The Mythical Traffic Matrix

Ahmed Mansy and Constantine Dovrolis
College of Computing
Georgia Institute of Technology
{amansy,dovrolis}@cc.gatech.edu

In the last few years significant progress has been made in the area of network capacity provisioning and traffic engineering. Most of the proposed methods in the literature assume knowledge of the intradomain traffic matrix. Several methodologies have been proposed for estimating traffic matrices. For instance, the authors in [1] propose a technique that combines flow level measurements collected at all ingress links with reachability information about the egress links. An information theoretic approach to estimate traffic matrices was presented in [3], where the ill-posed traffic matrix estimation problem is solved using regularization based on entropy penalization. The *tomogravity* model, presented in [4], combined the gravity model and standard tomography methods to estimate the traffic matrix.

As for traffic engineering, several routing techniques have been developed to achieve near-optimal routing taking into account some knowledge about the traffic matrix. In [8], the authors show that optimizing the OSPF link weights for a given set of demands is NP-hard. Assuming that they have an estimate of the traffic matrix, the authors proposed a heuristic for setting the link weights to optimize network performance. In [9], the same authors extended their earlier work to include IS-IS and to handle limited variations in the traffic demands due to either periodic changes or link failures. The notion of *oblivious* routing was introduced in [10], where the authors presented a routing technique that is robust to variations in the demands. Trying to deal with the dynamic nature of the traffic demand in an IP network, the authors in [7] proposed the idea of maintaining a set of routes that are good for a number of different traffic matrices and they found an optimal set of routes to minimize the expected cost. [11] proposed a class of algorithms that optimize the expected case while providing worst-case guarantees for unexpected scenarios.

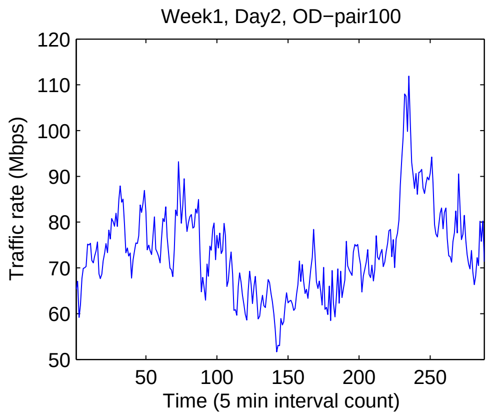

Figure 1: A time series for an OD flow from the Abilene dataset.

Traffic matrix estimation techniques work under the assumption that the Origin-Destination (OD) demands are stationary and statistically well-behaved (i.e., limited variance and short-range correlations). Otherwise, it will not be possible to do robust traffic engineering or capacity provisioning. We have analyzed publicly available datasets from the Abilene and Geant networks (see Figures 1 and 2) to examine the dynamic behavior of traffic matrix time series. These datasets give the amount of traffic in each OD flow in 5-minute (or 15-minute) intervals over a period of several weeks. Our main qualitative observation is that the traffic matrix time series are highly non-stationary, even in the timescales of a few hours. Frequent level-shifts, spikes that last for minutes or hours, and gradually increasing/decreasing rates are the rule rather than the exception. Even if we focus on time periods that can be modeled as stationary, we observed that the OD flow time series exhibit heavy-tailed variations and Long Range Dependency (LRD), meaning that attempts to statistically bound the

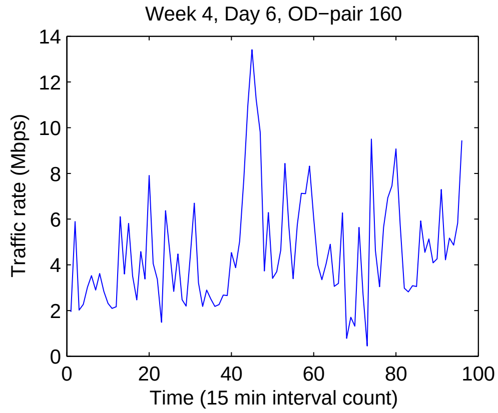

Figure 2: A time series for an OD flow from the Geant dataset.

variability of the traffic matrix can be problematic.

We believe that the behavior we observed at the Abilene and Geant datasets is true in all IP networks. Actually, our conjecture is that the great variability in the traffic matrix time series is an invariant, i.e., a fundamental property that is always true. The reasons behind this invariant are related to the heavy-tailed nature of the flow size distribution, the LRD characteristics of the traffic on any single link, and the large variability in the per-flow throughput. Consequently, as opposed to telephone networks, where the concept of a traffic matrix is well-established, we believe that traffic matrices are not a particularly useful concept in IP networks.

Further, the problem with traffic matrices in IP networks is not only due to temporal correlations and burstiness. Additionally, the presence of spatial correlations between OD flows that have the same ingress or egress node can also be a serious issue. Suppose that a given content provider suddenly becomes very popular. In that case, all the traffic matrix elements that correspond to that ingress node will increase at the same time. Such spatially-correlated traffic bursts can cause an even bigger problem to traffic engineering or capacity provisioning than the presence of independent “noise” in individual OD flows.

But if traffic engineering cannot rely on traffic matrices, then how is it possible to provide performance-based SLAs, avoid congested hot spots, or do load balancing in IP networks? We advocate the use of measurement-driven dynamic routing. In other words, routers can monitor the performance and available capacity in edge-to-edge paths and/or individual links, and adjust intradomain routing to meet the given SLA or load balancing objectives. This approach does not require traffic matrix estimates and precomputation of link weights. Our position is that the adaptive nature of measurement-driven routing makes this approach more appropriate for the highly unpredictable nature of traffic matrices in IP networks. In the past, measurement-driven routing proposals have been easily dismissed arguing that they will lead to instability, race conditions, oscillations, etc. Even though naive algorithms are indeed prone to such pathologies, we believe that it is possible to design stable and efficient measurement-based routing schemes. This is the objective of our ongoing research.

References

[1] P. 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, 2001.

[2] A. Medina, N. Taft, K. Salamatian, S. Bhattacharyya and C. Diot, "Traffic Matrix Estimation:Existing Techniques and New Directions." ACM SIGCOMM, 2002.

[3] Y. Zhang, M. Roughan, C. Lund, and D. Donoho, "An Information-Theoretic Approach to Traffic Matrix Estimation". ACM SIGCOMM, 2003.

[4] Y. Zhang, M. Roughan, N. Duffield and A. Greenberg, "Fast Accurate Computation of Large-Scale IP Traffic Matrices from Link Loads", ACM SIGMETRICS, 2003.

[5] V. Erramilli, M. Crovella, and N. Taft, "An Independent-Connection Model for Traffic Matrices",Internet Measurement Conference, 2006.

[6] A. Soule, A. Lakhina, N. Taft, K. Papagiannaki, K. Salamatian, A. Nucci, M. Crovella, and C. Diot, "Traffic Matrices: Balancing Measurements, Inference and Modeling", ACM SIGMETRICS, 2005.

[7] C. Zhang, Y. Liu, W. Gong, J. Kurose, R. Moll, and D. Towsley, "On Optimal Routing with Multiple Traffic Matrices", in Proc. of the 24th IEEE Conf. on Computer Communications (INFOCOM), 2005.

[8] B. fortz, and M. Thorup, "Internet Traffic Engineering by Optimizing OSPF Weights", in Proc. 19th IEEE Conf. on Computer Communications (INFOCOM), 2000.

[9] B. fortz, and M. Thorup, "Optimizing OSPF/IS-IS Weights in a Changing World", in IEEE Journal on Selected Areas in Communications, 2002.

[10] D. Applegate, and E. Cohen, "Making Intra-Domain Routing Roust to Changing and Uncertain Traffic Demands: Understanding Fundamental Tradoffs", ACM SIGCOMM, 2003.

[11] H. Wand, H. Xie, L. Qiu, Y. Yang, Y. Zhang, and A. Greenberg, "COPE: Traffic Engineering in Dynamic Networks", ACM SIGCOMM, 2006.