

Statistical Aspects of the Analysis of Data Networks

Lorraine Denby, James M. Landwehr, Colin L. Mallows, Jean Meloche, John Tuck

Data Analysis Research, Avaya Labs
Basking Ridge, NJ 07920

Bowei Xi

Department of Statistics, Purdue University
West Lafayette, IN 47907

George Michailidis and Vijayan N. Nair

Department of Statistics, University of Michigan
Ann Arbor, MI 48109

Abstract

Assessing and monitoring the performance of computer and communications networks is an important problem for network engineers. There has been a considerable amount of work on tools and techniques for data collection, modeling, and analysis within the network research community. The goal of this paper is to present an overview of the engineering problems and statistical issues, describe recent research developments, and summarize ongoing work and areas for further research. While there are many interesting issues related to network analysis, our focus here is on estimating and monitoring network Quality-of-Service parameters. We discuss methods for estimating edge-level parameters from end-to-end path-level measurements, an important engineering problem that raises interesting statistical modeling issues. Other topics include network monitoring, network visualization, and discovering network topology. Data from a corporate network are used to illustrate the problems and techniques. As in any overview paper, the discussion is likely to be slanted towards our own research interests.

Key Words: Internet telephony, Network monitoring, Network tomography, Network visualization, Probe packets, Traceroute data, and Voice over IP.

1 Introduction

Modern computer and communications networks are complex and dynamical systems with millions of users and a broad range of applications. Business customers as well as individual users expect consistently high levels of quality, especially when they are running complex applications such as telephony and video. So one has to continuously collect, model, and analyze relevant data in order to assess and monitor network performance. There are many interesting statistical issues and challenges involved in

doing this. The goal of the present paper is to provide a background of the engineering problems and an overview of the statistical issues, recent developments, and on-going research. The paper includes a review of existing results and a discussion of some new results. Predictably, however, the discussion emphasizes our own interests.

Internet Protocol telephony (IP telephony, also known as Voice over Internet Protocol or VoIP) is an important application and provides a basis for much of our discussion in this paper. IP telephony typically involves a pair of IP phones that send streams of packets that carry voice traffic. At the sending phone, the packets are sent with regularity (for example, every 20 milliseconds) and each packet contains a segment of voice. At the receiving phone, the packets do not arrive with the same regularity because of unpredictable events in the network. The packets can be dropped by network routers when queues are full, they can be delayed due to competing traffic, or they can arrive out of order. Packet loss and the lack of regularity in the packet stream at the receiving phone can result in poor sound quality.

Data that relate to the performance of IP telephony on the network in Figure 1 will be used to illustrate various concepts throughout the paper. [See also Bearden et al. (2002a and 2002b) and Karacali et al.(2004) for other studies based on data collected from this network.] This figure shows parts of the Avaya corporate network including 37 special communication endpoints that are labeled in the figure. [The network also includes thousands of regular endpoints (e.g., users) that connect and disconnect at different nodes that are not shown here.] The circles represent network routers, and the triangles represent the special communication endpoints. The three agglomerations correspond to the Asia Pacific region, the Europe/Middle East region, and the Americas. [Note that one of the edges (close to Dubai) is shaded in black; we will return to this when we discuss network monitoring techniques in Section 4.2.]

Network engineers designing and running networks (either within corporate and campus environments or for Internet Service Providers) need data collection and analysis tools to assess Quality-of-Service (QoS) measures such as packet loss rates, delays, and jitter, and for doing network bandwidth calculations. This information is needed for several important reasons: monitoring network performance and utilization over time; drilling into problems and finding their causes; detecting congestion; planning capacity and network provisioning; and for ensuring compliance with service level agreements. The needs are especially critical with real-time applications that require very high and sustained levels of quality, such as VoIP, video streaming, video-conferencing, and on-line games.

This is a challenging problem for various reasons. First, the size of network (often much larger than the one in Figure 1) can limit the type of analyses that can be performed in practice. Second, networks are evolving entities and QoS characteristics can change rapidly, for example, as a result of load or as a result of an automatic process that is attempting to circumvent some local network problem. Finally, network engineers often do not have access to all the relevant components in the network, for example, a node that belongs to a different administrative domain or an edge that belongs to an Internet Service Provider.

Traditional approaches to network analysis have relied on detailed queueing models at the indi-

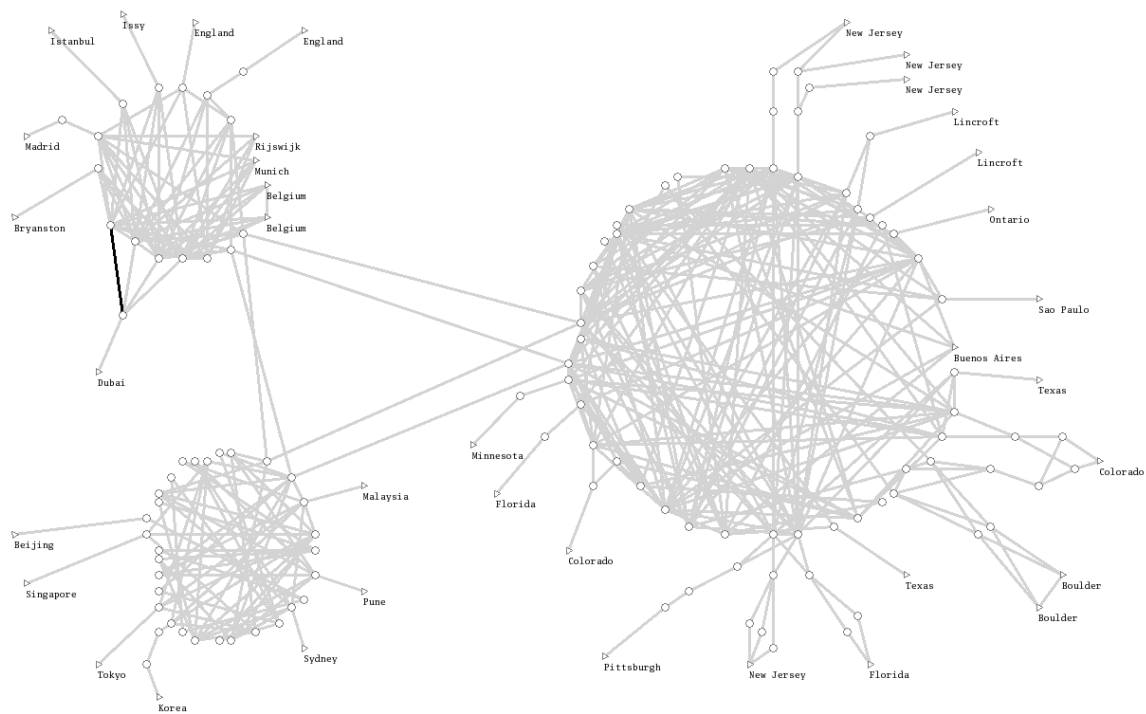


Figure 1: The Avaya Corporate Network

vidual router level. But such “local” modeling will not adequately capture the complexities and dynamic behavior of modern networks, including the fact that end-to-end results can be affected by interactions between non-adjacent network components. Expanding such local models to incorporate the behavior of even a moderately-sized network adequately would be impractical.

There have been considerable efforts in recent years to identify appropriate sources of data and develop methods of data collection, analysis tools, and techniques to address these problems. This has led to the emergence of *network tomography*, a term first introduced by Vardi (1996) in the context of estimating the origin-destination (O-D) traffic matrix. Here one is interested in estimating the intensities of traffic flowing between the O-D pairs in the network. This information is important in network management, capacity planning, and provisioning. The challenge arises from the fact that only aggregate data on traffic counts at individual nodes (e.g., number of packets going through a router) are available. The inverse problem is to recover distributions of traffic between all O-D pairs from the aggregate counts. See Castro et al. (2004) for a review of this literature and for an overview of developments in network tomography.

The focus in this paper is on estimating and monitoring the QoS of the network by actively injecting test packets, or probes, into the network from nodes located on the periphery and collecting performance measurements on the injected probes. A primary goal is to estimate edge-level QoS char-

acteristics from the end-to-end measurements. This inverse problem has sometimes been called *active network tomography*. Our overall purpose is to discuss a number of interesting statistical issues that arise in this area, including data collection, data quality, inference for the inverse problems, network monitoring, and the analysis of large-scale networks (graphs).

The paper is organized as follows. Section 2 gives background on data networks, protocols, and network utilities that can be used for our purposes to provide measurements. Section 3 deals with estimation of the edge-level parameters from end-to-end path-level measurements, which is a very interesting inverse problem with many facets. It introduces a novel approach for edge level estimation from end-to-end statistics based on a penalized likelihood which combines both traceroute and RTP data. Section 4 discusses several related topics: methods for monitoring network performance over time and using the information to diagnose the edges or subnetworks where problems occur; and using visualizations to assess network performance. Section 5 briefly describes several issues involved in discovery, analysis, and visualization of network topology. Throughout the paper, relevant literature and references are noted.

2 Background and Data Sources

2.1 Network Preliminaries




Figure 2: Left panel: Schematic of the UNC Network; Right Panel: Topology of the UNC Network for the VOIP Study.

A network can be represented by a directed graph $G = (V, E)$ where V is the set of nodes (routers) and E is the set of edges. Even though data networks can be specified at several different layers, this paper concentrates on the IP layer or Layer 3, in which a path consists of the source and destination nodes and the intervening IP routers (for a detailed exposition on network layers and protocols, see Peterson and Davie (2003)).

The left panel in Figure 2 is an example; it shows part of the University of North Carolina, Chapel Hill network with 19 end points and 4 internal nodes. When data are sent from one location to another in a packet-switched network, the file's content is first broken into pieces, called packets. Additional information such as origin-destination information, reassembly instructions (such as sequence numbers), and error-correcting features are added to the packet. The origin-destination information is used by the network routers and switches to deliver the packets to the intended recipient through the use of some network protocol. The right panel in Figure 2 shows the paths from the Sitterson node (node 0) to all the other endpoints on this UNC network. The paths are nothing more than a series of ordered edges that indicate the transmission route from a sender A to a receiver B.

There are several different protocols that control data transfer between communication endpoints, with each protocol fulfilling a particular need. For example, the Transmission Control Protocol (TCP) guarantees reliable and in-order delivery of data from sender to receiver. On the other hand, the lightweight User Datagram Protocol (UDP) does not provide such guarantees. The Internet Control Message Protocol (ICMP) is used primarily to report error messages through the network. The ICMP defines several types of packets, including the *echo request* and *echo reply* messages that are commonly used by the *ping* utility, and the *time-to-live exceeded* and *port unreachable* messages that are critical parts of the traceroute utility (see Section 2.2.1). The protocols can also rely on one another. For example, the well-known Hyper Text Transfer Protocol (HTTP) is built on top of the TCP protocol which is built on top of the Internet Protocol (IP). The Real-time Transport Protocol (or RTP) for delivering audio (such as VoIP) and video over the Internet is built on top of the UDP protocol since for these applications it is not necessary to resend a packet that was dropped.

Many data communications and protocols involve roundtrip communications. This can be asymmetric, i.e., the list of devices from point *A* to point *B* might not simply consist of the reversed list of devices from point *B* to point *A*. Even when network engineers intend to create a network that has symmetric paths (for a variety of reasons that include simplicity), the routing protocols can create temporary asymmetries. Furthermore, even if the routing is symmetric, there are still two paths to consider and the performance on these two paths can be quite different.

Packets arriving at a router or node are queued, awaiting their transmission to the next router according to the packet's protocol as handled by the router. Physically, a queue consists of a block of computer memory that temporarily stores the packets. If the queue (memory) is full when a packet arrives, it is discarded. Otherwise, it waits until it reaches the front of the queue and is forwarded to the next router on the way to its destination. This queuing mechanism is responsible for observed packet losses and, to a large extent, for packet delays.

One of the challenges in analyzing network performance arises from the fact that the topologies can change over time. Network paths are determined by cooperating routers on the network, and the routers use more or less standard routing protocols. Network engineers can also manually impose static routes as desired although such interventions are not common. The routing protocols work by exchanging messages that can result in changes to the paths on the network; these changes can take place within

seconds or minutes depending on the configurations of the routing protocols.

2.2 Data

2.2.1 Ping and Traceroute Data

There are several existing data collection utilities that can be used to collect data on the performance of network connections and remote computers. We discuss here the Ping and Traceroute utilities and their usefulness.

The Ping utility is a tool used by system administrators to check if a remote computer is operating and to determine network connectivity. The source computer sends an ICMP packet to the remote computer's IP address. If the destination computer is up and the network connections are fine, it receives a return ICMP packet. Ping also provides information on the number of hops between the two computers and the amount of time it takes for a packet to make the complete trip. Thus, one can collect data on roundtrip times and delays.

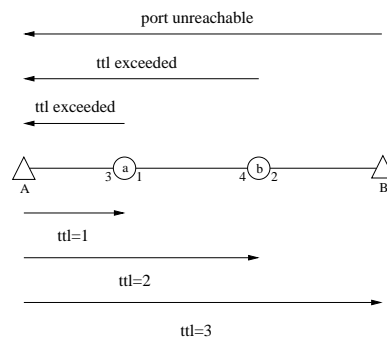


Figure 3: Traceroute Session

The Traceroute utility is another source although it is used mostly to identify network topology (see Section 5). Traceroute sends UDP packets from the source, and it exploits the time-to-live (TTL) field of a packet to determine the route that the packet takes to its destination (Figure 3). Every IP packet has a TTL field that takes values between 0 and 255. When a router receives an IP packet, it decrements this TTL field and forwards the packet to its destination according to its routing table. If, however, the TTL field was already 0, the router sends back an ICMP packet, indicating TTL exceeded, to the source. Traceroute packets are sent at increasing values of TTL, starting with 1, until the destination is actually reached.

The source actually sends the traceroute packets to some invalid port at the destination. When the destination receives a packet destined for an invalid port, an ICMP packet indicating “Port unreachable” is sent back to the source to indicate the error. The source then knows the destination was reached. All the previous packets failed to reach the destination because the TTL was too small and the source received

a *TTL exceeded* message from each of the intervening routers between the source and the destination, in the order in which they appear. Figure 3 illustrates a traceroute session.

Many things can go wrong with traceroute and ping. Some routers are configured to not send any ICMP messages or to not forward them. In addition, traceroute can produce false paths in the presence of per-packet load balancing which sends each successive packet on a different path. Traceroute does not directly identify the routers but only the IP addresses. Routers have many IP addresses, one for each of their interfaces. When multiple paths are collected with traceroute, a given router may appear under different IP addresses in different paths. Thus, the traceroute and ping data can provide incomplete or inaccurate information.

The biggest drawback with these data, however, is that their protocols are different from the applications of interest, and hence the network routers may treat these packets differently. Thus, their performances are not necessarily good surrogates for the applications one is studying such as VoIP.

2.2.2 Injected Probe Data

A direct approach for application-sensitive monitoring is to actively inject “probe” packets that mimic the particular application (VoIP, FTP, etc.) into the network and measure various characteristics of end-to-end performance. The idea of injecting probe packets originated in the MINC project (Caceres et al. 1999). For more information on transmission mechanisms, see Peterson and Davie (2003).

The right panel in Figure 2 shows the topology that was used to collect data on the UNC campus network for VoIP readiness (Lawrence, Michailidis, and Nair, 2006). In this case, the packets were sent from a single source node (Sitterson) to all the other end-points on the periphery of the network. Other schemes are also common such as sending packets from each node on the periphery to every other node (as in the Avaya network example to be discussed later).

The characteristics of interest include loss rates and delays of the probe packets. These can be one-way (node A to B) or round-trip measurements (node A to B and back to A). The measurement of one-way delays is difficult because it involves a synchronization of clocks at the sender and the receiver. This is an important problem, but we will not discuss it here due to space limitations. See Adhikari et al. (2003), Paxson (1998), Moon, Skelly and Towsely (1999), Zhang, Liu and Xia (2002), Jeske and Sampath (2003) and Jeske and Chakravartty (2006) for several algorithms in the literature.

There are many advantages to actively injecting probe packets to study network performance. First, probes can be sent proactively to monitor the network over time and detect developing problems before any adverse effect is experienced. Second, these probes measure QoS metrics on end-to-end performance for the applications of interest. These measures often depend on long-range interactions (in terms of the network topology) on the network. An attempt at direct evaluation of all such interactions would be prohibitively expensive and quite unnecessary since only a specific set of interactions is of

relevance to any given application.




Figure 4: ICMP and UDP Round Trip Times

The key issue here is that the packets mimic the application of interest. In particular, the fields of the IP headers of the injected packets should be indistinguishable from those used by the actual application. In the presence of devices that perform deep packet inspection, the IP payload may also need to be carefully crafted. To illustrate this, consider Figure 4 which shows roundtrip times for 100 UDP packets and 100 ping (ICMP) packets sent during the same period to compare their performance. The solid circles represent ICMP packets and the X's indicate UDP packets, both in milliseconds (ms). All 100 UDP packets made the round trip while only 73 of the ping packets completed it. The lost ICMP packets are indicated as vertical segments at the top of the figure. The UDP round-trip times did not exceed 200 ms while 27 of the ping round trip times were above 400 ms. [In this case, there was a traffic shaper on the path between the source and the destination that occasionally withheld some echo request packets (ICMP) for a multiple of 0.42 second delays.] The lesson to be drawn is that easily obtained measurements such as ping or traceroute are not necessarily adequate surrogates for performance of other types of traffic.

Traceroute data are useful in identifying paths (connected edges over which the packets are transmitted for a source to a destination) and discovering the network topology. The problem of topology discovery is discussed in Section 5. We will discuss the use of traceroute data as auxiliary information for estimating edge-level QoS parameters in Section 3.5.

We will restrict attention in this paper to performance assessment based on injected packets. However, if one had access to network elements, such as routers, one could address these problems by directly sampling packets. Estimation of network flow characteristics based on sampled data has recently been studied in the literature (see Claffy et al. (1999), Duffield (2004), Duffield et al. (2005), and Yang and Michailidis (2007)).

2.2.3 Avaya Network: Data Collection and Summary

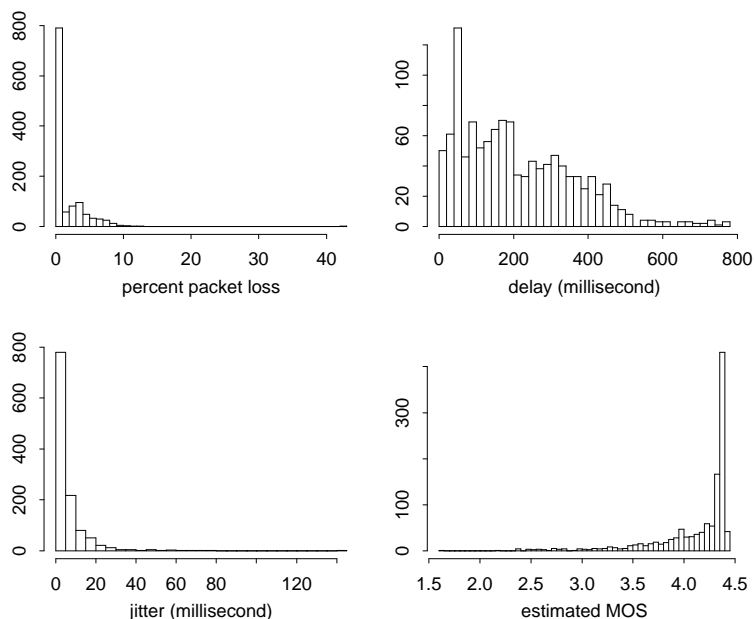


Figure 5: Summary of End-to-End Performance Data

This section describes a study of the network in Figure 1 and summarizes the data. The network had 37 communication endpoints and so there were $N = 37 \times 36 = 1332$ end-to-end pairs. For each pair, we sent an RTP stream of 100 packets from the source node to the destination. These packets were simply bounced back to the sender, leading to roundtrip delays and losses. The performance associated with these packets provided data on end-to-end network delay, jitter and loss. Delay is caused primarily by the queues (congestion) at the routers. One can keep track of the delays for each packet, but in this study, we recorded only the median delay for the 100 packets. Jitter is a measure of variability which is given here by the inter-quartile range of the packet round-trip times. Losses occur when the buffer is full and the packets are dropped. We computed loss as the percentage of packets (of the 100 packets sent) that did not make it back to the source within 5 seconds. (The worst network round trip time observed in the study was a bit less than 1 second). Finally, eMOS (estimated mean opinion score, ITU P.800.1 (2006)) is a common qualitative measure of overall voice telephony performance. It is derived as a function of delay, loss, and jitter, and the values range from 1 to 5 with 5 being best.

The data for the 1332 pairs were collected in 74 rounds of 18 pairs at a time in a random order during an 8 minute period. An endpoint was used in at most one pair in each round. In order to explore the temporal aspects of our problem, we collected data over a period of two weeks. Figure 5 shows the distribution of responses for the 1332 pairs for one time slice and gives an overall picture of the quality of the network for this time period. One can also analyze the data for each end-to-end pair over time (for the entire two weeks) to study temporal variation, time-of-day-effect, etc. Such end-to-end data are very

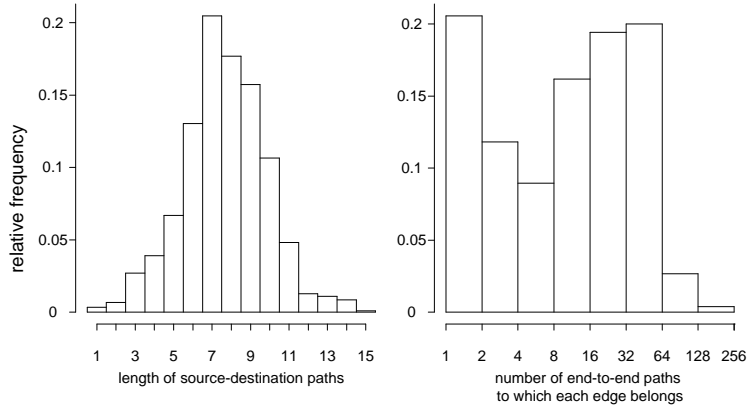


Figure 6: Left panel: Path Length Distribution for the First Routing Matrix; Right Panel: Edge Distribution for the First Routing Matrix

useful for monitoring network performance.

In addition to network performance, we are also interested in various features of the network itself, how they change over time, etc. For example, the left panel of Figure 6 shows the distribution of the path lengths (number of edges) for the 1332 end-point pairs for the first time slice. The lengths varied from 1 to 15 edges with average path length of 8. The right panel shows how the edges were distributed over the paths, i.e., the number of source-destination paths to which each edge belongs. The median of this distribution is 12. About 20.6% of the edges belonged to a single path while 3 edges belonged to over 100 paths, indicating that the performances of these edges are critical for the network.

3 Estimating Edge Parameters from End-to-end Data: Active Network Tomography

This section considers the problem of partitioning the end-to-end measurements to estimate loss rates and delay distributions of the edges. We use a combination of artificial and real situations and data collected on the Avaya network to illustrate the problems as well as to discuss the modeling and analysis results.

3.1 Problem Formulation

For simplicity, we start with the delay estimation problem for the toy network in the left panel of Figure 7. First, consider the case with a single source node 0 that sends probes to the two receiver nodes 2 and 3. Let X_1 , X_2 , and X_3 be the one-way delays associated with the edges as shown in the figure. For the moment, assume the delays are fixed numbers. Let $Y_{(0,2)}$ and $Y_{(0,3)}$ denote the end-to-end delays




Figure 7: Left panel: Two-Layer Binary Tree; Right panel: Four-Layer Symmetric Binary Tree.

associated with the paths $\langle 0, 2 \rangle$ and $\langle 0, 3 \rangle$ respectively. Then, $Y_{\langle 0,2 \rangle} = X_1 + X_2$ and $Y_{\langle 0,3 \rangle} = X_1 + X_3$. Letting \mathbf{y} and \mathbf{x} be the corresponding vectors, we can write $\mathbf{y} = \mathbf{R}\mathbf{x}$ where

$$\mathbf{R} = \begin{pmatrix} 1 & 1 & 0 \\ 1 & 0 & 1 \end{pmatrix}. \quad (1)$$

\mathbf{R} is the routing matrix with $R_{ij} = 1$ if the i -th path contains the j -th edge and zero otherwise.

Now, the delays are not fixed and will vary from probe to probe. We will assume a stochastic model where they are stationary and spatially independent, i.e., the X_j 's are independent. [We will revisit these assumptions later in the section.] Suppose we send M probes from 0 to each receiver node 2 and 3 for a total of $2M$ probes. Let μ_j be the mean delay and $\epsilon_{mj} = X_{mj} - \mu_j$, where X_{mj} is the delay at edge j experienced by the m -th probe. Letting \mathbf{y} be the $2M \times 1$ vector of end-to-end delays, we can write $\mathbf{y} = \mathbf{R}\boldsymbol{\mu} + \boldsymbol{\epsilon}$. Now the routing matrix \mathbf{R} is $2M \times 3$ with each row in equation 1 repeated M times. In our analyses, we have replaced the 100 individual packet delays by a single summary measure, viz., the sample median. The goal then is to estimate the edge-level delays μ_1, μ_2 and μ_3 from the end-to-end delay data. It is clear that we cannot estimate all three edge parameters in this situation. We will return to this estimability problem and related issues.

One can use a different probing scheme which sends probes from each end-point to all other end-points as shown in the left panel of Figure 7. We now have six link-level parameters corresponding to the two directions of the edges. There are also six end-to-end delays $Y_{\langle 0,2 \rangle}, Y_{\langle 0,3 \rangle}, Y_{\langle 2,0 \rangle}, Y_{\langle 2,3 \rangle}, Y_{\langle 3,0 \rangle}, Y_{\langle 3,2 \rangle}$. The routing matrix \mathbf{R} is now given by

$$\mathbf{R} = \begin{pmatrix} 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 1 & 0 & 0 & 1 & 0 \end{pmatrix}. \quad (2)$$

Again, we have a linear model of the form $\mathbf{y} = \mathbf{R}\boldsymbol{\mu} + \boldsymbol{\epsilon}$.

The same formulation works in general for other networks. For the UNC VoIP study, the probes were sent from a single source node (Sitterson) to all the other end-nodes using the topology on the right panel in Figure 2. For the Avaya network in Figure 1, probes were sent from each of the 37 communication end-point to all other 36 end points for a total of 1332 end-to-end pairs. There were 525 edges, so the dimension of the routing matrix was 1332×525 .

In general, we have a linear inverse problem of the form $\mathbf{y} = \mathbf{R}\boldsymbol{\mu} + \boldsymbol{\epsilon}$ where the goal is to estimate the mean link-level delays from the end-to-end data. [See Castro et al. (2004) and references therein for a review on linear inverse problems in network tomography.] If the routing matrix \mathbf{R} is of full rank, this is straightforward and can be solved using least squares, subject to the constraint that the μ 's have to be non-negative. One can also use weighted least-squares that incorporates the variance-covariance structure of the error terms to get more efficient estimators.

In most cases, however, the routing matrix is not of full rank. For example, the rank of \mathbf{R} in equation (1) is 2 while we have 3 edge parameters; in equation (2), the rank is 5 while we have 6 edges. For the network in Figure 1, the rank of \mathbf{R} is 331 while there were 525 edges. Part of the degeneracy arises from a ‘‘chaining’’ phenomenon where some edges are completely confounded with others (where an edge has only one child and the parent and child cannot be separated). If we remove such degeneracies, \mathbf{R} has 454 columns (edges) but its rank is still 331. Thus, only a subspace spanned by the edge parameters is estimable. We discuss alternative probing schemes and the use of auxiliary data to address the estimability problem in the next few sections.

While estimation of mean delays is a linear inverse problem, inference for delay *distributions* is not. Specifically, let F_j be the distribution of X_j in the toy example in the left panel of Figure 7. We want to estimate F_j , $j = 1, \dots, 6$ from the end-to-end data. This is a non-linear, deconvolution problem that is embedded within a graph.

The loss estimation problem can be transformed (approximately) into a linear problem. Consider again the left panel of Figure 7. The X_j 's are now binary with $X_j = 1$ if the packet is not lost on that edge and $X_j = 0$ if it is. The end-to-end data are $Y_{\langle 0,2 \rangle} = X_1 X_2$ and $Y_{\langle 0,3 \rangle} = X_1 X_3$, which are also binary. Let $\alpha_j = E(X_j)$, the probability of a successful transmission through edge j . Then $E(Y_{\langle 0,2 \rangle}) = \alpha_1 \alpha_2$ and $E(Y_{\langle 0,3 \rangle}) = \alpha_1 \alpha_3$. Suppose we send M probes each from node 0 to nodes 2 and 3, and let $Z_{\langle 0,2 \rangle}$ be the proportion of the M probes that successfully reach node 2; $Z_{\langle 0,3 \rangle}$ is similarly defined. Let $Y = \log(Z)$ and $\beta_j = \log(\alpha_j)$. Then, we can write an approximate linear model $\mathbf{y} = \mathbf{R}\boldsymbol{\beta} + \boldsymbol{\epsilon}$ where \mathbf{R} is the same routing matrix as before. [This is approximate because the mean of the error terms is not identically zero but tends to zero as the number of probes $M \rightarrow \infty$.] Estimability issues for the loss rates are thus similar to those for estimating mean delays.

As mentioned in Section 2, there are other QoS measures of interest such as jitter and eMOS, but we will not consider them here.

The estimation methods to be discussed below assume, as is commonly done in the literature,

that the losses and delays at the edges are temporally stationary and spatially independent. Temporal stationarity is reasonable as the probes are sent within an order of seconds, so the parameters will not vary much locally in time. By repeating the experiments over a period of time, we can estimate the QoS parameters as a function of time. A more troublesome assumption is that performance at different edges are independent of each other. This is unlikely to hold as the subnetwork being studied can be part of a bigger network with *cross traffic* that might influence behavior at several nodes simultaneously. There have been attempts to study this aspect through the use of network simulators (see Cáceres, et al. (1999)). However, the nature of the dependence is specific to the network being studied, and it appears difficult to obtain results that are generally applicable.

The rest of this section deals with several important and interesting statistical problems that arise in the context of estimating the loss rates and delays at the edge level. These include questions of estimability, design of probing experiments, inference methods for loss rates and delay distributions, and the use of traceroute data as auxiliary information for estimating edge parameters. We describe theoretical issues, previous research in the literature, and some new ideas that are applied to the probe experiments using the Avaya corporate network.

3.2 Estimability

The probing schemes we have discussed so far sends packets from a source node to a set of receiver nodes by transmitting the packets to one receiver at a time. This is called a *unicast* scheme. As we have seen, this leads to a routing matrix that is not of full rank, so one cannot estimate all the edge parameters. There is an alternative called the *multicast* scheme that gets around this problem by sending the packets simultaneously to a specified set of receiver nodes. Suppose we want to send a packet from node 0 to receiver nodes 2 and 3 simultaneously for the network in the left panel of Figure 7. The source node 0 sends a packet to node 1 where it is duplicated and forwarded to both nodes 2 and 3.

Consider again the loss problem, and suppose we use the multicast scheme to transmit N packets from node 0 to nodes 2 and 3 simultaneously. The resulting sufficient statistics can be expressed as a 4-tuple: $N_{11}, N_{10}, N_{01}, N_{00}$ which are, respectively, the number of transmissions that reached both 2 and 3, 2 but not 3, 3 but not 2, and none. This 4-tuple has a multinomial distribution with parameters $(\pi_1, \pi_2, \pi_3, \pi_4)$ where $\pi_1 = \alpha_1\alpha_2\alpha_3$, $\pi_2 = \alpha_1\alpha_2(1 - \alpha_3)$, $\pi_3 = \alpha_1(1 - \alpha_2)\alpha_3$, and $\pi_4 = 1 - \alpha_1 + \alpha_1(1 - \alpha_2)(1 - \alpha_3)$. It can be seen that one can now estimate all three edge-level parameters from this scheme. The higher-order information from the multicast scheme (in the form of the shared experience of the top edge) is critical for estimability.

Even under this multicast transmission, mean delays or delay distributions are in general not estimable. To see this, consider again the toy example in the left panel of Figure 7. Let the distribution of delay X_j be $\text{Normal}(\mu_j, 1)$, for $j = 1, 2, 3$. Suppose, as before, we send packets from node 0 to receiver nodes 2 and 3 simultaneously. Then, the observed end-to-end delays are bivariate normal; the mean vector has elements $\mu_1 + \mu_2$ and $\mu_1 + \mu_3$ and the variance-covariance matrix has diagonal elements

equal to 2 and off-diagonal elements equal to 1. It is clear that one cannot estimate all the μ_j 's from these bivariate normal data. Similar problems also exist if we send packets from all end-points to each other in this toy example or if we have larger networks.

Chen, Cao, and Bu (2007) showed that, in general, one can estimate all moments except the first of the delay distributions using multicast data. More information or constraints are needed to estimate the first moment. One such case is when the higher moments are a function of the first moment, leading to additional estimating equations. To see this, suppose delay distributions, $F_j(x)$'s, satisfy the property that $Var(X_j) = \mu^2(X_j)$. Then, for the two-layer tree in the left panel of Figure 7, $Cov(Y_{\langle 0,2 \rangle}, Y_{\langle 0,3 \rangle}) = Var(X_1) = \mu^2(X_1)$, giving us an estimating equation for μ_1 , the mean delay of the first edge. Another situation where we can estimate the delay distributions is when the edge delays have point mass at zero, i.e., the delays can be zero with positive probability. Consider again the left panel of Figure 7 and the subset of end-to-end measurements $Y_{\langle 0,2 \rangle} = 0$ and $Y_{\langle 0,3 \rangle} = x$ for some positive value x . This implies that the delay at edges 1 and 2 were both 0, and hence $Y_{\langle 0,3 \rangle} = X_3$. We can use these observations (for the various values of x) to estimate the delay distribution of edge 3. Similarly, the subset of data with $Y_{\langle 0,2 \rangle} = x$ and $Y_{\langle 0,3 \rangle} = 0$ for various values of x can be used to estimate distribution at edge 2. Thus, all three edge-level distributions are estimable.

A major practical problem with the use of multicast schemes is that multicast support is not mandatory under IPv4, so many networks do not have multicast enabled by default (see de Goyan (1998)). There have been proposals in the literature to use back-to-back unicast schemes, where packets are sent within nanoseconds of each other to two or more receivers, to mimic multicast transmissions (Tsang, Coates, and Nowak (2003)).

3.3 Design of Probing Experiments

Injecting probe packets into the network can add a significant amount of traffic, so we have to carefully design probing experiments in terms of how much data to inject, when, for which pairs of end-points, etc. A full multicast scheme sending packets from many end-points to many other end-points in the network can generate more traffic than is desirable or necessary. One may also want to inject traffic in different parts of the network with different intensities over time.

We noted earlier that the higher-order information about the shared edges in multicast schemes was critical for estimating the internal edges. It turns out that second-order information (pairs of receiver nodes at a time or bicast schemes) is sufficient for this purpose. To understand this, consider the right panel of Figure 7 and restrict attention to the three-layer tree with receiver nodes 4, 5, 6, 7. A full multicast involves sending the packets from node 0 to all four receiver nodes simultaneously. This results in a $2^4 - 1 = 15$ -dimensional multinomial experiment. Instead, it suffices to send packets to the pairs $\langle 4, 5 \rangle$, $\langle 6, 7 \rangle$, and $\langle 5, 6 \rangle$, as we will see below. Note that each of these bicast schemes results in only a 4-dimensional multinomial experiment. In general, one can use a combination of unicast and bicast schemes to estimate all the edge-level parameters. This is particularly appealing if we are using back-to-

back unicast schemes to mimic multicast schemes. Back-to-back schemes are most effective for pairs of receiver nodes and are less likely to mimic a multicast scheme as the number of receiver nodes increases.

A general class of flexible probing experiments (called flexicast) that is aimed at addressing the above problem was introduced in Xi et al. (2006) and Lawrence et al. (2006). It consists of a combination of k -cast schemes for different values of k with each k -cast scheme aimed at studying a subnetwork. Each of the k -cast schemes by itself will not necessarily allow us to estimate the edge-level parameters of that subnetwork. The data have to be combined across the various k -cast schemes to estimate the edge-level parameters. This class of experiments has to satisfy some simple conditions in order for all the edge-level parameters to be estimable. Xi et al. (2006) studied this problem for single-source tree topologies and showed that the following conditions are necessary and sufficient for identifiability of loss rates: (a) all receiver nodes are covered; and (b) for each internal node in the tree, there is a k -cast scheme that splits at that internal node. Lawrence et al. (2006) and Lawrence et al. (2007) showed these conditions are also necessary and sufficient provided the delay distribution is discrete or if the higher-order moments are a function of the first moment.

As an example, consider again the three-layer binary tree. Suppose we used an experiment with the two bicast schemes $\langle 4, 5 \rangle$ and $\langle 6, 7 \rangle$. We have covered all receiver nodes, and the first splits at node 2 and the second splits at node 3. However, there is no split at node 1 indicating that not all of the internal nodes are estimable. The following experiment with 3 bicasts, $\langle 4, 5 \rangle$, $\langle 6, 7 \rangle$, and $\langle 5, 6 \rangle$, satisfy the identifiability conditions as the third pair splits at node 1. Note that we can replace the third pair $\langle 5, 6 \rangle$ with another one, such as $\langle 4, 7 \rangle$, which also splits at node 1, indicating that the schemes satisfying the conditions are not unique.

A related question is how to allocate the total probe budget, of say N probes, among the various k -cast schemes. This can be formulated as an optimization problem using criteria in the optimal design literature. See Xi et al. (2006) for discussion in the context of loss rates and single-source topologies. It turns out that the optimal allocations depend on the unknown edge-level parameters.

A more interesting question relates to estimability with multisource topologies as in the Avaya network. For example, how should we supplement the unicast data (all end-to-end pairs) with a minimal number of bicast schemes (or back-to-back schemes) in order to estimate all the internal edge parameters? In other words, given a routing matrix that is degenerate, can we characterize explicitly what probing experiments are needed to resolve the degeneracy? This is part of a more general question on studying the estimability problem with multisource topologies.

3.4 Inference Using Injected Probe Data

We assume that the probing schemes satisfy the identifiability conditions so that all the edges are estimable. By treating the unobservable edge-level data as missing data, one can use the EM algorithm to compute the MLEs in both cases, and this has been done extensively in network tomography applications

(see Coates et al. (2002) and Duffield et al. (2002)). A different approach for estimating loss rates was introduced in Caceres et al. (1999), where the sufficient statistics of the data are first calculated and then a solution to the likelihood function is obtained, by solving a set of polynomial equations. It can be shown that this approach leads to asymptotically ML estimates, but their performance in finite samples can be inferior.

3.4.1 Delay Distributions

To keep things simple, we consider only situations with a single source, but the same ideas apply to multistate situations. Let X_k denote the (unobservable) delay on edge k , p_r be the path from node 0 to receiver node r , and let $Y_r = \sum_{k \in p_r} X_k$, the cumulative delay accumulated along this path. We observe end-to-end delays consisting of Y_r for all the receiver nodes.

Lawrence et al. (2007) examine inference for mean delays $\mu_j = E(X_j)$ under the model where $Var(X_j) \propto \mu_j^\phi$ for some $\phi > 0$. Maximum-likelihood estimation is still intractable even for simple parametric models, so they propose and study the properties of moment-based methods. For example, the covariance terms $Cov(Y_r, Y_s)$ provide additional estimating equations for estimating the edge-level mean delays.

The more general case of estimating delay *distributions* has been studied in network application assuming a discrete distribution. A simple and fast algorithm was developed in Lo Presti et al. (2002), but it is quite inefficient. Liang and Yu (2003) proposed a pseudo-likelihood estimation method with multicast data. This involves using only the one and two-dimensional data and ignoring the higher-order information for computational simplicity. Lawrence et al. (2006) studied maximum likelihood estimation and the behavior of the EM algorithm for general flexicast schemes. Again, the EM algorithm is a reasonable technique for computing the MLEs. However, the complexity of the EM algorithm, in particular computing conditional expectations of the internal edge delays for each bin, is prohibitive for all but fairly small-sized networks. To deal with larger networks, Lawrence et al. (2006) developed a *grafting* method which fits “local” EMs to the subtrees defined by the k -cast schemes and then combines the estimates through a fixed point algorithm. This hybrid algorithm is fast and has reasonable statistical efficiency compared to the full MLE.

See Chen et al.(2007) for delay estimation using Fourier methods, and Tsang et al. (2003) and Shih and Hero (2003) for inference under other models.

3.4.2 Loss Rates

The structure of the EM-algorithm for loss estimation was studied in Xi et al. (2006). It works well for small to moderate-sized topologies but becomes computationally infeasible as the number of edges gets large. We describe next a new, scalable algorithm based on least squares that leads to fast estimation of

loss rates (Michailidis, Nair, and Xi (2006)).

For concreteness, consider the three-layer symmetric binary tree (right panel of Figure 7 with just three layers and receiver nodes 4,5,6, and 7), and suppose we use a full multicast experiment to all the receivers (4, 5, 6, 7). There are 16 possible outcomes

$$(1, 1, 1, 1), (1, 1, 1, 0), (1, 1, 0, 0), \dots, (0, 0, 0, 0).$$

Denote the corresponding number of events for each outcome by $N_{(1,1,1,1)}, N_{(1,1,1,0)}$ and so on. We can ignore the last one as there are only 15 linearly independent results. Consider the one-to-one transformation of these 15 events to the following: $(1, 1, 1, 1), (1, 1, 1, +), (1, 1, +, +), \dots$ where a '+' indicates either a '1' or a '0'. The new outcomes are obtained by replacing all the '0's with '+'s. Let $M_{(i,j,k,l)}$ denote the number of these outcomes. Now, if N denotes the total number of probes for the experiment, we can write $E(M_{i,j,k,l})$ as N times the product of appropriate link-level α 's. For instance, $E(M_{(1,1,1,+)}) = N\gamma_{(1,1,1,+)}$ where $\gamma_{(1,1,1,+)} = \alpha_1\alpha_2\alpha_3\alpha_4\alpha_5\alpha_6$. Similarly, $E(M_{(1,+,+,1)}) = N\gamma_{(1,+,+,1)}$ where $\gamma_{(1,+,+,1)} = \alpha_1\alpha_2\alpha_3\alpha_4\alpha_7$. The expectations for the other M 's can be similarly written as a product of a suitable subset of the α_k 's. This naturally suggests fitting a log-linear model to the estimated probabilities $Y_{(i_1, \dots, i_k)} = \log(M_{(i_1, \dots, i_k)}/N)$.

Formally, suppose we have single source topology with P edges and M receiver nodes and we send probes from the source to all the receiver nodes using a multicast scheme. Then, there are $2^M - 1$ outcomes. Let Y be the $2^M - 1$ column vector containing the logarithms of the estimated probabilities, \mathbf{R} be the $(2^M - 1) \times P$ routing binary matrix, β is a P column vector of regression coefficients with $\beta_j = \log(\alpha_j)$ and ϵ a column vector of "error" terms with $E(\epsilon\epsilon') = \mathbf{V}$. We then have the (approximate) linear model $\mathbf{y} = \mathbf{R}\beta + \epsilon$, similar to the formulation in Section 3.1. Now we are also using the higher-dimensional outcomes, which results in a full-rank routing matrix \mathbf{R} . [Castro et al. (2004) also mention a linear model in terms of conditional probabilities, but the present formulation is more efficient.]

We can then estimate the parameters in the linear model using least-squares methods. The ordinary least squares estimate β is given by

$$\hat{\beta}_O = (\mathbf{R}'\mathbf{R})^{-1}\mathbf{R}'\mathbf{y}.$$

However, the error terms have unequal variances and are correlated, so it is more efficient to use generalized least squares. The form of \mathbf{V} , the variance-covariance matrix of \mathbf{y} , can be obtained in terms of the probabilities $\gamma_{(1,1,1,+)}, \gamma_{(1,+,+,1)}$, etc. For example,

$$Var(Y_{(1,1,1,+)}) = \frac{\gamma_{(1,1,1,+)}(1 - \gamma_{(1,1,1,+)})}{N\gamma_{(1,1,1,+)}^2}$$

and

$$Cov(Y_{(1,1,1,+)}, Y_{(1,+,+,1)}) = \frac{\gamma_{(1,1,1,1)} - \gamma_{(1,1,1,+)}\gamma_{(1,+,+,1)}}{N\gamma_{(1,1,1,+)}\gamma_{(1,+,+,1)}}.$$

A simple, non-iterative generalized least squares estimator can be obtained as

$$\hat{\beta}_G = (\mathbf{R}'\hat{\mathbf{V}}^{-1}\mathbf{R})^{-1}\mathbf{R}'\hat{\mathbf{V}}^{-1}\mathbf{y},$$

where $\hat{\mathbf{V}}$ is obtained using the method of moments estimates of γ . However, this simple plug-in estimate of V can perform poorly in small samples. A more efficient alternative is the iteratively reweighted least squares (IRWLS) estimator

$$\hat{\beta}_I = (\mathbf{R}'\bar{\mathbf{V}}^{-1}\mathbf{R})^{-1}\mathbf{R}'\bar{\mathbf{V}}^{-1}\mathbf{y},$$

where $\bar{\mathbf{V}}$ is based on the estimated values of α from the past iteration. Recall that the γ 's are products of appropriate subsets of the α 's. We found the IRWLS estimators to be numerically very close to the MLEs even in relatively small samples.

LS estimation with multicast experiments can be computationally expensive for large networks as the number of rows in the routing matrix \mathbf{R} increases exponentially with the size of the network. The flexicast experiments discussed earlier are more attractive in this case. If we use a minimal number of bicasts that satisfy the identifiability condition discussed in the last section, the number of rows of the entire routing matrix \mathbf{R} in a flexicast experiment is linear in the size of the receiver set, as opposed to exponential for multicast experiments.

The LS estimation approach extends readily to flexicast schemes and multisource topologies. Specifically, for each k -cast scheme h , write the corresponding log-linear model $\mathbf{y}^h = \mathbf{R}^h\beta + \epsilon^h$. By stacking together the data and the routing matrices for all the k -cast schemes in the flexicast experiment, we get a combined linear model $\mathbf{y} = \mathbf{R}\beta + \epsilon$. Since the different k -cast schemes are independent, the variance-covariance matrix of the error term given by

$$\mathbf{V} = \begin{bmatrix} V^1 & 0 & \dots & \dots & 0 \\ 0 & V^2 & \dots & \dots & 0 \\ \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & \dots & \dots & V^H \end{bmatrix}.$$

Another advantage of these LS-based schemes is that there is an explicit expression for the (asymptotic) variance-covariance matrix of the estimators, leading to easy construction of standard errors and hypothesis tests. The LS algorithms also lend themselves to easy updating as additional data become available.




Figure 8: Relative Efficiency of the LS Estimators vs IRWLS

We now describe some properties of these LS estimators. See Michailidis et al. (2006) for more details. The GLS and IRWLS estimates are consistent, asymptotically normal and fully efficient, i.e., $(\mathbf{R}'\mathbf{V}^{-1}\mathbf{R})^{-1} = \mathcal{I}^{-1}(\beta)$, where $\mathcal{I}^{-1}(\beta)$ denotes the inverse of the Fisher information matrix of β . However, in small samples, the GLS estimators do not perform as well as the IRWLS estimators.

Figure 8 shows the asymptotic relative efficiencies of the IRWLS estimators compared to the LS estimators for selected edges for a four-layer tree (see right panel of Figure 7). The y-axis in Figure 8 are the relative efficiencies (ratio of asymptotic variances of the LS estimators and the IRWLS estimators) for four edges (1, 3, 7, and 15) as a function of the probability of successful transmission which varies from 0.75 to 0.99. The computations were done under a multicast scheme and assuming the probabilities for all the edges are the same. We see that the LS estimators can be quite inefficient, so the use of iterative schemes can have big pay off.

One issue that we have not discussed is the large number of parameters (loss or delay) to be estimated in a typical network and that these parameters vary over time. So some type of regularization on shrinkage is warranted. This should take into account the network topology and other relevant information. Regularization has been used in network tomography although in a different context. Zhang et al. (2003) use a *gravity model* in the origin-destination estimation problem to get around the estimability issues. See Liang, Taft, and Yu (2006) for a variation of this approach.

3.5 Combining RTP Probing with Traceroute Data

VoIP and other real-time applications are sent using RTP, so an important need is to understand how RTP streams are handled by the network. For this purpose we want the probes to be RTP streams, or to be able to predict the behavior of RTP streams. Moreover, in corporate or large networks the estimability problems are often severe, and multicast probing is generally not available. Thus, we consider an alternative method for estimating the edge parameters based on combining unicast end-to-end RTP tests with additional data arising from traceroute tests. The ideas are also applicable to combining other types of testing data. We restrict attention to mean delays in this section and focus on roundtrip times. Similar discussion holds for one-way delays and loss data.

The idea is that on the one hand RTP data are available for end-to-end pairs but not directly for all edges. On the other hand, traceroute data are available for all edges, but, for reasons discussed in Section 2.2.2, they might not reliably mimic the behavior of RTP traffic at all times on all edges. We discuss various ways to combine the two sources of information to estimate the edge parameters relevant for VoIP.

To be concrete, consider again the simple two-layer network in the left panel of Figure 7. Recall that the traceroute involves a sequence of probes with roundtrip times to successive routers along each path. So, for the path $\langle 0, 2 \rangle$, we get the traceroute data $Z_1 = X_{1m}^{tr} + X_{6m}^{tr}$ and $Z_2 = X_{1n}^{tr} + X_{2n}^{tr} + X_{4n}^{tr} + X_{6n}^{tr}$ for packets m and n . For path $\langle 0, 3 \rangle$, we get $Z_3 = X_{1p}^{tr} + X_{6p}^{tr}$ and $Z_4 = X_{1r}^{tr} + X_{3r}^{tr} + X_{5r}^{tr} + X_{6r}^{tr}$

for packets p and r . A similar set of data is observed for the other end-to-end pairs $\langle 2, 0 \rangle$, $\langle 2, 3 \rangle$, $\langle 3, 0 \rangle$ and $\langle 3, 2 \rangle$. Letting $E(X_j^{tr}) = \mu_j^{tr}$, we can write a linear model for the traceroute delay data as before as $\mathbf{z} = \mathbf{R}^{tr} \mathbf{x} + \delta$. The routing matrix is now given by

$$\mathbf{R}^{tr} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 1 \\ 1 & 1 & 0 & 1 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 1 & 0 & 1 & 1 \\ 0 & 1 & 0 & 1 & 0 & 0 \\ 1 & 1 & 0 & 1 & 0 & 1 \\ 0 & 1 & 0 & 1 & 0 & 0 \\ 0 & 1 & 1 & 1 & 1 & 0 \\ 0 & 0 & 1 & 0 & 1 & 0 \\ 1 & 0 & 1 & 0 & 1 & 1 \\ 0 & 0 & 1 & 0 & 1 & 0 \\ 0 & 1 & 1 & 1 & 1 & 0 \end{pmatrix}. \quad (3)$$

This matrix has full rank ($= 6$), so we can estimate all six mean delay parameters for traceroute performance, μ_j^{tr} s, directly. The parameters must be estimated under the non-negativity constraint that $\mu_j^{tr} \geq 0$. The constrained least squares algorithm (NNLS) is an option but it can be time consuming with large networks. The pool-adjacent-violators algorithm (Barlow et al. 1972) is a faster alternative; we will denote it as $\hat{\mu}_{\text{pava}}^{tr}$. This method pools adjacent points that violate the monotonicity constraint, averages them, and iterates until the sequence is monotonic. This provides an estimate for each edge in each traceroute path. If an edge is included in more than one path, we take as the edge estimate the median of these values. In a sense, $\hat{\mu}_{\text{pava}}^{tr}$ amounts to taking care of the non-negativity constraint on a path-by-path basis rather than taking care of it globally as does NNLS. We have compared the two methods on data collected on the Avaya network (described in Section 2.2.3) and found them to perform similarly for the most part. For computational reasons, we will focus on $\hat{\mu}_{\text{pava}}^{tr}$ in the rest of this section.

Our goal is estimation of RTP delays. Below we investigate the possibility of using $\hat{\mu}_{\text{pava}}^{tr}$ as an estimate for μ^{rtp} . We also assess the relationship between the traceroute and RTP data in order to combine them. We have roundtrip delay times for all end-to-end pairs from both data sources. Figure 9 plots the differences for the 1332 cases with end-to-end pairs grouped according to the number of edges between the source and destination. Most of the differences are near zero with 63% within 2 ms, 25% between 2 and 10 ms, and 12% > 10 ms. As expected from engineering considerations, traceroute is equal to or, in a small but not negligible set of cases, larger than the RTP measurement, but rarely smaller. The differences are substantial enough so that we cannot simply take the traceroute results as valid surrogates for RTP delays. We now discuss several heuristic methods that use the traceroute data as auxiliary information to estimate RTP performance.

A straightforward approach is to rescale the estimated traceroute delays $\hat{\mu}_{\text{pava}}^{tr}$ so that the estimated delay summed over all the edges in the end-to-end path matches the RTP delay measured on the same end-to-end path. When an edge is a member of more than one end-to-end path, we first scale on each




Figure 9: Comparing End-to-End Traceroute and RTP.

path and then take the median over all the paths that contain this edge. We call this estimator $\hat{\mu}_{\text{cal}}^{rtp}$ where “cal” stands for “calibration”.

An alternative is a penalization framework that uses the intermediate values $\hat{\mu}_{\text{pava}}^{tr}$ to estimate the μ^{rtp} . This can be obtained by minimizing

$$\|\mathbf{y} - \mathbf{R}\mu^{rtp}\|^2 + \lambda\|\mu^{rtp} - \hat{\mu}_{\text{pava}}^{tr}\|^2 \quad (4)$$

where λ is a weighting parameter that is specified. The resulting estimator, $\hat{\mu}_{\text{pen}}^{rtp}$ can be obtained using least squares. We note, in passing, that regularization methods have been used in network tomography in other contexts (

One can also directly minimize

$$\|\mathbf{y} - \mathbf{R}^{rtp}\mu^{rtp}\|^2 + \lambda\|\mathbf{z} - \mathbf{R}^{tr}\mu^{rtp}\|^2. \quad (5)$$

which is just WLS with λ standing for $Var(Y)/Var(Z)$. Since traceroute delays can be substantially larger than RTP delays, it makes sense to take λ small, or even to consider the limit as $\lambda \rightarrow 0$. In this limiting case, we are estimating all estimable linear functions from the RTP data alone, while the traceroute data is used solely to disambiguate the degeneracies. This approach goes some way to allowing for the possibility that the traceroute observations are not estimating the same parameters as the RTP data.

Figure 10 compares the three edge-level estimators: $\hat{\mu}_{\text{pava}}^{\text{tr}}$, $\hat{\mu}_{\text{cal}}^{\text{rtp}}$, and $\hat{\mu}_{\text{pen}}^{\text{rtp}}$ from equation (4) with $\lambda = .001$. Other small values of λ would produce nearly identical fits. The scatter diagrams show the values of $\hat{\mu}_{\text{pava}}^{\text{tr}} - \hat{\mu}_{\text{cal}}^{\text{rtp}}$ against $\hat{\mu}_{\text{cal}}^{\text{rtp}}$ and $\hat{\mu}_{\text{pen}}^{\text{rtp}} - \hat{\mu}_{\text{cal}}^{\text{rtp}}$ against $\hat{\mu}_{\text{cal}}^{\text{rtp}}$ for each of the 525 edges. The first two estimates, in the left panel, are in fairly close agreement: these estimates both have 114 coefficients (22%) which are exactly zero; they differ by at most 40 ms; and 99% of them are within 10 ms of each other. These estimators are quite different from the penalized estimator $\hat{\mu}_{\text{pen}}^{\text{rtp}}$ as seen in the right panel.

Figure 11 provides a different comparison of the same three estimators. The scatter diagrams display residuals against fitted values, and the mean sum of squared residuals are 207, 217, and 76 respectively. Using the second penalized approach, equation (5), also leads to a mean sum of squared residuals of 76. The penalized estimators are clearly superior, and $\hat{\mu}_{\text{pen}}^{\text{rtp}}$ from equation (4) involves much less computation than using equation (5).




Figure 10: Comparing $\hat{\mu}_{\text{cal}}^{\text{rtp}}$ to $\hat{\mu}_{\text{pava}}^{\text{tr}}$ and $\hat{\mu}_{\text{pen}}^{\text{rtp}}$

An important practical problem not addressed here is computing the estimators in real time in a distributed network environment. The penalized estimation method seems to provide a good balance of accuracy and computational ease, but further research is needed.

A more general formulation for combining the two sources is a calibration model. Recall that in the calibration problem, we have *inexpensive* but less reliable measurements from one source and *expensive*, reliable, but limited measurements from another. In our case, the traceroute is the less reliable source and the RTP test is the more reliable one. We have end-to-end delay measurements for the 1332 pairs from both sources, but only traceroute data for the edges. So we can develop a calibration model that relates the two data sources from the end-to-end measurements and use the model to calibrate the

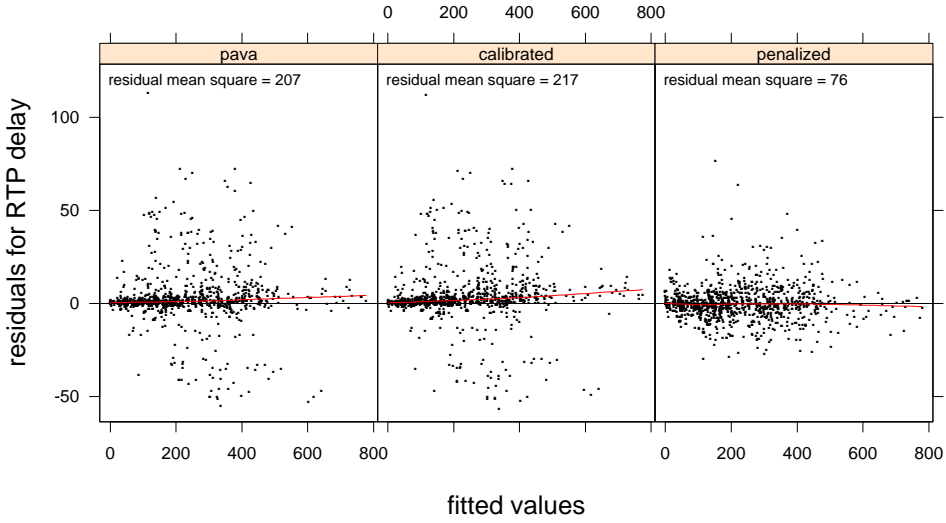


Figure 11: Comparing $\hat{\mu}_{\text{cal}}^{\text{rtp}}$ to $\hat{\mu}_{\text{pava}}^{\text{tr}}$ and $\hat{\mu}_{\text{pen}}^{\text{rtp}}$

edge-level traceroute data. Using the existing notation, we can write

$$Y_i^{\text{tr}} = f(Y_i^{\text{rtp}}, \mathbf{R}|p_i)$$

for $i = 1, \dots, n$. The functional form of $f(\cdot)$ captures the relationship between the end-to-end RTP data Y_i^{rtp} and the end-to-end traceroute data. This relationship may depend on the path p_i ; for instance through the number of edges in the path. There are many ways one could specify and fit $f(\cdot)$. Once this is done and the model is fit, one can combine the information from traceroute and RTP data to estimate the edge parameters.

4 Monitoring Network Performance: Detecting and Diagnosing Changes

This section considers methods for monitoring network performance, detecting degradation in performance levels, and diagnosing where the problems occur. Given the high-level of quality, network monitoring often amounts to filtering through large amounts of irrelevant data to determine unusual behavior that is important.

4.1 Monitoring Techniques

Our primary interest is in end-to-end performance, so for the purposes of monitoring, we can focus attention on path characteristics. There is an extensive literature on process monitoring techniques and

change-point detection methods which can be used to monitor end-to-end performance (see, for example, Basseville and Nikiforov (1993) and Stoumbos et al. (2000)). Some preliminary results for monitoring loss rates using exponentially-weighted moving average (EWMA) techniques can be found in Xi et al. (2006).

A challenge in implementing these techniques to network monitoring is that there is usually a large number of paths to monitor (for example, 1332 end-to-end pairs in the moderate-sized Avaya network). Monitoring a large number will result in a high rate of false alarms; if we control the overall false alarm rate, the power of detection will be small. It is important to keep the number of paths to be monitored to a reasonable number (say 20-30) using engineering knowledge about the network topology and other critical business information. There are also key questions about the parameters to monitor, such as loss rates, mean delays, probability of large delays, or some overall measure such as eMOS.

In cases where we can estimate the edge-level parameters (for example, using one of the techniques discussed in Section 3), we can also monitor these parameters directly. A statistical comparison of the relative efficiencies of monitoring techniques based on end-to-end paths versus edge parameters is being done in Yang (2006). However, this involves solving the inverse problems at each stage which can be computationally expensive.

While detecting and diagnosing changes in performance, it is important to distinguish between changes that may be statistically significant from changes that are practically important for the network and its applications. For example, there can be a substantial change in the mean delay of one edge, but if this edge is part of a short path or if the other edges in the path have very small delays, the overall path-level performance can still be quite adequate.

Consider Figure 12 which shows a simple topology with three endpoints (A, B, C) involved in end-to-end communications that are handled by four routers (a, b, c, d). Suppose edge-level estimation resulted in the estimates (in milliseconds) next to each edge in the figure. Assume that the threshold beyond which end-to-end communication is inadequate is 200 ms. The edge c-d has the largest latency, 180 ms, and thus might stand out from an edge-level estimation or monitoring perspective. But it is used in only the B-C communication, for which it is adequate. So the c-d edge does not constitute a problem from this perspective. On the other hand, the A-B communication is inadequate since it has an accumulated latency of 210 ms. It is the serial utilization of the a-b-c edges that generates the problem. For example, if the delay on a-b had increased from 80 to 100 ms, users might experience a greater impact than if c-d had increased from 120 to 180 ms.

Thus, monitoring and diagnosing a network's performance must incorporate fundamentally the design and desired behavior of the network and use these factors as a basis for any alarms, alerts, or diagnostic statements. For the example of Figure 12, it could be argued that the network as designated is not capable of meeting the requirement for this application. To understand what are possible root causes for end-to-end problems, we must also incorporate the paths into the analysis directly.




Figure 12: Root Cause Analysis

4.2 Diagnosing Changes: Root-Cause Analysis

With the above example in mind, we propose the following approach for root-cause analysis. Let Y_i be an indicator of success and failure for the i -th test, where $Y_i = 1$ means the test fails and $Y_i = 0$ corresponds to the test being successful. This indicator could represent, for example, whether or not the delay for the i -th test is adequate relative to its own threshold. As in Section 3.1, let \mathbf{R} be the relevant routing matrix where $R_{ij} = 1$ if the j -th edge is involved in the i -th test and zero otherwise, for $j = 1, \dots, K$ and K is the total number of edges. (The approach can also be generalized to incorporate not just topological factors such as edges, but also other factors such as software version or codec, etc.)

Let π_i be the probability the i -th test fails. For the j -th edge, let p_j be the probability that performance on this edge “fails”; that is, the j -th edge performs poorly enough so that the whole end-to-end test fails. Furthermore, a test has some probability of failing that is not necessarily associated with any specific edge, so let p denote this background probability of failure. That is, p is the probability the test fails although none of the edges directly “causes” the failure. Then $Y_i \sim \text{Bernoulli}(\pi_i)$. For the end-to-end test to be successful it must be successful for each traversed edge and also for the background, so

$$(1 - \pi_i) = (1 - p) \times \prod_{\{j: R_{ij}=1\}} (1 - p_j).$$

The estimation of the p_j 's is difficult mostly because of the large number of boundary conditions, but a straightforward approach appears to work well when there are few failures. We define the edge culpation ratio

$$I_j = \frac{\text{cardinality}\{i : y_i = 1 \text{ and } R_{ij} = 1\}}{\text{cardinality}\{i : y_i = 0 \text{ and } R_{ij} = 1\}}.$$

That is, I_j is the odds ratio of the number of tests crossing edge j that failed to the number of tests crossing this edge that succeeded.

We call this approach “culpation”, and the inculpation set is the set of edges for which the culpation ratio I_j is above some specified threshold. Such edges are inculpated as those that seem to contribute to a high probability of failure for tests traversing them. The inculpation set is simple to obtain and requires processing that can be distributed easily across the network. Indeed, the edge memberships in the routing matrix \mathbf{R} can be determined and stored independently of the test results; distributing the

culpation estimation process amounts to partitioning \mathbf{R} into subsets of columns (i.e., edges) that are handled separately. Moreover, the culpation approach can also be applied to interactions among edges by expanding the routing matrix \mathbf{R} and including products between the columns, as with analysis of variance. One needs to be judicious, however, and only include interactions among edges that make some sense in the network, as otherwise the problem could grow unmanageable quickly.

We now illustrate how the inculpation method works using the estimated MOS statistic (see Sect. 2.2.3) as the end-to-end test, thus incorporating all of delay, loss, and jitter. We take the threshold for test failure to be 3.5 or lower, where 4.0 or higher is interpreted roughly as “toll quality” voice transmission. Figure 1 shows the inculpation set corresponding to the edges $\{j : I_j > 0.8\}$ on a gray scaled as dark and the excluded edges as light. There is only one dark edge, and it is involved in some of the test traffic in and out of the Dubai site, but not to all destinations from Dubai. The threshold 0.8 was selected as a value large enough so that the set contained only one edge. Decreasing the threshold results in including more edges. By varying the threshold and examining how the inculpation set changes, we can get an idea of edges that seem to have the most problematic behavior. In this case, lowering the threshold adds further edges near Dubai.

4.3 Visualization Tools for Monitoring and Diagnosis

The formal techniques discussed so far are insufficient in understanding and diagnosing problems in large, complex networks. We have developed and used visualization tools in an iterative manner to supplement the formal techniques (Adhikari et al., 2006; Adhikari et al., 2007). A static medium such as a paper obviously constraints our ability to illustrate the features of a visualization tool.

The first challenge is to represent the topology and layout of a large network as a graph. This problem is discussed in more detail in the next section. Given such a graph, we have found animation to be useful in studying changes in performance over time. For the topology in Figure 1, one can color edges that exceed performance thresholds in (say) red and visualize the changes over time. Such animations will reveal coincidences over time and space that would otherwise be difficult to detect. One difficulty with this is that the network topology can change over time with nodes and edges being added and deleted, which in turn affects connectivity patterns and other characteristics that might be shown on the network graph.

One useful feature is to display additional information through mouse-over features. This could include static information that is too extensive to display all the time in the plot. Because the topology graph can get very cluttered, we have found it better to reserve a separate panel in the plotting region to display such information than to display it directly adjacent to the item that is moused-over. For nodes, information could be items such as IP addresses for a router, or its name, product type, etc. For edges, the additional information could be a statistic measuring recent performance on that edge, such as delay.

An analysis of edges can hide information about end-to-end paths by visually emphasizing edges

rather than paths. So it is important to reintegrate end-to-end path information into the visualization. A feature we have developed for dealing with this is to click on two endpoints, after which the display highlights (e.g., colors differently) the path (or paths) taken by traffic between these two endpoints. The associated numerical results and graph coloring can also be restricted to calculations using only data from the end-to-end paths between the two endpoints.




Figure 13: Pop-up Time Series for All Paths Touching the Edge

Another feature is a pop-up time series plot that helps with all three aspects of drill-down, temporal behavior, and paths. Clicking on an edge creates a new plot, such as the one in Figure 13, which shows end-to-end test results for delay for a period of time that traversed this edge, regardless of their end-to-end paths. In Figure 13, time is shown on the x -axis with the most recent points on the right, so animating this over time gives the appearance of the points scrolling to the left. The label for the y -axis is shown on the right, which is unconventional but useful here since we typically are most interested in reading off numerical values for the most recent points. Furthermore, there are check boxes on the bottom that indicate all the end-to-end paths that pertain to these data, and by checking each box, or not, we include or exclude in the plot the data points from that path. Sometimes a “banding” feature is visible in the plot, whereby the points fall in several separated horizontal bands. By checking and unchecking the boxes interesting behavior can often be associated with specific end-to-end paths. Color can also be used to identify characteristics of the points, such as whether or not they contributed to triggering an alarm for the edge.

Clearly, there are many other possibilities for visualizing network performance and network data. Developing and implementing visualization tools that are useful to the network engineer or analyst requires careful consideration of many aspects, including user preferences and interactions. Some innovative visual displays for packet header data are given in Solka et al. (2000).

5 Analysis of Network Topology

This section briefly describes several research issues associated with the analysis of network topology.

5.1 Topology Discovery

As has been evident throughout the paper, knowledge of the network topology is essential for assessing and monitoring its performance. However, this information is often difficult to obtain. Furthermore, the topology can change frequently as network devices go up or come down.

There are standard engineering methods to get topological information provided we have the necessary access. The most straightforward is to obtain the routing table of all the routers of interest. This information would give the connectivity graph of the network under consideration and can be obtained through using a certain protocol, SNMP. However, this is a laborious task because of policies that restrict such router queries for security reasons. In addition, SNMP access usually receives low priority and hence discovery based on this mechanism can take a long time (e.g. hours) for a large network. Further, if routing protocols are active in the course of the discovery process, it can produce an inconsistent set of routing tables from which it would not be possible to recover some of the network paths.

Another approach is to use the IP Record Route mechanism that consists of setting a single bit in the IP header. Routers that process packets with the IP Record Route bit set will record their IP address in the IP header before sending the packet on its way to the destination. The IP header can record up to 9 IP addresses, including the source IP address. However, there are several issues that make IP Record Route difficult to exploit in practice. First, IP Record Route fails when there are more than 8 intervening routers, which is common. Second, many devices routinely drop packets that have such special processing bits set. Third and most importantly, many devices handle the IP Record Route packets in an entirely different way than regular packets, often sending such packets on a different interface than regular packets would have been sent. So the recorded route may not be the one actually used.

The approach we used to discover the Avaya network topology in Figure 1 was based on traceroute data. We used traceroute data sent from all the end-to-end pairs and reconstructed the topology by merging the underlying spanning trees. Given the problems noted with traceroute data, this is a “best-guess” effort at constructing the topology. See also Achlioptas et al. (2005) and Dallasta et al. (2006) for biases in using traceroute data.

There have also been methods proposed in the literature for discovering a single-source topology using multicast (or back-to-back unicast) probe data. Given a source and a set of destinations, the problem is to identify a tree topology that best fits the data from among all topologies. One can define a similarity measure based on loss rates or delay times (nodes that share a longer common path should be more correlated than those with small paths). One can then use a clustering algorithm to group the nodes together and determine a topology. Duffield et al. (2002) showed that this strategy can completely identify a binary tree topology. An alternative approach based on a random search strategy for locating an optimal tree topology is discussed in Coates et al. (2002). This paper also proposed a different probing method (called *sandwich probing*) that induces higher correlation and also gets around the clock-synchronization problem. An alternative methodology based on measurements from end-to-end delay covariances is presented in Duffield and Lo Presti (2004). Finally, a discussion of how the topology discovery problem can be cast as a maximum likelihood estimation one and some additional references can be found in Castro et al. (2004)

The topology discovery problem as formulated above is computationally rather difficult, and there is still no satisfactory solution. It has a similar flavor to that of identifying phylogenies in genetics, for which various clustering algorithms can be useful, although the phylogenetic problem has more structure due to evolutionary considerations.

5.2 Representing, Visualizing, and Summarizing Networks

Representing a large and complex network as a graph is a challenging problem. The graph layout needs to combine some aspects of logical and physical reality. Achieving this by manually moving points around can be very time consuming and tedious. Automatic layout procedures have received much research attention in the last few years (see Michailidis (2006) for a review).

The automatic layout problem is defined as follows: given a set of nodes connected by a set of edges, identify the positions of the nodes in some space and calculate the curves that connect them. Most graph drawing techniques use straight lines to connect the nodes and Euclidean space, although other choices such as lattices or hyperbolic space have proved useful in some application areas. The majority of scalable graph drawing algorithms use either an embedding model or an adjacency model (Michailidis (2006)). In the former approach, path length distances are defined between the nodes, which are subsequently approximated by Euclidean distances derived from low (usually 2 or 3) dimensional configurations. Multidimensional scaling and its variations (Buja and Swayne (2002)) are examples of this approach. In the adjacency model, the emphasis is to place close together in Euclidean space nodes that are connected together and a popular algorithm turns out to use an eigenvalue decomposition of the Laplacian matrix of the underlying graph.

While automatic layout algorithms are necessary to get started, our experience is that there is generally detailed engineering, geographic, or network information that can help to modify the layout and make it more interpretable for the user. We have found it beneficial to combine some automatic

layout to get started with capabilities to manually modify the results.

A related topic is that the network can be so large or dense that all of it cannot be displayed conveniently and informatively on the computer screen. Thinking of the virtual display as a large canvas, features are needed to enable the user to navigate as desired across different regions of the canvas, and to zoom in or out. A companion topic involves compression techniques to automatically reduce the size of the graph while at the same time enhancing semantically relevant information, such as the presence of highly connected nodes (hubs) and clusters, or the preservation of the shape of the node degree distribution. Several compression schemes that utilize node degree or shortest path importance, or node similarity measures are discussed in Adler and Mitzenmacher (2001) and Gilbert and Levchenko (2004).

As we have noted, changes in the network topology over time add to the difficulties in representing and visualizing the network. Technical challenges include finding an appropriate data representation for dynamically evolving graphs, tracking those changes over time, updating its structure, and visualizing its evolution. An overview of some relevant issues are discussed in Cortes, Pregibon, and Volinsky (2003) and Eppstein, Galil, and Italiano (1999).

6 Concluding Remarks

We have provided a review of several interesting statistical issues that arise in the context of assessing and monitoring network performance as well as in characterizing the properties of networks. While there has been considerable work in this area, mostly in the network community, there are still many interesting and challenging statistical problems. However, the research issues change rapidly with advances in technology, so it is important for statisticians to identify and collaborate with network engineers who are closely tied to the technology and problems.

Acknowledgements

The authors would like to thank two anonymous referees for useful suggestions. The research of Michailidis and Nair was supported in part by NSF-DMS grant 0500535.

References

- [1] Achlioptas, D., Clauset, A., Kempe, D. and Moore, C. (2005), "On the bias of traceroute sampling, or: Why almost every network looks like it has a power law," *Proceedings of the 2005 Symposium on the Theory of Computation (STOC)*, Baltimore, MD, 694-703

- [2] Adhikari, A., Denby, L., Mallows, C. and Meloche, J. (2003), "Measuring Network One-way Transit Time," *Avaya Labs Research Technical Report ALR-2003-051*.
- [3] Adhikari, A., Bianco, S. V., Denby, L., Mallows, C. L., Meloche, J., Rao, B., Sullivan, S. M., and Vardi, Y. (2006), "Distributed monitoring and analysis system for network traffic," United States Patent, 7,031,264.
- [4] Adhikari, A., Denby, L., Landwehr, J. M., and Meloche, J. (2007), "Using data network metrics, graphics, and topology to explore network characteristics," in *Statistical Inverse Problems*, R. Liu. W. Strawderman, C. H. Zhang (editors), IMS Lecture Note Series.
- [5] Adler, M and Mitzenmacher, M. (2001), "Towards compressing Web graphs," *Proceedings of the IEEE Data Compression Conference*, Snowbird, UT, 203-212
- [6] Barlow, R. E., Bartholomew, D. J., Bremner, J. M., and Brunk, H. D. (1972), *Statistical Inference under Order Restrictions*, New York: Wiley, p. 13.
- [7] Basseville, M. and Nikiforov, (1993) *Detection of Abrupt Changes: Theory and Applications*, Prentice Hall.
- [8] Bearden, M., Denby, L., Karacali, B., Meloche, J. and Stott, D. T. (2002a), "Assessing Network Readiness for IP Telephony," *Proc. of the 2002 IEEE Intl. Conf. on Communications (ICC-02)*.
- [9] Bearden, M., Denby, L., Karacali, B., Meloche, J. and Stott, D. T. (2002b), "Experiences with Evaluating Network QoS for IP Telephony," *Proc. of the Tenth Intl. Workshop on Quality of Service, (IWQoS 2002)*, 259-268.
- [10] Bestavros, A., Byers, J. and Harfoush, K. (2002), "Inference and labeling of metric-induced network topologies," *Proceedings IEEE Infocom*, New York, NY, 628-637
- [11] Buja, A. and Swayne, D. (2002), "Visualization methodology for multidimensional scaling," *Journal of Classification*, 19, 7-43
- [12] Cáceres, R., Duffield, N.G., Horowitz, J. and Towsley, D. (1999), "Multicast Based Inference of Network Internal Loss Characteristics," *IEEE Transactions on Information Theory*, 45, 2462-2480
- [13] Castro, R., Coates, M, Liang, G., Nowak, R. and Yu, B. (2004), "Network Tomography: Recent Developments," *Statistical Science*, 19, 499-517
- [14] Castro, R., Tsang, Y. and Nowak, R. (2004), "Likelihood based hierarchical clustering," *IEEE Transactions on Signal Processing*, 52, 2308-2321
- [15] Chen A., Cao, J. and Bu, T. (2007), "Network tomography: Identifiability and Fourier domain estimation," *IEEE Infocom* (to appear).
- [16] Claffy, K.C., Polyzos, G.C. and Braun, H.W. (1993), "Application of sampling methodologies to network traffic characterization," *Proceedings ACM SIGCOMM*, 13-17

- [17] Coates, M., Castro, R., Nowak, R., Gadhiok, M., King R. and Tsang, Y. (2002), "Maximum likelihood topology identification from edge based unicast measurements," *Proceedings ACM Sigmetrics*, Marina del Rey, CA, 11-20
- [18] Cortes, C., Pregibon, D. and Volinsky, C. (2003), "Computational methods for dynamic graphs," *Journal of Computational and Graphical Statistics*, 12, 950-970
- [19] DallAsta, L., Alvarez-Hamelin, I., Barrat, A., Vazquez, A. and Vespignani, A. (2006), "Exploring networks with traceroute-like probes: theory and simulations," *Theoretical Computer Science*, 355, 6-24
- [20] Duffield, N.G., Horowitz, J., Lo Presti, F. and Towsley, D. (2002), "Multicast topology inference from measured end-to-end loss," *IEEE Transactions on Information Theory*, 48, 26-45
- [21] Duffield, N.G. (2004), Sampling for passive Internet measurement: A Review, *Statistical Science*, 19, 472-498
- [22] Duffield, N.G. and Lo Presti, F. (2004), Network Tomography from Measured End-to-End Delay Covariance, *IEEE/ACM Transactions on Networking*, 12, 978-992
- [23] Duffield, N.G., Lund, C. and Thorup, M. (2005), "Estimating flow distributions from sampled flow statistics," *IEEE/ACM Transactions on Networking*, 13, 325-336
- [24] de Goyan, Juan-Mariano (1998), <http://tldp.org/HOWTO/Multicast-HOWTO.html>.
- [25] Eppstein, D., Galil, Z. and Italiano, G. F. (1999), "Dynamic graph algorithms," in *Algorithms and Theoretical Computing Handbook*, Atallah (ed.), CRC Press, Boca Raton
- [26] Gilbert, A. C. and Levchenko, K. (2004), "Compressing network graphs," *Proceedings of the KDD Conference*, Seattle, WA
- [27] ITU P.800.1 (2006), "Mean opinion score (MOS) terminology," <http://www.itu.int/rec/T-REC-P.800.1-200607-I/en>
- [28] Jeske, D. R. and Sampath, A. (2003), "Estimation of Clock Offset Using Bootstrap Bias-Correction Techniques," *Technometrics*, 45, 256-26
- [29] Jeske, D. R. and Chakravartty, A. (2006), "Effectiveness of bootstrap correction in the context of clock offset estimators," *Technometrics*, 48, 530-8.
- [30] Karacali, B. and Denby, L. and Meloche, J. (2004), "Scalable Network Assessment for IP Telephony," *Proc. of the 2004 IEEE Intl. Conf. on Communications (ICC-04)*.
- [31] Lawrence, E., Michailidis, G., Nair, V. N. (2006), "Network Delay Tomography Using Felxicast Experiments," *Journal of the Royal Statistical Society*, B, pp. 785-813.
- [32] Lawrence, E., Michailidis, G., Nair, V. N. (2007) "Statistical Inverse Problems in Active Network Tomography," in *Statistical Inverse Problems*, R. Liu., W. Strawderman, C. H. Zhang (editors), IMS Lecture Note Series.

- [33] Lawrence, E., Michailidis, G., Nair, V. N. and Xi, B. W. (2006), "Network Tomography: A Review and Recent Developments" , in *Frontiers in Statistics*, Fan and Koul (eds.), Imperial College Press, London, pp. 345-366
- [34] Lo Presti, F., Duffield, N. G., Horowitz, J., Towsley, D. (2002), "Multicast-based inference of network-internal delay distributions," *IEEE/ACM Transaction Networking*, Volume 10, Number 6, pp 761-775.
- [35] Liang, G. and Yu, B. (2003), "Maximum Psuedo-likelihood Estimation in Network Tomography," *IEEE Transactions on Signal Processing*, Volume 512, Number 8, pp. 2043-2053.
- [36] Liang, G., Taft, N. and Bin, Y. (2006), "A fast lightweight approach to origin-destination IP traffic estimation using partial measurements," *IEEE Transactions on Information Theory*, 52, 2634-2648
- [37] Michailidis, G. (2006), "Data visualization through their graph representations," in *Handbook of Computational Statistics: Data Visualization*, Chen, Hardle and Unwin (eds.), Springer Verlag, Heidelberg (to appear).
- [38] Michailidis, G., Nair, V. and Xi, B. W. (2007), "Fast Edtimation Methods for Estimation of Loss Rates in Active Network Tomography," (preprint).
- [39] Moon, S. B., Skelly, P. and Towsley, D. (1999), "Estimation and Removal of Clock Skew from Network Delay Measurements," *Proceedings of IEEE Infocom*, New York, NY, 227-234
- [40] Paxson, V. (1997), "End-to-end routing behavior in the Internet," *IEEE/ACM Transactions on Networking*, 5, 601-615
- [41] Paxson, V. (1998), "On Calibrating Measurements of Packet Transit Times," *Proceedings of the ACM Sigmetrics*, Madison, WI
- [42] Peterson, L. and Davie, B. (2003), *Computer Networks: A Systems Approach*, Morgan Kaufmann, San Francisco, CA
- [43] Shih M. F., and Hero, A. O., (2003) "Unicast-based inference of network link dely distributions with finite mixture models," *IEEE Transactions on Singal Processing*, 51, 2219-2228.
- [44] Solka, J. L., Marchette, D. J. and Wallet, B. (2000), Statistical visualization methods in intrusion detection, *Computing Science and Statistics*, 32, 16-24
- [45] Stoumbos, Z., Reynolds, M. R., Ryan, T. P., and Woodall, W. H. (2000), "The State of Statistical Process Control as We Proceed into the 21st Century," *Journal of the American Statistical Association*, 95, 992-998.
- [46] Tsang, Y., Coates, M. and Nowak, R. D. (2003), "Network Delay Tomography," *EEE Transactions on Signal Processing*, 8, pp 2125-2135.
- [47] Vardi, Y., (1996) "Network tomography: Estimating source-destination traffic intensities from source data," *JASA*, 91, 365-377.

- [48] Xi, B. W., Michailidis, G. and Nair, V. (2006) “Estimating Network Loss Rates Using Active Tomography,” *JASA, T&M*, 1430-1449.
- [49] Yang, L. and Michailidis, G. (2007), “Sample Based Estimation of Network Traffic Flow Characteristics,” *Proceedings of Infocom 2007*
- [50] Yang, X. Y. (2007) *Network Monitoring*, Unpublished Ph. D. Dissertation, University of Michigan, Ann Arbor (in preparation).
- [51] Zhang, L., Liu, Z. and Xia, C. H. (2002), “Clock Synchronization Algorithms for Network Measurements,” *Proceedings Infocom*, New York, NY
- [52] Zhang, Y., Roughan, M., Lund, C. and Donoho, D. (2003), “An information theoretic approach to traffic matrix estimation,” *ACM SIGCOMM Proceedings*, Karlsruhe, Germany.