

# Experience in Measuring Backbone Traffic Variability: Models, Metrics, Measurements and Meaning

Matthew Roughan, Albert Greenberg, Charles Kalmanek, Michael Rumsewicz,  
Jennifer Yates and Yin Zhang

*Abstract*—Understanding the variability of Internet traffic in backbone networks is essential to better plan and manage existing networks, as well as to design next generation networks. However, most traffic analyses that might be used to approach this problem are based on detailed packet or flow level measurements, which are usually not available throughout a large network. As a result there is a poor understanding of backbone traffic variability, and its impact on network operations (e.g. on capacity planning or traffic engineering).

This paper introduces a metric for measuring backbone traffic variability that is grounded on simple but powerful traffic theory. What sets this metric apart, however, is that we present a method for making practical measurements of the metric using widely available SNMP traffic measurements. Furthermore, we use a novel method to overcome the major limitation of SNMP measurements – that they only provide link statistics. The method, based on a “gravity model”, derives an approximate traffic matrix from the SNMP data. In addition to simulations, we use more than 1 year’s worth of SNMP data from an operational IP network of about 1000 nodes to test our methods. We also delve into the degree and sources of variability in real backbone traffic, providing insight into the true nature of traffic variability.

Despite a significant amount of research addressing Internet traffic models [1], [2], [3], [4], there is not yet widespread agreement about the characteristics of backbone Internet traffic. This problem is exacerbated by exaggerated reports on Internet traffic growth and variability [5], [6], by the challenges associated with Internet traffic measurements [7], and a lack of understanding of the applicability of results such as the discovery of self-similarity in traffic [1], [2], [3]. For instance, in [5], dire claims are made on the basis of the notion that large volumes of traffic *slosh* around the Internet in a highly irregular way.

Obtaining the data necessary to develop an accurate and current view of backbone traffic requires significant investment in measurement infrastructure. Nonetheless, understanding Internet backbone traffic is crucial for evolving the Internet architecture, doing capacity planning, traffic engineering, and meeting service level agreements. In particular, our investigation was specifically motivated by the question: to what extent does traffic variability justify the

need for a re-configurable optical network below the IP layer to provide bandwidth management. Such an optical network would allow IP routers equipped with the appropriate interfaces to request additional point-to-point capacity when needed, and to reconfigure existing capacity between routers [8], [9], [10]. Routers might need additional capacity due to congestion resulting from any of a number of causes: major events (September 11th), re-routing events triggered by failures, transient overloads due to Denial of Service (DoS) attacks or flash crowds, or externally induced traffic shifts from peer networks. Alternatively, we can view this problem through the lens of over-provisioning, namely to what extent does the IP layer need to be over-provisioned to meet its service layer agreements with high reliability.

We address the problem of backbone traffic variability by looking at aggregate link statistics collected via SNMP on a large ISP backbone. From these statistics it is clear that the traffic has both daily and weekly periodic components, as well as a longer-term trend. Superimposed on top of these components are shorter time scale stochastic variations. Given these characteristics, we develop a simple, but powerful stochastic model for backbone traffic (based on the Norros model [11]), and then use that model to derive an empirical metric referred to here as the peakedness parameter, that provides a measure of the traffic variability. We believe that this metric will be useful to network operators in both architecture evolution and traffic management, e.g., allowing network operators to determine whether (or when) it makes sense to layer IP over a re-configurable optical network, assisting in provisioning backbone capacity, tuning OSPF links weights, etc. An important feature of this model is parsimony – only one parameter is required to describe the most important features of the stochastic variation in the traffic, and this parameter can be estimated from standard SNMP traffic measurements.

We apply this stochastic traffic model in the context of a large backbone network. Ideally, we would obtain a backbone traffic matrix using detailed flow measurements on network access routers, as in [7], and fit the traffic to the stochastic model. However, as mentioned in [7], many access routers are not currently able to continuously collect

M. Roughan, A. Greenberg, C. Kalmanek, J. Yates and Y. Zhang are with AT&T Labs – Research, 180 Park Av., Florham Park, NJ, 07932. {roughan,albert,crk,jyates,yzhang}@research.att.com.

M. Rumsewicz is with Telic Australia, mpr@telicaustralia.com

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

IMW'02, Nov. 6-8, 2002, Marseille, France

Copyright 2002 ACM ISBN 1-58113-603-X/02/0011 ...\$5.00

flow statistics. A number of innovative methods have been proposed for deriving traffic matrices [12], [13], [14], [15], but the limitations of SNMP data, and the size of the network make these quite difficult to apply here.

Instead, we use a *gravity model* to analyze the SNMP aggregate link statistics to derive a backbone region to region traffic matrix. 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. In a geographic gravity model for cities, for example, the relative strength of the interaction between two cities is proportional to the product of the populations divided by the distance squared. Such models provide surprisingly accurate estimates of telephone traffic exchanged between areas codes (see [16] and the references therein). In our gravity model for backbone traffic, we compute the fraction of the total traffic entering (leaving) the backbone to (from) each region or Point-of-Presence (PoP). For each PoP, we then take this fraction to be the fraction of traffic sourced (sunked) from every other PoP to that PoP. This gives an approximation of the traffic matrix. While the gravity model does not capture details of the actual traffic demands, the impact of peering policies on traffic flow, etc., it is relatively simple to compute and is surprisingly accurate.

A major insight of this paper is that traffic is predominantly regular and predictable, though it does have a significant stochastic component. The results show that although backbone traffic is highly non-stationary, showing significant daily and weekly variations. This periodic variation shows that traffic engineering based on long term averages is not sufficient. A key insight of this paper is that large deviations from traffic predictions are rare. The observed cases are due to large, transient events, such as flash crowds, network failures, and natural disasters. Most normal variation has peakedness parameter  $a$  in the range 0.5-3.0 Mbs for 5 minute SNMP measurements. This value of  $a$  appears to represent relatively stable traffic, however we note that  $a$  can be significantly larger even when we exclude obvious transient events. At the very least this provides a realistic set of parameter values for simulations of backbone traffic.

This paper also demonstrates that gravity models are a natural and powerful starting point for deriving traffic matrices from link statistics. Given today's difficulty in obtaining flow-level measurement data at every edge router in a large backbone, we believe that gravity models are a pragmatic tool for providing traffic matrices until we have more reliable sources of detailed data. We are pursuing generalizations of the gravity model that will allow us to differentiate traffic with finer grained detail. We also plan

refinements of the traffic matrix by combining the gravity model approach with flow level measurements where available.

Our original motivation for this work was to look at the potential benefits of building IP backbones on top of a re-configurable optical network. Though the diurnal variations in traffic are significant, these are tightly coupled across the North American continent, and so do not present an opportunity for temporal sharing of capacity. Furthermore, the stability of the stochastic component of backbone traffic suggests that the case for a re-configurable optical network layer based solely on IP traffic variations is weak. However, it may still make sense to use a re-configurable optical network to deal with traffic load changes resulting from IP layer re-routing due to failures [17]. This is a subject of on-going work.

#### ACKNOWLEDGMENTS

We would like to thank Fred True, Joel Gottlieb, and Tim Griffin for their work in collecting, and managing the data used here, and Bob Doverspike, and Panita Pongpaibool for stimulating discussion of optical reconfiguration.

#### REFERENCES

- [1] W. E. Leland, M. S. Taqqu, W. Willinger, and D. V. Wilson, "On the self-similar nature of Ethernet traffic (extended version)," *IEEE/ACM Transactions on Networking*, vol. 2, pp. 1-15, Feb 1994.
- [2] W. Willinger, V. Paxson, and M. S. Taqqu, "Self-similarity and heavy tails: Structural modeling of network traffic," in *A Practical Guide to Heavy Tails: Statistical Techniques and Applications* (R. Adler, R. Feldman, and M. S. Taqqu, eds.), pp. 27-53, Birkhauser, Boston, 1998.
- [3] V. Paxson, "Empirically-derived analytic models of wide-area TCP connections," *IEEE/ACM Transactions on Networking*, vol. 2, no. 4, pp. 316-336, 1994.
- [4] J. Cao, W. S. Cleveland, D. Lin, and D. X. Sun, "The effect of statistical multiplexing on internet packet traffic: Theory and empirical study," tech. rep., Bell Labs, 2001. Available at <http://cm.bell-labs.com/cm/ms/departments/sia/wsc/publish.html>.
- [5] B. S. Arnaud, "Current optical network designs may be flawed," *Optical Networks Magazine*, vol. 2, March/April 2001.
- [6] J. Mooney, "Gaining the edge in flexible metro service provisioning," *Lightwave*, February 2001. Available at [http://www.sycamorenet.com/corporate/articles/reprintedarticle.asp?comm&and=live&news\\_item\\_id=437](http://www.sycamorenet.com/corporate/articles/reprintedarticle.asp?comm&and=live&news_item_id=437).
- [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. An earlier version appeared in Sigcomm'00.
- [8] J. Y. Wei, "IP over WDM network traffic engineering approaches," in *invited talk Optical Fiber Commun. Conf. (OFC)*, 2002.
- [9] D. Awduche and Y. Rekhter, "Multiprotocol lambda switching: combining mpls traffic engineering control with optical crossconnects," *IEEE Communications Magazine*, vol. 39, pp. 111-116, March 2001.
- [10] D. Banerjee, "Wavelength routed optical networks: linear formulation, resource budgeting tradeoffs and a reconfiguration study," in *IEEE INFOCOM'97*, pp. 269-276, 1997.
- [11] I. Norros, "A storage model with self-similar input," *Queueing Systems*, vol. 16, pp. 387-396, 1994.
- [12] Y. Vardi, "Network tomography," *Journal of the American Statistical Association*, March 1996.
- [13] C. Tebaldi and M. West, "Bayesian inference on network traffic," *Journal of the American Statistical Association*, June 1998.
- [14] J. Cao, D. Davis, S. V. Wiel, and B. Yu, "Time-varying network tomography," *Journal of the American Statistical Association*, December 2000.
- [15] N. G. Duffield and M. Grossglauser, "Trajectory sampling for direct traffic observation," *IEEE/ACM Trans. on Networking*, vol. 9, pp. 280-292, June 2001.
- [16] J. Kowalski and B. Warfield, "Modeling traffic demand between nodes in a telecommunications network," in *ATNAC'95*, 1995.
- [17] P. Pongpaibool, R. Doverspike, M. Roughan, and J. Gottlieb, "Handling ip traffic surges via optical layer reconfiguration," in *to appear in OFC*, 2002.