A Statistical Framework for Efficient Monitoring of End-to-End Network Properties

David Chua
Mathematics & Statistics Dept.
Boston University
dchua@math.bu.edu

Eric D. Kolaczyk
Mathematics & Statistics Dept.
Boston Univeristy
kolaczyk@math.bu.edu

Mark Crovella
Computer Science Dept.
Boston University
crovella@cs.bu.edu

ABSTRACT

Network service providers and customers are often concerned with aggregate performance measures that span multiple network paths. Unfortunately, forming such network-wide measures can be difficult, due to the issues of scale involved. As a result, it is of interest to explore the feasibility of methods that dramatically reduce the number of paths measured in such situations while maintaining acceptable accuracy.

In previous work [4] we have proposed a statistical framework for efficiently addressing this problem. The key to our method lies in the observation and exploitation of the fact that network paths show significant redundancy (sharing of common links).

We now make three contributions in [3]: (1) we generalize the framework to make it more immediately applicable to network measurements encountered in practice; (2) we demonstrate that the observed path redundancy upon which our method is based is robust to variation in key network conditions and characteristics, including the presence of link failures; and (3) we show how the framework may be applied to address three practical problems of interest to network providers and customers, using data from an operating network.

Categories and Subject Descriptors: C.2.3 [Computer Communication Networks]: Network Operations—*network monitoring*; C.4 [Performance of Systems]: Measurement Techniques.

General Terms: Measurement, Performance.

Keywords: Algorithms, Statistical analysis, Networking.

1. INTRODUCTION

In many situations it is important to obtain a network-wide view of path metrics such as latency and packet loss rate. For example, in overlay networks regular measurement of path properties is used to select alternate routes. At the IP level, path property measurements can be used to monitor network health, assess user experience, and choose between alternate providers, among other applications. Typical examples of systems that perform such measurements include the NLANR AMP project, the RIPE Test-Traffic Project, and the Internet End-to-end Performance Monitoring project.

Unfortunately, extending such efforts to large networks can be difficult, because the number of network paths grows as the square of the number of network endpoints. Initial work in this area has found that it is possible to reduce the number of end-to-end measurements to the number of "virtual links" (identifiable link subsets)—which typically grows more slowly than the the number of paths—and yet still recover the complete set of end-to-end path properties exactly [2, 5].

This result is based on a linear algebraic analysis of routing matrices, as pioneered in [5]. A routing matrix is a binary matrix that specifies which links appear in which end-to-end paths. Such a matrix G has size (# paths) × (# links), and $G_{i,j} = 1$ if and only if link j is found along the route taken by path i. The rank of G, which is generally equal to the number of independent paths in the network, tends to be much smaller than the total number of paths. Since a maximal set of such independent paths can be used to reconstruct any other path in the network, it is sufficient to monitor only this set. Algorithms for choosing such a set have been developed [5, 2], based on linear algebraic methods of subset selection.

In previous work [4] we have proposed a general approach to extending this strategy. First, we note that routing matrices encountered in practice generally show significant sharing of links between paths. One implication is that routing matrices have small *effective* rank compared to their actual rank—a small set of eigenvalues of $G^T G$ tend to be much larger than the rest. Next, we propose a statistical framework for approximating summary path metrics which can be used to exploit the property of low effective rank of routing matrices.

The statistical framework we have developed draws on linear model theory. Linear model theory is concerned with the statistical inference of quantities that are related through some underlying linear relationship. In this case, we are concerned with the inference of path properties $y = [y_1, y_2, \dots]^T$, that are related through their linear dependence $y = Gx$ on link properties $x = [x_1, x_2, \dots]^T$. Linear model theory provides a well-developed set of tools that allow one in many cases to construct and analyze procedures for making optimal linear inferences.

While our previous work has developed a basic framework for attacking the problem, there are a number of hurdles, both theoretical and practical, separating our simple model from application to realistic network settings. Our goal in [3] is to eliminate those hurdles. We do so by a combination of analysis, experimentation, and application to real network data. In the process, we demonstrate the power of the

SIGMETRICS'05, June 6–10, 2005, Banff, Alberta, Canada.
ACM 1-59593-022-1/05/0006.

resulting framework for solving problems of real interest to network operators and users.

2. EFFICIENT MONITORING

Our first contribution concerns the generality of our methods. Our initial, simple analysis relied on a model in which link metrics are assumed to have equal variance. This property does not hold in practice; in reality, the variance of link delay (for example) can vary considerably from link to link, depending on phenomena such as localized congestion. It is important to understand how the structure of our proposed solution changes when link metrics have unequal variance, as occurs in practice. We analyze this more general case and show how to incorporate unequal variances into our framework.

3. ROBUSTNESS OF PATH REDUNDANCY

The opportunity for efficient network-wide measurement derives from the small effective rank of routing matrices. However the factors that determine this phenomenon are not fully understood. Thus it is important to ask whether the low effective rank of routing matrices is a robust property. This property can be affected in at least two ways: first, if links fail, the degree of link sharing among paths is affected; and second, if link metrics have unequal variance, the benefits of link sharing may be reduced. Evaluating how effective rank changes, we find that even in the presence of significant differences in variance across links, unequal link variance does not generally diminish the utility of our methods; even if a link or two goes down.

4. APPLICATIONS

Since our goal is to move our methods into practical use, and to demonstrate their utility, it is appropriate to apply them to data obtained from a real network. For that purpose we employ a large set of measurements from the NLANR Active Measurement Project (AMP). These are comprised of per-hop delay measurements; we select the subset of these measurements traversing the Abilene network. We assess our analytic assumptions as well as the actual performance of our methods using this data. For a detailed presentation and analysis of the results, please see [3].

4.1 Monitoring a Network-wide Average

The first problem is that of obtaining network-wide averages of per-path delay, as would be needed for a network "health" measure. We show that a very accurate estimate of average path delay (with relative error typically less than 1%) is obtained using as few as one-tenth of the paths needed for exact measurement. In addition, we show that our extended methods that incorporate unequal link variance, through the use of a diagonal covariance matrix, are more accurate than the simpler methods in our previous work.

4.2 Anomaly-detection

Our second problem concerns the detection of anomalies within the network. We are interested in detecting when average network performance exceeds a threshold, such as three standard deviations from its mean. Reasonable detection of such spikes in the delay time can be accomplished using fewer than one-third of the paths needed for exact measurement. In Figure 1, circles have been placed along the actual path delay time series at the epochs that were flagged as anomalies in the predicted time series. We identify most of the major spikes, resulting in a true positive rate of 81%, while the false positive rate is only 8%. Furthermore, most of these false positives seem to occur at lesser spikes in the actual delay data.

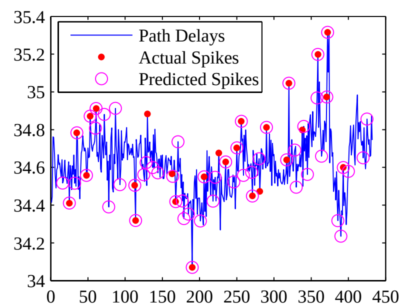

Figure 1: Predicted versus actual spikes in delay.

4.3 Subnetwork Comparison

Our final problem concerns the problem of selecting between two network ingress points. In this setting, a network customer or peer has multiple connection points to the network of interest and needs to choose between them for accessing a given set of destinations. Research such as [1] has shown that significant performance benefits can accrue when a customer chooses the best path from among a set of alternatives made available by multihoming. We show that only one-sixth the paths needed for exact results need to be measured in order to accurately choose between the alternatives. Comparing the average delays for paths entering the network in Chicago against Atlanta, the ingress point with a shorter mean delay is correctly predicted over 88% of the time with most errors occuring when the difference in delays is near zero.

5. ACKNOWLEDGEMENTS

The authors would like to thank Anukool Lakhina (Boston University, email: `anukool@cs.bu.edu`) for his help with collecting and processing the data.

6. REFERENCES

[1] A. Akella, S. Seshan, and A. Shaikh. Multihoming performance benefits: An experimental evaluation of practical enterprise startegies. In *USENIX*, 2004.

[2] Y. Chen, D. Bindel, and R. H. Katz. Tomography-based overlay network monitoring. In *Proceedings of the 2003 ACM SIGCOMM conference on Internet measurement*, pages 216–231. ACM Press, 2003.

[3] D. B. Chua, E. D. Kolaczyk, and M. Crovella. A statistical framework for efficient monitoring of end-to end network properties. http://arxiv.org/abs/cs.NI/0412037.

[4] D. B. Chua, E. D. Kolaczyk, and M. Crovella. Efficient estimation of end-to-end network properties. In *Proceedings of IEEE INFOCOM 2005*, 2005.

[5] Y. Shavitt, X. Sun, A. Wool, and B. Yener. Computing the unmeasured: An algebraic approach to internet mapping. In *IEEE INFOCOM 2001*, April 2001.