

Multicast Inference of Packet Delay Variance at Interior Network Links

N.G. Duffield¹

F. Lo Presti^{1,2}

¹AT&T Labs—Research ²University of Massachusetts

Abstract—End-to-end measurement is a common tool for network performance diagnosis, primarily because it can reflect user experience and typically requires minimal support from intervening network elements. Challenges in this approach are: (i) to identify the locale of performance degradation; and (ii) to perform measurements in a scalable manner for large and complex networks. In this paper we show how end-to-end delay measurements of multicast traffic can be used to estimate packet delay variance on each link of a logical multicast tree. The method does not depend on cooperation from intervening network elements; multicast probing is bandwidth efficient. We establish desirable statistical properties of the estimator, namely consistency and asymptotic normality. We evaluate the approach through model based and network simulations. The approach extends to the estimation of higher order moments of the link delay distribution.

Keywords—End-to-end measurement, queuing delay, estimation theory, multicast trees, network tomography.

I. INTRODUCTION

A. Background and Motivation.

Monitoring the performance of large communications networks and diagnosing the causes of its degradation is a challenging problem. There are two broad approaches to performance diagnosis. In the *internal* approach, direct measurements are made at or between network elements, e.g., of packet loss or delay. This approach has a number of potential limitations: it may not be available for general users; coverage may not span paths of interest; measurements may be disabled during period of high load; there are issues of scale in gathering and correlating the measurements in large networks; how should per hop measurements be composed to an end-to-end view?

This motivates *external* approaches, diagnosing the network through end-to-end measurements, without necessarily assuming the cooperation of network elements on the path. There has been much recent experimental work to understand the phenomenology of end-to-end performance (e.g., see [1], [2], [8], [21], [16], [23], [24], [26]); several measurement infrastructure projects are in development (including CAIDA [6], Felix [10], IPMA [12], NIMI [15], Surveyor [30]) with the aim to collect and analyze end-to-end measurements across a mesh of paths between a number of hosts. Standard diagnostic tools for IP networks, ping and traceroute report roundtrip loss and delay. A recent refinement of this approach, *pathchar* [13], estimates hop-by-hop link capacities, packet delay and loss rates. *pathchar* is still under evaluation; initial experience indicates many packets are required for inference leading to either high load of measurement traffic or long measurement intervals, although adaptive approaches can reduce this [9]. More broadly,

This work was sponsored in part by the DARPA and the Air Force Research Laboratory under agreement F30602-98-2-0238.

Address: AT&T Labs—Research, 180 Park Avenue, Florham Park, NJ 07932, USA; E-mail: {duffield,lopresti}@research.att.com

measurement approaches based on Time To Live (TTL) expiry require the cooperation of network elements in returning Internet Control Message Protocol (ICMP) messages. In future, encapsulation may hide TTL from higher layers that would see just a single hop between tunnel endpoints. Finally, the success of active measurement approaches to performance diagnosis may itself cause increased congestion if intensive probing techniques are widely adopted.

In response to some of these concerns, a multicast-based approach to active measurement has been proposed recently in [3], [4]. The idea is that correlation in performance seen on *intersecting* end-to-end paths can be used to draw inferences about the performance characteristics of their common portion, without cooperation from the network. Multicast traffic is well suited for this since a given packet only occurs once per link in the (logical) multicast tree. Characteristics such as loss and end-to-end delay seen at different endpoints are highly correlated. Another advantage is in scalability. Suppose packets are exchanged on a mesh of paths between a collection of N measurement hosts stationed in a network. With unicast the probe load on the network may grow proportionally to N^2 in some links of the network. With multicast the load grows proportionally only to N .

B. Contribution

In this paper we describe a method to infer the variance of internal link delays from measured end-to-end delays of multicast probe packets. It is assumed that the link delays are independent random variables, both spatially (i.e., between different links) and temporally (i.e., between different packets); later we discuss the impact of violation of these assumptions. The method rests on (generalizations of) the following observation. Consider the logical multicast topology of Figure 1(left), in which packets are multicast from the root 0 to receivers 1 and 2. D_i is the random delay on link i , and the source-to-leaf delays from the root 0 to the leaf nodes 1 and 2 are $X_1 = D_k + D_1$ and $X_2 = D_k + D_2$ respectively. Then a calculation that we perform later shows that, under the independence assumption,

$$\text{Var}(D_k) = \text{Cov}(X_1, X_2), \quad (1)$$

i.e., we express the variance of an internal link delay in terms of the covariance of the source-to-leaf delays. We can form an unbiased estimate of $\text{Cov}(X_1, X_2)$ directly from end-to-end measurements; this constitutes an unbiased estimate of $\text{Var}(D_k)$. The same method extends to higher order moments; when the node k has branching ratio m , we are able to estimate the first m moments of D_k ; see Figure 1(right). We specify the delay model in Section II and describe the basic moment estimators in Section III.




Fig. 1. LEFT: Two leaf tree. RIGHT: m -leaf tree.

Here we focus on estimation of the delay variance, either on individual links, or from the root to a given node. In Section IV we show how the above scheme can be used to obtain multiple unbiased estimates of the variance of the delay from the root to a given node k , one estimator for every pair of leaf nodes descended through different children of k . The estimates are consistent, i.e., they converge in probability to the true variance as the number of probes grows to infinity. Any convex combination of these estimators shares these properties; although the rate of convergence will be different in each case. This rate can be used to distinguish between the estimators. We show how to choose the weights of the combination in order to obtain the combination with the fastest asymptotic rate of convergence.

Packet loss reduces the number of packets available for delay estimation, hence increasing estimator variance. In Section V we quantify this for an estimation scheme that makes maximal use of information from surviving packets, using all packets reaching a given node pair for which a covariance estimator is calculated.

The model used here also assumes temporal independence, i.e., that delays between successive probe packets at a given node are not dependent. This can be arranged for by making the interprobe times greater than the queueing timescale. However, for a wide class of temporally dependent delay processes—we require only ergodicity—the consistency of the estimators is unaffected, i.e., they still converge to the true values as the number of probes grows to infinity. However, the rate of convergence may be slower.

In Section VI we report two types of simulation: (i) model simulations with packet delay chosen pseudo-randomly according to a given distribution; and (ii) ns [22] simulations that represent probe traffic mixed in with background traffic of TCP and UDP sessions; probe delay occurs as result of queueing against background traffic, probe loss is due to buffer overflow. The model simulations allow us to compare the theoretical prediction with a model in a controlled manner. We verify the accuracy of the delay variance estimator. The variance of the variance estimators over many simulation runs is conformant with the model; this verifies the benefit in accuracy of using the minimum variance estimator. The ns simulation allow us to investigate the performance of the inference method in a more realistic setting in which the independence assumption may not be exactly satis-

fied. We find that dependence between delays in different links is smaller when buffers are larger, and that inference is correspondingly more accurate. In a 12 node topology we found the typical error in estimation to be about 20% in some configurations, this based on a sample size of 1,000 probes. We believe this is sufficiently accurate to distinguish links with high delay variance. As far as we are aware there are no studies in deployed networks that measure delay correlation between different nodes. However, we believe that large and long-lasting spatial dependence is unlikely in a real network such as the Internet because of its traffic and link diversity.

C. Implementation Requirements

Since the data for delay inference comprises one-way packet delays, the primary requirement is the deployment of measurement hosts with synchronized clocks. (Actually, since delay covariances are invariant under time-shifts, the absolute times need not be synchronized, provided that the rates are identical). Using Global Positioning System (GPS) timing it is possible to make one-way delay measurements accurate to within tens of microseconds or better. GPS is currently used or planned in several of the measurement infrastructures mentioned earlier. The Network Time Protocol (NTP) [17] is more widely deployed, but provides accuracy in only the order a few tens of milliseconds, a resolution at least as coarse as the queueing delays in practice. An alternative approach to calibration and synchronization of clocks has been developed in [18], [25], [27].

Another requirement is to know the multicast topology. There is a multicast-based measurement tool, mtrace [19], already in use in the Internet. mtrace reports the route from a multicast source to a receiver, along with other information about that path such as per-hop loss and rate. Presently it does not support delay measurements. A potential drawback for larger topologies is that mtrace does not scale to large numbers of receivers as it needs to run once for each receiver to cover the entire multicast tree. In addition, mtrace relies on multicast routers responding to explicit measurement queries; the feature that can be administratively disabled. An alternative approach that is closely related to the work on multicast-based loss inference [3], [4] is to infer the logical multicast topology directly from measured probe statistics; see [5], [28]. The delay variance estimates of the present paper can also be used to infer topology.

This method does not require cooperation from the network.

D. Use of Delay Variance Estimates

Although prior work has characterized end-to-end delays [1], [21], [24], to the best of our knowledge there is no generally accepted model for per link delays in real networks. Without a model it is difficult to map a given inferred value of the link delay variance to a specific value of a quality metric, such as the probability of queuing delay exceeding a given value. Nevertheless, we believe that knowledge of the per link delay variance will be increasingly useful for the following reasons:

Model Development. The mapping problem just described will become easier upon development of delay models. We expect these to arise from two sources. The first is the development of measurement infrastructure projects in which selected links are instrumented for one-way delay measurements. The second is the development of multicast-based estimators for the link delay *distribution* from end-to-end measurements, using a more computationally intensive technique proposed in a companion paper [14]. We anticipate that this will allow the development of link delay distribution models, with the distribution inferred from network measurements.

Ordering. Identification of links with highest delay variance suggests candidate for links on which performance is degraded for delay sensitive applications.

Delay and Delay Variation. The variance of the packet delay (on a link or path) can be used to estimate or bound the variance of the interpacket delay variation. Let D^i be the delay encountered by packet i on a given link. The interpacket delay variation (or jitter) between packets i and $i+1$ on the link is $J^i = D^{i+1} - D^i$; a similar notion applies to end to end delay. Observe

$$\text{Var}(J^i) = \text{Var}(D^i) + \text{Var}(D^{i+1}) - 2\text{Cov}(D^i, D^{i+1}). \quad (2)$$

Assuming $D(\cdot)$ to be stationary, the first two terms on the RHS of (2) are equal, while under the assumption of temporal independence the last term is zero, and so $\text{Var}(J^i) = 2\text{Var}(D^i)$. Measurements of end-to-end delays in the Internet [1] show that end-to-end delays successive packets are only slightly dependent when the interpacket time is longer than the typical queuing timescales. Stronger dependence is found at shorter timescales: successive packets are more likely to queue together. With positive correlation between successive probe delays $\text{Cov}(D^i, D^{i+1}) > 0$; in this case $\text{Var}(J^i)$ is bounded above by $2\text{Var}(D^i)$, a quantity that we can estimate from end-to-end measurements.

Topology Inference. If the logical multicast topology is not initially known, it can be inferred from delay variances. This technique uses the estimated variance of the cumulative delay from the source to a given node. Consequently we shall be interested here in the estimation of cumulative delay variance as well as link delay variance.

II. THE TREE AND DELAY MODELS

We identify the physical multicast tree as comprising actual network elements (the nodes) and the communication links than join them. The logical multicast tree comprises the branch points of the physical tree, and the logical links between them.

The logical links comprise one or more physical links. Thus each node in the logical tree, except the leaf nodes and possibly the root, must have 2 or more children. We can construct the logical tree from the physical tree by deleting all links with one child and adjusting the links accordingly by directly joining its parent and child.

Let $\mathcal{T} = (V, L)$ denote a logical multicast tree with nodes V and links L . We identify one node, the root 0, with the source of probes, and $R \subset V$ will denote the set of leaf nodes (identified as the set of receivers). The set of children of node $j \in V$ is denoted by $d(j)$. Each node, k , apart from the root has a parent $f(k)$ such that $(j, k) \in L$. Define recursively the compositions $f^n = f \circ f^{n-1}$ with $f^1 = f$. Nodes are said to be siblings if they have the same parent. If $k = f^m(j)$ for some $m \in \mathbb{N}$ we say that j is descended for k (or equivalently that k is an ancestor of j) and write the corresponding partial order in V as $j \prec k$. $i \vee j$ will denote the minimal common ancestor of i and j in the \prec -ordering, i.e., the unique common ancestor that is \prec all others.

We associate each node k a random variable D_k taking values in the extended positive real line $\mathbb{R} = \mathbb{R}_+ \cup \{\infty\}$. By convention $D_0 = 0$. D_k is the random delay that would be encountered by a packet attempting to traverse the link $(f(k), k) \in L$. The value $D_k = \infty$ indicates the packet is lost on the link. The delay experienced on the path from the root 0 to a node k is $X_k = \sum_{j \prec k} D_j$. We assume that the D_k are independent. Let $\alpha_k = \mathbb{P}[D_k < \infty]$, the probability of successful transmission over link k .

III. NON-PARAMETRIC ESTIMATION OF DELAY DISTRIBUTION MOMENTS

In this section we present a class of non-parametric estimators of the delay distribution. We assume initially that all delays are finite: $\mathbb{P}[D_k = \infty] = 0$. Consider first a logical subtree formed by the root 0, and a non-leaf node k with two descendants 1 and 2 that are leaf nodes; see Figure 1(left). Now

$$\begin{aligned} \text{Cov}(X_1, X_2) &= \text{Cov}(X_k + (X_1 - X_k), X_k + (X_2 - X_k)) \\ &= \text{Cov}(X_1 - X_k, X_k) + \text{Cov}(X_2 - X_k, X_k) \\ &\quad + \text{Cov}(X_1 - X_k, X_2 - X_k) + \text{Var}(X_k) \\ &= \text{Var}(X_k), \end{aligned} \quad (3)$$

since by assumption of mutual independence of the link delays D_k , the random variables $X_k, X_1 - X_k$ and $X_2 - X_k$ are mutually independent. Hence any unbiased estimator of $\text{Cov}(X_1, X_2)$ is also an unbiased estimator of $\text{Var}(X_k)$. Let $X_1^{(i)}, X_2^{(i)}$, $i = 1, 2, \dots, n$ be measured end-to-end delays between the root 0 and leaf nodes 1 and 2 respectively. Abbreviate $\text{Cov}(X_j, X_k)$ by s_{jk} and write s_{kk} as s_k . We estimate s_k by a uniformly minimum variance unbiased estimator of s_{12} , namely \hat{s}_{12} where

$$\hat{s}_{ij} = \frac{1}{n-1} \left(\sum_{m=1}^n X_i^{(m)} X_j^{(m)} - \frac{1}{n} \sum_{m, m'=1}^n X_i^{(m)} X_j^{(m')} \right) \quad (4)$$

At a node with branching ratio m we are able to estimate the first m moments of the delay on the shared portion of the

path from the root; see Figure 1(right). The cumulant generating function of the m leaf delays $\mathbf{X} = (X_1, \dots, X_m)$ is defined for $\theta \in \mathbb{R}^m$ by

$$\lambda(\theta; \mathbf{X}) = \log E[\exp(\sum_{i=1}^m \theta_i X_i)]. \quad (5)$$

The cumulants are defined by partial differentiation w.r.t. the components θ_i (when derivatives exist): for indices $j_1, \dots, j_m \in \mathbb{Z}_+$ set

$$K^{j_1, \dots, j_m}(\mathbf{X}) = \left(\prod_{i=1}^m \frac{\partial^{j_i}}{\partial \theta_i^{j_i}} \right) \lambda(\theta; \mathbf{X}) \Big|_{\theta_i=0} \quad (6)$$

The first and second cumulants K^1 and K^2 of a single random variable are its mean and variance respectively. Knowing the cumulants of a set of random variables is equivalent to knowing their joint distribution. The cumulants of D_0 are related to those of the X_i as follows. Set $\mathbf{1} = (1, \dots, 1) \in \mathbb{R}^m$.

Theorem 1: $K^1(\mathbf{X}) = K^m(X_k)$. Hence any unbiased estimator of $K^1(\mathbf{X})$ is also an unbiased estimator of $K^m(X_k)$.

Proof: Observe $K^1(X_1, \dots, X_m) = K^1(X_1 - X_k, \dots, X_m - X_k) + K^1(\mathbf{1}X_k) = K^m(X_k)$. The first equality is because K is affine in each of its arguments, the second because the cumulant of a set of independent random variables is zero. ■

IV. DELAY VARIANCE ESTIMATION ON GENERAL TREES

In a general tree with receiver set R , let $Q(k) = \{\{i, j\} \subset R \mid i \vee j = k, \}$ be the set of distinct pairs of leaf-nodes whose \prec -least common ancestor is k . Any convex combination $\sum_{\{i,j\} \in Q(k)} \mu_{ij} \hat{s}_{ij}$ (i.e., with the $\mu_{ij} \geq 0$ and summing to 1) is also an unbiased estimator of s_k . An example the **uniform estimator**

$$\frac{1}{\#Q(k)} \sum_{\{i,j\} \in Q(k)} \hat{s}_{ij}. \quad (7)$$

One potential disadvantage with the uniform estimator is that high variance of one of the summands may lead to high estimator variance overall. This motivates choosing convex combinations that are functions of the end-to-end delays themselves in order to reduce variance. In this section we shall assume that all delays are finite with bounded fourth moments. We shall relax the finiteness assumption in Section V.

We formalize the notion of (possibly random) convex combinations of \hat{s}_{ij} through a **covariance aggregator**. For $S \subset R$ let $\mathcal{F}_n(S)$ denote the σ -algebra generated by the end-to-end delays $(X_k)_{k \in S}$ (i.e. the set of events that can be determined from knowing $(X_k)_{k \in S}$). A covariance aggregator μ is sequence $(\mu(n))_{n \in \mathbb{N}}$ of random vectors $\{\mu_{ij}(n) : \{i, j\} \in Q(k); k \in V \setminus R\}$ with $0 \leq \mu_{ij}(n) \leq 1$ and $\sum_{\{i,j\} \in Q(k)} \mu_{ij}(n) = 1$ for each $k \in V \setminus R$. We assume each μ_n to be $\mathcal{F}_n(R)$ -measurable, i.e., it is a function of the measured delays of the first n probes. We will usually suppress the explicit dependence on the number of probes n . Let $\hat{s} = \{\hat{s}_{ij}(n) : \{i, j\} \in Q(k); k \in V \setminus R\}$ be a family of estimators, $\hat{s}_{ij}(n)$ being an $\mathcal{F}_n(i, j)$ -measurable unbiased estimator of $\text{Var}(X_{i \vee j})$. Then we estimate $\text{Var}(X_k)$ by

$$V_k(\mu, \hat{s}) = \sum_{\{i,j\} \in Q(k)} \mu_{ij} \hat{s}_{ij} \quad (8)$$

A covariance aggregator is called **deterministic** if it does not depend on the $X^{(i)}$. We denote the set of such aggregators with indices in $Q(k)$ by \mathcal{D}_k . An example is the **uniform aggregator** that was used in the uniform estimator (7): $\mu_{ij} = (\#Q(k))^{-1}$, $i \vee j = k$. Define the covariance matrix

$$C_{(ij),(\ell m)} = \text{Cov}(Z_i Z_j, Z_\ell Z_m), \quad (9)$$

where $Z_i = X_i - E[X_i]$. We will use $C(k) = [C_{(ij),(\ell m)}]_{(ij),(\ell m) \in Q(k)}$ to denote the matrix obtained by letting the indices (ij) and (ℓm) in (9) run over $Q(k)$; this is a submatrix of the matrix $C^0(k)$ obtained by taking the indices unrestricted over $Q^0(k)$ defined as the set of binary subsets of $R(k)$.

A. Minimum Variance Estimation for Cumulative Delays

In the next theorem we characterize the asymptotic distribution of the \hat{s}_{ij} as $n \rightarrow \infty$, and give a form for the estimator $V_k(\mu, \hat{s})$ of minimum cumulative variance.

Theorem 2: (i) For each $k \in V \setminus R$ the random variables $\{\sqrt{n}(\hat{s}_{ij} - s_k) \mid \{i, j\} \in Q(k)\}$ converge in distribution as $n \rightarrow \infty$ to a multivariate Gaussian random variable with mean 0 and covariance matrix $C(k)$. Hence the \hat{s}_{ij} are consistent estimators of s_k and so is $V(\mu, \hat{s})$. For any deterministic covariance aggregator μ , $\sqrt{n}(V_k(\mu, \hat{s}) - s_k)$ converges in distribution as $n \rightarrow \infty$ to a Gaussian random variable of mean zero and variance $\mu \cdot C(k) \cdot \mu$.

(ii) The minimal asymptotic variance $\inf_{\mu \in \mathcal{D}_k} \mu \cdot C(k) \cdot \mu$ is achieved when

$$\mu_{ij} = \mu_{ij}^*(C(k)) := (C(k)^{-1} \cdot \mathbf{1})_{(ij)} / \mathbf{1} \cdot C(k)^{-1} \cdot \mathbf{1} \quad (10)$$

where $C(k)^{-1}$ denotes the inverse matrix of $C(k)$ and $\mathbf{1}_{(ij)} = 1$, $\{i, j\} \in Q(k)$. The corresponding asymptotic variance of the variance estimator is $(\mathbf{1} \cdot C(k)^{-1} \cdot \mathbf{1})^{-1}$.

Proof: (i) The proof follows from standard results in multivariate analysis; convergence to the stated Gaussian random variable follows by Corollary 1.2.18 in [20]

(ii) Since the μ_{ij} sum to 1, the proof follows by considering the constrained minimization of $\mu \cdot C(k) \cdot \mu - 2\beta \mu \cdot \mathbf{1}$ with Lagrange multiplier β . As a covariance matrix, $C(k)$ is positive definite and hence invertible; minimization of the convex function of μ takes place at the stationary point $\mu = \beta C(k)^{-1} \cdot \mathbf{1}$. This yields $\mu^*(C(k))$ upon normalization. The corresponding minimal asymptotic variance is $\mu^*(C(k)) \cdot C(k) \cdot \mu^*(C(k)) = (\mathbf{1} \cdot C(k)^{-1} \cdot \mathbf{1})^{-1}$. ■

Operationally, the coefficients μ_{ij} of the minimum variance estimator $V_k(\mu^*(C(k)), \hat{s})$ of Theorem 2 are to be calculated from an estimate of the covariance matrix $C(k)$. Let $Z_i^{(m)} = X_i^{(m)} - \frac{1}{n} \sum_{m=1}^n X_i^{(m)}$. Let $\hat{C}(k)$ denote the empirical covariance matrix with entries

$$\begin{aligned} \hat{C}(k)_{(ij), (i'j')} &= \frac{n^2}{(n-1)s} \left(\sum_{m=1}^n Z_i^{(m)} Z_j^{(m)} Z_{i'}^{(m)} Z_{j'}^{(m)} \right. \\ &\quad \left. - \frac{1}{n} \sum_{m=1}^n Z_i^{(m)} Z_j^{(m)} \sum_{m=1}^n Z_{i'}^{(m)} Z_{j'}^{(m)} \right) \quad (11) \end{aligned}$$

$\widehat{C}(k)$ is an estimator of $C(k)$. Estimating $\mu^*(C(k))$ by $\mu^*(\widehat{C}(k))$ and s_k by $V_k(\mu^*(\widehat{C}(k)), \widehat{s})$ potentially introduces bias and increases variance in the estimation of the s_k . However, the following Theorem shows that it is consistent and has the same asymptotic variance as $V_k(\mu^*(C), \widehat{s})$.

Theorem 3: $V_k(\mu^*(\widehat{C}(k)), \widehat{s})$ is a consistent estimator of s_k . $\sqrt{n}(V_k(\mu^*(\widehat{C}(k)), \widehat{s}) - s_k)$ converges in distribution to a Gaussian random variable of mean zero and variance $(\mathbf{1} \cdot C(k)^{-1} \cdot \mathbf{1})^{-1}$.

Proof: Clearly $\widehat{C}(k)$ converges almost surely to $C(k)$ as $n \rightarrow \infty$. Since matrix inversion is continuous on the set of strictly positive definite matrices, $\mu^*(\widehat{C}(k))$ converges almost surely (to $\mu^*(C(k))$); since each \widehat{s}_{ij} converges to $s_{ij} = s_k$, $V_k(\mu^*(\widehat{C}(k)), \widehat{s})$ is consistent.

By the δ -method (see e.g. [29]), $\sqrt{n}(V(\mu^*(\widehat{C}(k)), \widehat{s}) - s_k)$ converges to a Gaussian random variable with mean 0 and variance $\alpha \cdot C^0(k) \cdot \alpha$, where for $(\ell, m) \in Q^0(k)$,

$$\alpha_{\ell m} = \frac{\partial}{\partial s_{\ell m}} \sum_{\{i,j\} \in Q(k)} \mu_{ij}^*(C(k)) s_{ij}. \quad (12)$$

Differentiating,

$$\begin{aligned} \alpha_{\ell m} &= \mu_{\ell m}^*(C(k)) \chi_{Q(k)}(\{\ell, m\}) \\ &+ \sum_{\{i,j\} \in Q(k)} s_{ij} \frac{\partial}{\partial s_{\ell m}} \mu_{ij}^*(C(k)), \end{aligned} \quad (13)$$

where $\chi_{Q(k)}$ denotes the indicator function of the set $Q(k)$. But $s_{ij} = s_{i \vee j} = s_k$ for $\{i, j\}$ in $Q(k)$ and so is constant in the sum. Since the μ_{ij}^* sum to 1, the sum in (13) is zero. Hence $\alpha \cdot C^0(k) \cdot \alpha = \mu^*(C(k)) \cdot C(k) \cdot \mu^*(C(k))$. ■

B. Minimum Variance Estimation for Link Delays

We can estimate the link delay variance as the difference of two cumulative variances since

$$\text{Var}(X_k) = \text{Var}(X_{f(k)} + D_k) = \text{Var}(X_{f(k)}) + \text{Var}(D_k), \quad (14)$$

by the independence assumption on link delays. An unbiased estimator of $r_k := \text{Var}(D_k)$ is $V_k(\mu^*(C(k)), \widehat{s}) - V_{f(k)}(\mu^*(C(f(k))), \widehat{s})$. We now show that joint optimization of the aggregators in V_k and $V_{f(k)}$ will result in an estimator of lower variance.

Given a pair $\mu = (\mu(k), \mu(f(k))) \in \mathcal{D}_k \times \mathcal{D}_{f(k)}$ of deterministic covariance aggregators with indices in $Q(k)$ and $Q(f(k))$ respectively, we can form an unbiased estimate of r_k as

$$W_k(\mu, \widehat{s}) := V_k(\mu(k), \widehat{s}) - V_{f(k)}(\mu(f(k)), \widehat{s}) \quad (15)$$

Let $C'(k)$ denote the $\#Q(k) + \#Q(f(k))$ dimensional matrix written in block form

$$C'(k) = \begin{pmatrix} C(k) & C(k, f(k)) \\ C(k, f(k))^T & C(f(k)) \end{pmatrix}, \quad (16)$$

where $C(k, f(k))$ is the $\#Q(k) \times \#Q(f(k))$ matrix of covariances $[C_{(ij), (\ell m)}]_{(ij) \in Q(k), (\ell m) \in Q(f(k))}$. Then statements analogous to Theorem 2(ii) follow straightforwardly, using parallel arguments. In particular $\sqrt{n}(W_k(\mu, \widehat{s}) - r_k)$ converges to a

Gaussian random variable of mean 0 and variance $\mu \cdot C'(k)^{-1} \mu$ and the minimum over deterministic aggregators of the asymptotic variance takes the value $(c_1 + c_2 + 2c_3)/(c_1 c_2 - c_3^2)$ where $c_1 = \mathbf{1}_k \cdot C(k)^{-1} \cdot \mathbf{1}_k$, $c_2 = \mathbf{1}_{f(k)} \cdot C(f(k))^{-1} \cdot \mathbf{1}_{f(k)}$ and $c_3 = \mathbf{1}_{f(k)} \cdot C(k, f(k))^{-1} \cdot \mathbf{1}_k$. (Here the subscripts on $\mathbf{1}_k, \mathbf{1}_{f(k)}$ distinguish the subspaces in which these vectors live).

C. Criteria for Assessing Inference Reliability

In sections IV-A and IV-B we derived expressions for the variances of estimates of the cumulative and link delays respectively. For a given delay variance estimate, we can associate its variance by using the plug-in estimator for the corresponding analytic expression. This enables use to find confidence intervals for the estimates that will be asymptotically accurate for large n . For example, if we use n probes to form the estimate $V_k(\mu^*(\widehat{C}(k)), \widehat{s})$, we associate with this a variance σ^2/n where $\sigma^2 = (\mathbf{1} \cdot \widehat{C}(k)^{-1} \cdot \mathbf{1})^{-1}$. We write confidence limits for the estimate as

$$V_k(\mu^*(\widehat{C}(k)), \widehat{s}) \pm z_{\delta/2} \sigma/n, \quad (17)$$

where $z_{\delta/2}$ denotes the number that cuts off an area $\delta/2$ in the right tail of the standard normal distribution. This is used for a confidence interval of level $1 - \delta$.

V. IMPACT OF LOSS ON ESTIMATOR VARIANCE

We relax the assumption of finite delays. We identify infinite delays with packet loss, although the same results would hold were we to treat as lost any packet with source to leaf delay greater than some finite value. The link and cumulative delay random variables will be denoted by D'_k and X'_k respectively each possibly taking the value ∞ . We use D_k to be the distribution of D'_k conditional on $D'_k < \infty$, and similarly for X_k . We assume throughout that the D_k have finite fourth moments. Since we are interested in delay variance, we want to estimate $\text{Var}(X_k)$ and $\text{Var}(D_k)$ even in the presence of packet loss. For estimation, the effect of packet loss is to reduce the number of delay samples available, and hence to increase the variability of the estimates. A simple way to apply the foregoing theory is to restrict attention to only those packets that reach every receiver (or at least at every receiver in $R(k)$ when estimating s_k). A disadvantage of this approach is that it does not scale well as the topology grows: assuming link loss rates to be bounded away from zero, the proportion of packets reaching all receivers in a tree decays geometrically fast in the number of links in the tree.

An alternative that wastes less data is to calculate pairwise estimates of \widehat{s}_{ij} that use all packets received at i and j . Let us formalize this. For a subset of receivers $S \subset V$ define $I_n(S) = \{i \in \{1, 2, \dots, n\} \mid X_j^{(i)} < \infty \forall j \in S\}$: the subset of the first n probes that are received at all nodes in S ; set $N_n(S) = \#I_n(S)$. We will sometimes write $I_n(i_1, \dots, i_r)$ for $I_n(\{i_1, \dots, i_r\})$, and similarly for N_n . For $S \subset R$ let $V(S)$ be the set of nodes in the minimal tree spanning 0 and S . Set $B(S) = \prod_{i \in V(S)} \alpha_i$, where α_i is the probability of successful transmission over link k . Clearly $n^{-1} N_n(S)$ converges almost surely to $B(S)$ as $n \rightarrow \infty$. Estimator variance can be reduced by using all packets in $I_n(i, j)$ to estimate s_{ij} , not just those in




Fig. 2. TREES USED IN SIMULATIONS. LEFT: 8-leaf binary tree for model simulations ; RIGHT: Heterogeneous 7-leaf tree for ns simulation.

$I_n(R(i \vee j))$. Define

$$\hat{v}_{ij} = \frac{1}{N} \left(\sum_m X_i^{(m)} X_j^{(m)} - \frac{1}{N} \sum_{m,m'} X_i^{(m)} X_j^{(m')} \right) \quad (18)$$

where N abbreviates $N_n(i, j)$ and in the sums m, m' run over $I_n(i, j)$. \hat{v}_{ij} is an unbiased estimate on s_{ij} . Analogously to the previous results we have:

Theorem 4: (i) For each $k \in V \setminus R$ the random variables $\{\sqrt{n}(\hat{v}_{ij} - s_k) \mid \{i, j\} \in Q(k)\}$ converge in distribution as $n \rightarrow \infty$ to a multivariate Gaussian random variable with mean 0 and covariance matrix $G(k)_{(ij),(\ell m)} = C(k)_{(ij),(\ell m)} B(i, j, \ell, m) / (B(i, j) B(\ell, m))$. Hence the \hat{v}_{ij} are consistent estimators of s_k and so is $V_k(\mu, \hat{v})$ for any deterministic covariance aggregator μ . For any deterministic covariance aggregator μ , $\sqrt{n}(V_k(\mu, \hat{v}) - s_k)$ converges in distribution as $n \rightarrow \infty$ to a Gaussian random variable of mean zero and variance $\mu \cdot G(k) \cdot \mu$.

(ii) The minimal asymptotic variance $\inf_{\mu \in \mathcal{D}_k} \mu \cdot G(k) \cdot \mu$ is achieved when $\mu = \mu^*(G)$; the corresponding minimal asymptotic variance is $(\mathbf{1} \cdot G(k)^{-1} \cdot \mathbf{1})^{-1}$.

(iii) $V_k(\mu^*(\hat{G}), \hat{v})$ has the same asymptotic properties as $V_k(\mu^*(G), \hat{v})$ where the estimated covariance \hat{G} is defined by

$$\frac{N_n(i, j) N_n(k, \ell)}{N_n(i, j, k, \ell)} \hat{G}_{(ij), (k\ell)} = \sum_m Z_i^{(m)} Z_j^{(m)} Z_k^{(m)} Z_\ell^{(m)} - \frac{1}{N_n(i, j, k, \ell)} \sum_m Z_i^{(m)} Z_j^{(m)} \sum_{m'} Z_k^{(m')} Z_\ell^{(m')} \quad (19)$$

where the sums run over $I_n(i, j, k, \ell)$.

The corresponding version of the minimum variance link delay variance estimator follows by replacing C by G and \hat{s} by \hat{v} throughout Section IV-B.

VI. SIMULATION EVALUATION

We conducted two types of simulation: (i) model simulation with packet delay chosen pseudo-randomly according to a given distribution; and (ii) ns [22] simulations that represent probe

traffic mixed in with background traffic of TCP and UDP sessions; probe delay occurs as a result of queueing against background traffic, probe loss is due to buffer overflow. The model simulations allow us to compare the theoretical prediction with a model in a controlled manner; their purpose is to show that the statistical properties of the estimators conform to the model used. The ns simulation allow us to investigate the performance of the inference method in a more realistic setting in which the model assumption of independence may not be exactly satisfied. Their purpose is to investigate conformance of the predicted delay variances with those occurring in the network interior.

A. Model Simulations

The model simulation used an 8 leaf binary tree (see Figure 2(left)); delays were exponentially distributed. The delay variances were heterogeneous: leaf links 8 and 15 had delay variance 10, all other links had delay variance 1. Losses were not modeled. This heterogeneity was chosen in order to evaluate the advantages of the minimum variance estimator. We present a representative set of results from experiments for the link delay variance W and the cumulative delay variance V .

A.1 Convergence

Figure 3 shows the variance of the cumulative delay variance from sources to nodes $k = 1, 6, 10, 15$ in Figure 2(left), plotted as a function of the number of probes. On the left is the theoretical variance $\text{Var}(V_k(\mu^*(C), \hat{s}))$; on the right the empirical variance from 100 samples of $V_k(\mu^*(C), \hat{s})$ found by simulation. Observe in both cases the decay of the variance towards 0 as the number of probes increases; furthermore the experimental variance is very close to the theoretical values over the range of probe numbers.

Figure 4 shows detail from a single simulation; sample paths of the link variance estimator $W_k(\mu, \hat{s})$ for links $k = 1, 3, 5, 10$ as function of the number of probes, for up to 10,000 probes. On the left figure, the aggregator μ is uniform, on the right, the minimum variance aggregator $\mu^*(\hat{C})$. Observe in both cases that the estimate approaches the model value, 1, as the number of probes increases.

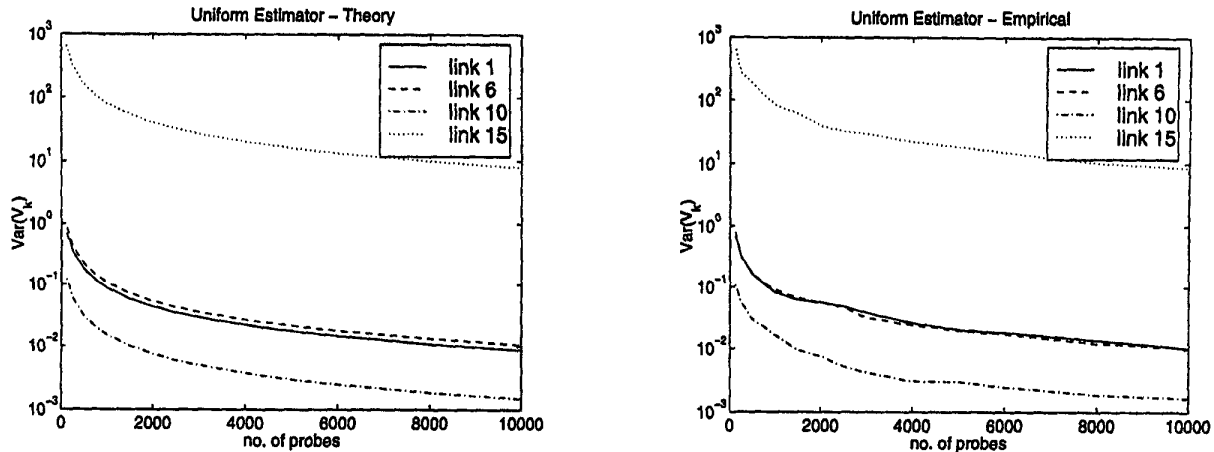


Fig. 3. VARIANCE OF ESTIMATED VARIANCE. Cumulative Delay Variance to nodes 1,6,10,15 in Figure 2(left). LEFT: Calculate Variance; RIGHT: Empirical Variance from 100 simulations.




Fig. 4. SAMPLE PATHS OF LINK VARIANCE ESTIMATORS. For up to 10,000 probes. LEFT: Uniform Estimator, RIGHT: Minimum Variance Estimator

Weight μ_{ij}	Link pairs (i, j)
0.000018	(8,15)
0.001213	(8,12) (8,13) (10,15) (11,15)
0.001811	(8,14) (9,15)
0.081286	(10,12) (10,13) (11,12) (11,13)
0.121322	(9,12) (9,13) (10,14) (11,14)
0.181077	(9,14)

TABLE I

WEIGHTS FOR MINIMUM VARIANCE ESTIMATOR. Topology of Figure 2. Links 8 and 15 have ten times variance of others.

A.2 Variance Reduction

In Figure 4, convergence is tighter for the minimum variance estimator (on the right) than in the uniform case; this is particularly apparent in the left region of each plot, corresponding to smaller numbers of probes. The difference is particularly evident for link 1 (which has 2 high variance links as descendants, 8 and 15) and link 3, which has link 15 as a descendant. The variance of the estimators W_k for both these links is decreased in the minimum variance estimator, relative to the uniform estimator,

by reducing the weight μ_{ij} when i or j corresponds to a high variance link. This is particularly striking in the minimum variance estimator for link 1; we tabulate the weights $\mu_{ij}(C(1))$ in Table I. The weight for the pair (8, 15) of high variance links is 10^{-4} times the highest weight, that for pair (9, 14).

To see the statistics of estimator variation reduction, we display in Figure 5 the ratio of the standard deviation of the uniform estimator to the standard deviation of the minimum variance estimator, and a function of the number of probes. This is shown on the left for the cumulative variance, and on the right for the link delay variance. For the cumulative variance we display only for links, 1, 2 and 3; the other internal links the uniform and minimum variance estimators are identical because there is only one term in the sum for V . For the link variance we display only for links 1–7; the other links terminate at receivers and so have only one term in both V_k and $V_{f(k)}$. The figures show that the reduction in variance is roughly uniform across a range of experiment length up to 10,000 probes. The standard deviation was roughly halved for the cumulative delay variance, and between 0.3 and 0.5 for the link delay variance. Reduction was somewhat greater




Fig. 5. VARIANCE REDUCTION Ratio of standard deviation the uniform estimator to standard deviation minimum variance estimator. LEFT: cumulative delay variance, RIGHT: link delay variance.

for the standard deviation of the link delay variance, except for nodes 4 and 7. These nodes have only two descendants, one of which terminates a high variance link; there is no flexibility to avoid the high variance of the first term of $W_k = V_k - V_{f(k)}$.

B. Network Simulation

B.1 Methodology

The ns simulations used the topology in Figure 2(right). We arranged for some heterogeneity between the edges and the center of the tree in order to mimic the difference between the core and edges of a large WAN, with the interior of the tree having higher capacity (5Mb/sec) and latency (50ms) than at the edge (1Mb/sec and 10ms). Each node had a finite buffer capacity; packet losses were due to drops for the tail of the buffer. We used buffer capacities of 4 and 20 packets in two different sets of experiments. The cross traffic comprised 66 FTP sessions over TCP, and 29 UDP traffic sources following an exponential on-off model; there were on average around 8 background traffic sources per link. In each simulation we used the source-to-leaf delays of probes as data to infer delay variance per internal link by and also from the source to a given internal node. Since the simulations exhibit packet loss, the inference was performed using the algorithms described in Section V. We compared the inferred values W_k with the actual delay variance for probes on internal links that was observed during the simulation run. The comparison was performed over each link in Figure 2(right) for 100 simulation runs.

B.2 Comparison of Inferred and Actual Delay Variance

Figure 6 shows scatter plots of 1200 pairs of (inferred, actual) link delay variance, based on 1000 probes, on the left with buffer capacity of 4 packets, on the right with buffer capacity 20 packets. Also shown is the line through the origin at gradient 1; a point on this line would indicate an instance of perfect inference. In the scatter plots we differentiate between predictions using the uniform estimator, and those using the minimum variance estimator.

Taking each plot separately we observe that inference is

slightly more accurate for the minimum variance estimator than the uniform estimator, the difference being more evident for the smaller buffer size. Comparing the plots we see that inference is more accurate when for the simulated network with larger buffer capacities, particularly for small delay variances. A small number of inferred values were negative; this occurred for some links of high bandwidth for which queueing delays were small. Estimation of the link delay variance as the difference between the variance of the cumulative delays (see (15)) is sensitive to estimation errors. Nevertheless, the estimation error is sufficiently small that it would not impair identification of those links with the largest delay variance. Furthermore, in practice we can avoid the worst small variance estimation errors by eliminating estimates that are not significantly different from zero according at some confidence level. Similar to (17), these are the estimates W_k from n probes for which $W_k < z_\delta \sigma / \sqrt{n}$, the quantity $1 - \delta$ is the desired (one-sided) confidence level, and σ^2 is the appropriate asymptotic variance expressed in terms of the estimated covariance.

We attribute bias of inference to departures of the delay process from the independence assumption of the model. We calculated the off-diagonal elements of the correlation matrix of the actual link delays. For buffer size 4 the mean value was 0.071, the maximum 0.51. For buffer size 20 the mean was 0.021, the maximum 0.17. Thus correlations were more pronounced for the smaller buffer size, leading to greater violation of the inference assumptions, and evidently great inference inaccuracy. We found that bias was more pronounced in the inference of cumulative delay, see Figure 7. Bias was less evident for the link delay variance. Since this is expressed as a difference of estimated cumulative delay variance, consistent bias in the latter quantities should cancel somewhat in subtraction. Conversely, small delay variances are better estimated for for the cumulative than the link case.

In order to quantify the accuracy of inference we define a metric for evaluating estimator accuracy. If w and \hat{w} are the actual and inferred delay variances (either cumulative to a link or at the

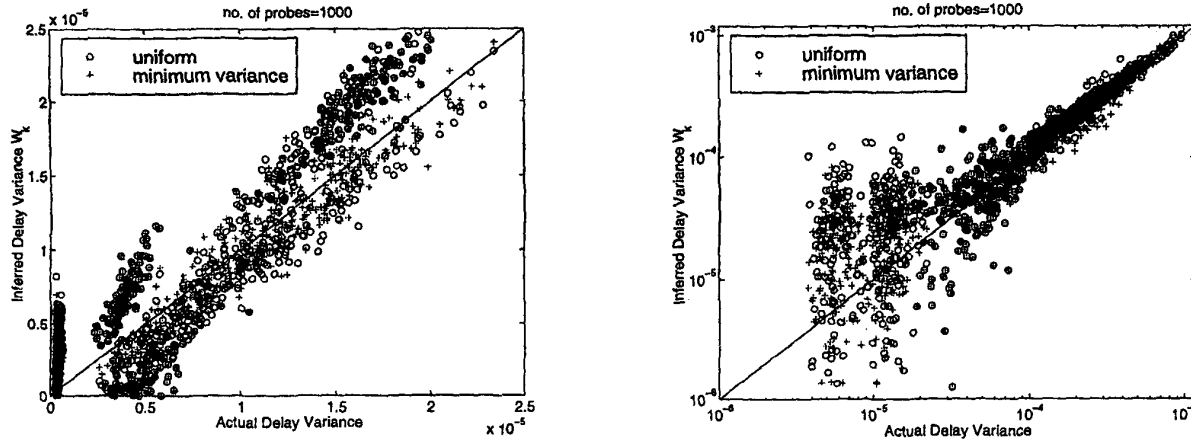


Fig. 6. SCATTER PLOTS FOR LINK DELAY INFERENCE. Pairs of (actual,inferred) over 12 links in 100 experiments, each using 1000 probes LEFT: nodes have 4 packet buffers; RIGHT: nodes have 20 packet buffers.




Fig. 7. SCATTER PLOTS FOR CUMULATIVE DELAY INFERENCE. Pairs of (actual,inferred) over 12 links in 100 experiments. LEFT: nodes have 4 packet buffers; RIGHT: nodes have 20 packet buffers.

link itself) we form their *error factor*

$$F(\hat{w}, w) = \max \left\{ \frac{\hat{w}}{w}, \frac{w}{\hat{w}} \right\}. \quad (20)$$

For example, if \hat{w} is either twice or half w , their error factor is 2. As a robust summary statistic to capture the center of the distribution of error factors, we use the two-sided quartile-weighted median (QWM)

$$(Q_{.25} + 2Q_{.5} + Q_{.75})/4, \quad (21)$$

where Q_p denotes the p^{th} quantile of a given set of error factors.

In Table II we display the QWM of error factors for link variance estimation. Small or negative inferred variances were omitted, the quantity z being the number of standard deviations characterizing the confidence interval about 0. $z = 0$ corresponds to rejecting only negative inferred variances. Ruling out these small variances decreases the QWM of the error factor: the smaller variances typically have higher error factor. (For $z = 2$, buffer = 4, it happens that the 75th-percentile of the error factor

z	buffer = 4 pkts.		buffer = 20 pkts.	
	unif.	min. var.	unif.	min. var.
0	2.31	2.06	1.32	1.32
2	1.56	1.76	1.23	1.23

TABLE II
QUARTILE WEIGHTED MEDIAN ERROR FACTORS FOR INFERENCE ON 1000 PACKETS. Link delay variance estimation, according to number of standard deviations z in confidence level to avoid small variances. Errors are smaller for minimum variance estimator than uniform estimator, and also with increased buffer capacity.

distribution is larger for the minimum variance estimator, but for other nearby quantiles it is smaller). For large buffer sizes the error factors are noticeably smaller for the minimum variance estimator; the difference in accuracy between the uniform and minimum variance estimator is smaller too. We found no great advantage in increasing the number of probes to 10,000 since bias becomes a larger relative part of the errors.

VII. CONCLUSIONS AND FURTHER WORK

In this paper we have proposed a novel technique for the inference from end-to-end measurements of the variance of the delay encountered by multicast packets on an internal link. The cooperation of intervening network nodes is not required.

We constructed a convex family of variance estimators and found the estimator of minimal asymptotic variance. Evaluating the minimal variance estimator comes at some computational cost, namely, the inversion of the covariance matrix \hat{C} . In work to be reported elsewhere, we show how this computation may be considerably simplified for binary trees, although at the cost of increasing estimator variance somewhat. Another approach is to compromise between the computational simplicity of the uniform estimator and variance reduction. An example would be to set $\mu_{ij} = 0$ for $\{i, j\}$ in some subset of $Q(k)$ in which the measures end-to-end variances \hat{s}_i are high. It remains to develop a robust approach along these lines.

The ns experiments showed, in the best cases, typical errors of about 20% in estimation of the delay variance using 1,000 probes. We observe that using a 40 bytes probe every 100ms, the load on the network is less than 4kb/sec and the measurements can be completed within 2 minutes.

We found inference to be more accurate in networks with larger buffers; there was smaller correlation between delays at different nodes and hence closer conformance to the underlying model. It appears that the larger buffers admit a greater diversity of connections through a node over queueing timescales, diluting the correlation seen between delays at successive nodes. We believe that diversity of traffic in real networks such as the Internet makes large and long lasting correlations unlikely. Furthermore the introduction of Random Early Detection (RED) [11] policies in Internet routers may help reduce dependence; evidence for this comes from related work on internal link loss inference [4], where the introduction of RED was found to increase accuracy of inference relative to networks with a Drop from Tail packet discard mechanism.

Acknowledgment

We thank Joseph Horowitz for his suggestion to find the minimum variance estimator.

REFERENCES

- [1] J. Bolot, "Characterizing End-to-End Packet Delay and Loss in the Internet." *Journal of High-Speed Network*, vol. 2 n. 3, pp. 289-298, Dec. 1993.
- [2] J.-C. Bolot and A. Vega Garcia "The case for FEC-based error control for packet audio in the Internet" to appear in *ACM Multimedia Systems*.
- [3] R. Caceres, N.G. Duffield, J. Horowitz and D. Towsley, "Multicast-Based Inference of Network Internal Loss Characteristics" to appear in *IEEE Trans. of Information Theory*.
- [4] R. Caceres, N.G. Duffield, J. Horowitz, D. Towsley and T. Bu, "Multicast-Based Inference of Network Internal Loss Characteristics: Accuracy of Packet Estimation" *Proc. of Infocom '99*, New York, NY, Mar. 1999.
- [5] R. Caceres, N.G. Duffield, J. Horowitz, F. Lo Presti and D. Towsley, "Loss-Based Inference of Multicast Network Topology" *IEEE Conference on Decision and Control*, 1999, to appear.
- [6] CAIDA: Cooperative Association for Internet Data Analysis. For more information see <http://www.caida.org>
- [7] K. Claffy, G. Polyzos and H.-W. Braun, "Measurements Considerations for Assessing Unidirectional Latencies", *Internetworking: Research and Experience*, Vol. 4, no. 3, pp. 121-132, Sept. 1993.
- [8] R. L. Carter and M. E. Crovella, "Measuring Bottleneck Link Speed in Packet-Switched Networks," *PERFORMANCE '96*, Oct. 1996.
- [9] A. Downey, "Using pathchar to estimate Internet link characteristics, On Proceedings ACM SIGCOMM'99, Cambridge, MA.
- [10] Felix: Independent Monitoring for Network Survivability. For more information see <ftp://ftp.bellcore.com/pub/mwgf/felix/index.html>
- [11] S. Floyd and V. Jacobson, "Random Early Detection Gateways for Congestion Avoidance," *IEEE/ACM Transactions on Networking*, Vol. 1, no. 4, August 1993.
- [12] IPMA: Internet Performance Measurement and Analysis. For more information see <http://www.merit.edu/ipma>
- [13] V. Jacobson, Pathchar - A Tool to Infer Characteristics of Internet paths. For more information see <ftp://ftp.ee.lbl.gov/pathchar>
- [14] F. Lo Presti and N.G. Duffield, "Multicast-Based Inference of Network-Internal Delay Distributions", Preprint AT&T Laboratories and University of Massachusetts.
- [15] J. Mahdavi, V. Paxson, A. Adams, M. Mathis, "Creating a Scalable Architecture for Internet Measurement," to appear in *Proc. INET '98*.
- [16] M. Mathis and J. Mahdavi, "Diagnosing Internet Congestion with a Transport Layer Performance Tool," *Proc. INET '96*, Montreal, June 1996.
- [17] D. Mills, "Network Time Protocol (Version 3): Specification, Implementation and Analysis", *RFC 1305*, Network Information Center, SRI International, Menlo Park, CA, Mar. 1992.
- [18] S. Moon, P. Skelly and D. Towsley, "Estimation and Removal of Clock Skew from Network Delay Measurements" *Proc. of Infocom '99*, New York, NY, Mar. 1999.
- [19] mtrace - Print multicast path from a source to a receiver. For more information see <ftp://ftp.parc.xerox.com/pub/net-research/ipmulti>
- [20] R.J. Muirhead, "Aspects of Multivariate Statistical Theory", Wiley, New York, 1982.
- [21] A. Mukherjee, "On the Dynamics and Significance of Low Frequency Components of Internet Load", *Internetworking: Research and Experience*, Vol. 5, pp. 163-205, Dec. 1994.
- [22] ns - Network Simulator. For more information see <http://www-mash.cs.berkeley.edu/ns/ns.html>
- [23] V. Paxson, "End-to-End Routing Behavior in the Internet," *Proc. SIGCOMM '96*, Stanford, Aug. 1996.
- [24] V. Paxson, "End-to-End Internet Packet Dynamics," *Proc. SIGCOMM 1997*, Cannes, France, pp. 139-152, Sept. 1997.
- [25] V. Paxson, "Measurements and Analysis of End-to-End Internet Dynamics," Ph.D. Dissertation, University of California, Berkeley, Apr. 1997.
- [26] V. Paxson, "Automated Packet Trace Analysis of TCP Implementations," *Proc. SIGCOMM 1997*, Cannes, France, 167-179, Sept. 1997.
- [27] V. Paxson, "On calibrating measurements of Packet Transit Times", *Proc. of SIGMETRICS '98*, Madison, June 1998.
- [28] S. Ratnasamy and S. McCanne, "Inference of Multicast Routing Tree Topologies and Bottleneck Bandwidths using End-to-end Measurements", *Proceedings IEEE Infocom '99*, New York, NY, Mar. 1999.
- [29] M.J. Schervish, "Theory of Statistics", Springer, New York, 1995.
- [30] Surveyor. For more information see <http://io.advanced.org/surveyor/>