GRE Encapsulated Multicast Probing: A Scalable Technique for Measuring One-Way Loss

[Extended Abstract]

Yu Gu[‡], Lee Breslau[*], Nick Duffield[*], Subhabrata Sen[*]
[*]AT&T Labs – Research,Florham Park, NJ 07932
[‡]Dept. of Comp. Science,University of Massachusetts, Amherst
{breslau,duffield,sen}@research.att.com, yugu@cs.umass.edu

ABSTRACT

We develop techniques for estimating one-way loss from a measurement host to network routers which exploit commonly implemented features on commercial routers and do not require any new router capabilities. The work addresses the problem of scalably performing one-way loss measurements across specific network paths.

Categories and Subject Descriptors

C.2.3 [**Network Operations**]: Network monitoring

General Terms

Measurement, Performance

Keywords

Measurement, Monitoring, Performance, Multicast, One-way loss

1. INTRODUCTION

Internet service providers increasingly wish to monitor the performance of customer traffic within their networks. One particularly sought-after capability is that of measuring the one-way performance along a specific path (or subpath) from a network element A to an element B in the provider backbone part of a provider-based VPN (Figure 1). Recent work has proposed using *remote virtual monitoring* [1, 2]. A *single* centralized monitoring host establishes virtual connections (e.g., via encapsulation) to routers in the network to enable remote monitoring of customer performance. In the context of VPNs, the monitor can send and receive either unicast (for unicast VPN) or multicast (for multicast VPN) traffic across the provider backbone. Within the provider network, this traffic appears like, and is treated the same as, the corresponding traffic sent by the customer into the VPN. As an illustration, in Figure 1, the monitor M *virtually attaches* to points A and B. M transmits probe packets to A along path P_{MA}. From A to B, the probes traverse the path P_{AB} that we want to monitor. Finally they are sent back from B to M along the path P_{BM}. By comparing what

SIGMETRICS'07, June 12–16, 2007, San Diego, California, USA.
ACM 978-1-59593-639-4/07/0006.

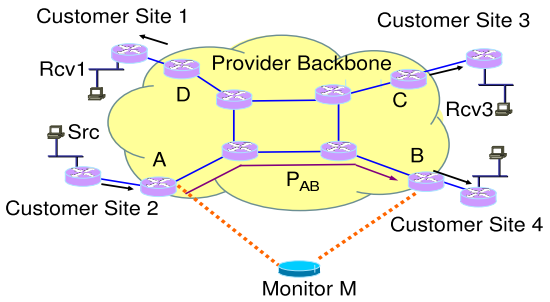

Figure 1: Monitoring one-way performance in VPNs

was sent to what was eventually recieved, M can measure the one-way performance along the directed path P_{MABM}. The key advantage of this approach is that it requires a *single* specialized monitor located anywhere, making it easy to deploy and manage such a system.

A critical component of these systems is the ability to factor out the performance contributions due to paths P_{MA} and P_{BM} from the overall performance observed on P_{MABM} to determine the performance on the path of interest P_{AB}. The measurement problem then devolves to measuring one-way loss for paths P_{MA} and P_{BM}, where one end of the path has a specialized measurement device and the other end is a production router. The existing remote virtual monitoring systems [1, 2] use round trip measurements between the monitoring device and the network routers and assign half the measured loss to each direction of the path. Since path properties are known to be asymmetric, we are motivated to explore more accurate techniques for factoring out the performance of the one-way paths.

2. CONTRIBUTIONS

We develop novel techniques for estimating one-way loss from a measurement host to network routers, which exploit commonly implemented features on commercial routers and do not require any new router capabilities. By exploiting these features, the expense of deploying special purpose measurement hosts can be avoided while overcoming the shortcomings of two-way measurements between a source measurement host and a router. In this extended abstract we overview key components of our approach and describe three specific measurement scenarios in which these techniques may be deployed. We point the reader to a full ver-

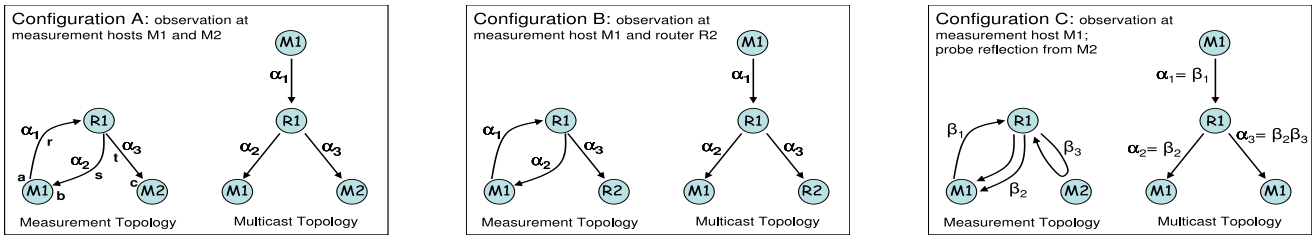

Figure 2: Three measurement configurations and associated logical multicast topologies

sion of the paper [5] for a detailed description and evaluation. There are three components. Firstly, GRE tunneling is used to control the paths followed by measurement traffic through the network. Secondly, innovative probing methods, coupled with standard measurement capabilities such as NetFlow and SNMP, are used to isolate the performance of groups or even individual measurement packets. Thirdly, we exploit and extend tomographic inference methods in order to extract the performance of probe traffic on customer paths within the network. This combination yields a powerful yet lightweight method to determine customer performance within the network.

We propose three specific measurement configurations; these are illustrated in Figure 2. In each figure α_i or β_i are the packet transmission probabilities associated with tunnels. The three configurations differ in the type of loss information that can be measured and for each we use suitable inference algorithms.

Configuration A: Probe Collection at Two Measurement Hosts. Two GRE tunnels r and s are created from $M1$ to $R1$, using distinct IP addresses at $M1$, denoted $M1(a)$ and $M1(b)$. A GRE tunnel t is created from $M2$ to $R1$ using IP address c on $M2$. Multicast routing is enabled on $R1$. Host $M1$ joins a multicast group through the tunnel s and host $M2$ joins the same multicast group through tunnel t. Host $M1$ launches probes into tunnel r, copies of which are multicast down tunnels s and t, to be collected at $M1(a)$ and $M2(c)$. In this configuration $M1$ and $M2$ can measure individual packets. Therefore inference can be performed using the individial packet MINC estimator [3].

Configuration B: Probe Collection at a Measurement Host and a Router. There are two variants of this configuration. In the first variant, the additional router $R2$ is positioned in the path between $R1$ and $M2$ in order to observe packets in transit. In the second variation, $R2$ substitutes for $M2$, and joins the multicast group using a static IGMP join, then observes packets destined to it. In this configuration, the ability to distinguish individual packets is limited by the router level measurement at $R2$, which is assumed to report only aggregates. Therefore we must use the aggregate MINC estimator [4] on packet aggregates reported by $R2$.

Configuration C: Observation at a Measurement Host and Reflection. This configuration is based upon configuration A, except that host $M2$, instead of measuring each packet, responds to its receipt by multicasting a second packet into tunnel t, subsequently to be received by $M1$. The configuration advantage of this approach is that only M1 collects data, thus avoiding coordination of measurement collection between multiple hosts. In Figure 2, β_3 is the transmission probability on the round-trip path $R1 \rightarrow M2 \rightarrow R1$. Here $M1$ can measure individual packets and inference can be performed using the reflected MINC estimator we developed as a variant of the standard MINC estimator.

We have implemented the measurement and inference algorithms in a prototype system and have tested the implementation in a research testbed. Concerning the performance of our estimators: the relative errors of all methods seem acceptable for a feasible measurement bandwidth over a period of about a minute. The results indicate that under the conditions tested, and except for some potential topological biases, estimator performance is sufficient to estimate loss rates of between 0.02% and 5% with no worse than at 7% error on average. The presence of temporal correlation in loss due to queueing and/or TCP source behavior does increase the variance of estimators that use multipacket event frequencies for estimation, but the resulting estimator accuracy is no worse than twice as bad expected. Spatial correlations due to background traffic processes, although having some impact on estimation accuracy, likewise did not unduly increase estimation variance.

3. CONCLUSIONS

We presented techniques for estimating one-way loss from a measurement host to network routers which do not require any new router capabilities. While our work is motivated by the remote virtual performance monitoring systems that have been proposed for VPNs [1, 2], we believe our results have wider applicability. As an example, the architecture in the remote virtual monitoring systems can be applied to generic IP network monitoring, which would also require a general solution to the problem of estimating one-way loss.

4. REFERENCES

[1] L. Breslau, C. Chase, N. Duffield, B. Fenner, Y. Mao, and S. Sen. Vmscope – a virtual multicast VPN performance monitor. *ACM SIGCOMM Workshop on Internet Network Management (INM)*, 2006.

[2] H. Burch and C. Chase. Monitoring link delays with one measurement host. *SIGMETRICS Performance Evaluation Review*, 33(3):10–17, 2005.

[3] R. Caceres, N. Duffield, J. Horowitz, and D. Towsley. Multicast-based inference of network-internal loss characteristic. *IEEE Transactions in Information Theory*, 45:2462–2480, 1999.

[4] N. Duffield, V. Arya, R. Bellino, T. Friedman, J. Horowitz, D. Towsley, and T. Turletti. Network tomography from aggregate loss reports. In *Performance 2005*, 2005.

[5] Y. Gu, N. G. Duffield, L. Breslau, and S. Sen. GRE encapsulated multicast probing: A scalable technique for measuring one-way loss. In submission, 2007.