*, Pittsburg, USA, August 2002.

[5] Y. Zhang, M. Roughan, N. Duffield, and A. Greenberg, "Fast accurate computation of large-scale IP traffic matrices from link loads," in *ACM SIGMETRICS*, San Diego, California, June 2003, pp. 206–217.

[6] Y. Zhang, M. Roughan, C. Lund, and D. Donoho, "An information-theoretic approach to traffic matrix estimation," in *ACM SIGCOMM*, Karlsruhe, Germany, August 2003, pp. 301–312.

[7] A. Feldmann, A. Greenberg, C. Lund, N. Reingold, J. Rexford, and F. True, "Deriving traffic demands for operational IP networks: Methodology and experience," *IEEE/ACM Transactions on Networking*, pp. 265–279, June 2001.

[8] M. Bertero, T. Poggio, and V. Torre., "Ill-posed problems in early vision." *Proc. of the IEEE*, vol. 76, no. 869–889, 1988.

[9] I. J. Craig and J. C. Brown, *Inverse Problems in Astronomy: A Guide to Inversion Strategies for Remotely Sensed Data.* Adam Hilger, Boston, 1986.

[10] A. Neumaier, "Solving ill-conditioned and singular linear systems: A tutorial on regularization." *SIAM Review*, vol. 40, no. 3, 1998.

[11] R. Parker, *Geophysical Inverse Theory.* Princeton University Press, Princeton, NJ, 1994.

[12] G. Wahba, *Statistical Decision Theory and Related Topics III.* Academic Press, 1982, vol. 2, ch. Constrained regularization for ill posed linear operator equations, with applications in meteorology and medicine., pp. 383–418.

[13] A. Nucci, R. Cruz, N. Taft, and C. Diot, "Design of IGP link weights for estimation of traffic matrices," in *IEEE Infocom*, Hong Kong, March 2004.

[14] M. Roughan, A. Greenberg, C. Kalmanek, M. Rumsewicz, J. Yates, and Y. Zhang, "Experience in measuring Internet backbone traffic variability: Models, metrics, measurements and meaning," in *Proceedings of the International Teletraffic Congress (ITC-18)*, Berlin, Germany, 2003, pp. 379–388.

[15] M. Roughan and J. Gottlieb, "Large-scale measurement and modeling of backbone Internet traffic," in *SPIE ITCom*, Boston, 2002.

[16] A. Lakhina, K. Papagiannaki, C. D. Mark Crovella, E. D. Kolaczyk, and N. Taft, "Structural analysis of network traffic flows," in *ACM SIGMETRICS / Performance*, 2004.

[17] A. Lakhina, M. Crovella, and C. Diot, "Diagnosing network-wide traffic anomalies," in *ACM SIGCOMM*, 2004.

[18] ——, "Characterization of network-wide anomalies in traffic flows," in *ACM/SIGCOMM Internet Measurement Conference*, 2004.

[19] B. Fortz, J. Rexford, and M. Thorup, "Traffic engineering with traditional IP routing protocols," *IEEE Communications Magazine*, vol. 40, no. 10, pp. 118–124, October 2002.

[20] J. Tinbergen, "Shaping the world economy: Suggestions for an international economic policy," The Twentieth Century Fund, 1962.

[21] P. Pyhnen, "A tentative model for the volume of trade between countries," *Weltwirtschaftliches Archive*, vol. 90, pp. 93–100, 1963.

[22] D. Lam, D. Cox, and J. Widom, "Teletraffic modeling for personal communications services," *IEEE Communications Magazine: Special Issues on Teletraffic Modeling Engineering and Management in Wireless and Broadband Networks*, vol. 35, no. 2, pp. 79–87, February 1997.

[23] J. Kowalski and B. Warfield, "Modeling traffic demand between nodes in a telecommunications network," in *ATNAC'95*, 1995.