

Remote Network Bottleneck Diagnosis

Paper Number 171

Earl Eiland (eee@nmt.edu)

Nathan Campbell (nate@nmt.edu)

Billy Byler (cougar@sec.org)

Harley Kozushko (hkozushk@nmt.edu)

Liang Xiaoguang (liangxiaoguang@yahoo.com)

Computer Science Department

New Mexico Institute of Mining and Technology

801 Leroy Place, Socorro, NM, USA

Phone: 1-505-835-5170 or 1-505-835-6008

Introduction:

By their very nature, networks are highly distributed systems. This means that nodes and subnets in any one system may be great distances apart. Managing such a distributed system has many challenges, not the least of which is maintaining acceptable performance over channels owned and/ or managed by third parties. Troubleshooting and preventative maintenance on such highly distributed systems depends upon the cooperation of these independent entities. Unfortunately, cooperation may be spotty or non-existent if these vendors believe the requested information is proprietary or increases their exposure to attack. Vendors may even provide mis-information rather than admit a problem exists. In this environment, system administrators need to be able to independently and remotely acquire or verify data about paths outside of their control. A recently developed technique, Network Spectroscopy [1], appears to provide the means to characterize paths at a level of detail not possible any other way. We apply Network Spectroscopy to the problem of remotely identifying the cause of a network bottleneck. Preliminary results discussed below, and shown in Graph 1, are encouraging.

Background:

Previous work established a method of locating a path bottleneck [2]. This is useful in that it allows routing around problem areas. Solving the problem, however, requires more information. Bandwidth bottlenecks can occur in two areas. Either the traffic load is greater than the channel capacity across a link in the path, or a node (router or switch) has insufficient processing or buffering capacity. Being able to pinpoint the location of a bottleneck is only part of analyzing a problem. The cause must be diagnosed before a cure can be affected.

Recently, network spectroscopy has emerged as a useful tool for fine grained analysis of paths across a network. As explained by the techniques architects [1],

“Spectroscopy is different from previous approaches in that it emphasizes extracting information from: (1) packet timing jitters, which most other techniques interpret as noise; or (2) fine-grained delay quantizations, such as cell or slot times in TDM (time-division multiplexed) infrastructures.”

Network spectroscopy has been used successfully to explore a wide range of remote diagnostic problems, such as characterizing DoS attacks [3] and identifying the sources of spurious DNS update requests from private networks [4]. A recent study examining ATM effects on DSL modem traffic [5] discovered timing jitters the researchers attributed to queuing effects .

This Papers Contribution:

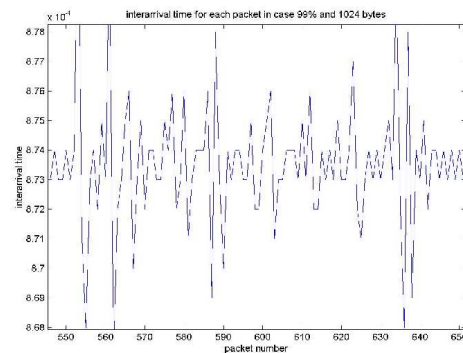

We believe the above-mentioned queuing effects can be exploited to remotely identify bottleneck causes. This work explores the use of Network spectroscopy as a tool to determine whether a bottleneck is due to insufficient channel capacity, or insufficient buffer capacity within a router or switch.

Method:

We have assembled a source-router-sink testbed. Tests are being run using deterministically and stochastically generated traffic to examine operating characteristics in three states: normal operation, and operation with buffer- and channel capacity-induced bottlenecks.

Ongoing Work:

The goal of this project is to identify time-series properties specific to buffer- and channel capacity-induced bottlenecks. To date, we have run fixed-rate, fixed-packet size tests. A sample of our raw time-series type data, shown in Graphs 1 and 2, indicate the existence of markedly different characteristics for “normal” and full-channel operation. Ultimately, we expect to estimate the power spectral densities for each of the three operating conditions, extract their dominant frequencies, and determine their distribution parameters.



Graph 1: Sequence of packet interarrival times on a full channel (99% usage).

Graph 2: Sequence of packet interarrival times on a “normal” channel (58% usage).

The packet number ranges were randomly selected. Differences in the range of interarrival times is due to the rates at which packets were sent to create the two traffic densities. High frequency seems to dominate in the high usage test (Graph 1), while a much lower frequency dominates in the “normal” usage test (Graph 2). Switch settings used minimize queuing effects.

References:

- [1] Andre Broido, R.K., Evi Nemeth, kc claffy. *Radon spectroscopy of packet delay*. in *IEEE High-Speed Networking Workshop*. 2003. San Diego, Cal: IEEE.
- [2] Dina Katabi, C.B., *Inferring Congestion Sharing and Link Characteristics from Packet Interarrival Times*. MIT-LCS-TR828. 2001, Mass Inst. of Tech., Laboratory for Computer Science: Cambridge, Massachusetts. pp. 13.
- [3] Alefiya Hussain, J.H., Christos Papadopoulos. *A Framework for Classifying Denial of Service Attacks*. in *Applications, technologies, architectures, and protocols for computer communications*. 2003. Karlsruhe, Germany: ACM.
- [4] Andre Broido, E.N., kc claffy. *Spectroscopy of private DNS Update Sources*. in *IEEE High-Speed Networking Workshop*. 2003. San Diego, California: IEEE.
- [5] King, Ryan, *ATM Induced Quantization of Delay in DSL Modem traffic*. 2003. <http://www.caida.org/~ryan/spect/dsl.html>. Cooperative Association for Internet Data Analysis, University of California's San Diego Supercomputer Center. San Diego, California .

