

Unicast-Based Inference of Network Link Delay Distributions With Finite Mixture Models

Meng-Fu Shih, *Student Member, IEEE*, and Alfred O. Hero, III, *Fellow, IEEE*

Abstract—Providers of high quality-of-service over telecommunication networks require accurate methods for remote measurement of link-level performance. Recent research in network tomography has demonstrated that it is possible to estimate internal link characteristics, e.g., link delays and packet losses, using unicast probing schemes in which probes are exchanged between several pairs of sites in the network. In this paper, we present a new method for estimation of internal link delay distributions using the end-to-end packet pair delay statistics gathered by back-to-back packet-pair unicast probes. Our method is based on a variant of the penalized maximum likelihood expectation-maximization (PML-EM) algorithm applied to an additive finite mixture model for the link delay probability density functions. The mixture model incorporates a combination of discrete and continuous components, and we use a minimum message length (MML) penalty for selection of model order. We present results of matlab and ns — 2 simulations to illustrate the promise of our network tomography algorithm for light cross-traffic scenarios.

Index Terms—EM algorithm, mixture models, MML penalties, network tomography, signal processing in networking.

I. INTRODUCTION

HEREIN, we address the problem of determining internal link-delay distributions from multiple end-to-end unicast packet probes that are sent across a network of sensors, routers, or terminals from many different edge node pairs. This is a sub-problem of network tomography [1], [2], which is concerned with reconstruction of unobserved states of a network from a set of indirect measurements of these states. For example, end-to-end network tomography allows a few agents at the edge of a network to gain important information on global network behavior without cooperation of internal nodes. This is especially useful when link parameters are inaccessible or when direct measurement of data traffic statistics are not supported by internal switches and routers [1]–[5]. The unicast tomography methods we present in this paper are applicable to this situation and are designed to perform well when the network is lightly loaded. While modification of our methods may also be applicable to wireless networks, here, we focus on the simpler case of wired networks, such as the Internet.

The causes of delays along a packet probe's path through the network can be separated into the sum of two types of delays:

constant link transmission delays and time-varying link processing delays. Link transmission delays are due to the propagation delays through the physical medium, e.g., a wire, or optical fiber. Link processing delays are due to a combination of router queueing, buffering and servicing delays that depend on factors such as the amount of cross-traffic at the router, the number of retransmits required over the link, and the integrity of router equipment and associated software. While transmission delays usually remain constant over a probing interval, processing delays are highly variable and are thus commonly modeled as random variables. Thus, it is generally impossible to recover the actual internal link delays that probes encounter along their end-to-end path. However, the determination of the statistical distribution of the internal link delays from multiple end-to-end delay measurements can be formulated as a statistical inverse problem whose solution yields estimates of the internal delay distribution [2], [6]. These estimates can be used by an autonomous system (AS), e.g., an Internet service provider (ISP), to evaluate its average quality of service (QoS) or to assess link performance of other, perhaps competing, ASs. When acquired over large portions of the network, link delay estimates can also be used for detecting network anomalies such as imminent link failures or coordinated denial of service (DoS) attacks.

The problem of empirically characterizing Internet link delay distributions has been looked at by several researchers; see, for example, [7]–[10]. A common observation is that when the link is lightly loaded, such as in the early morning, link delay scatterplots appear stationary. Furthermore, while much of the scatter appears spread out over a continuum of delay values, a non-negligible proportion of the delays appear to concentrate at one or more discrete values; see, for example, [10, Fig. 4]. This implies the existence of point masses in the time-averaged link delay distribution. The positions of these point masses vary according to factors such as length of packet, incoming and outgoing queue sizes of routers on the link, router configuration, deployment of firewalls, and the physical distance between routers [10].

In this paper, we propose to capture these empirically observed features by fitting hybrid continuous/discrete finite mixture models to the link delay distributions. While our algorithms are easily generalizable to multiple discrete point masses, for simplicity, we focus here on the case where the discrete component is a single point mass. Unlike purely continuous models the hybrid continuous/discrete model is identifiable and is justified under the lightly loaded scenario. In this scenario, there is a nonzero probability that a packet will encounter an empty queue in which case the packet delay is nonrandom due to fixed propagation and processing delays. While this is unlikely in a congested network, the model is valid for a number of common

Manuscript received October 7, 2002; revised April 7, 2003. This work was supported in part by a Department of EECS Graduate Fellowship at the University of Michigan and by ARO-DARPA MURI Grant DAAD19-02-1-0262. The associate editor coordinating the review of this paper and approving it for publication was Prof. Ilkka Norros.

The authors are with the Department of Electrical and Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109 USA (e-mail: mfshih@umich.edu; hero@eeecs.umich.edu).

Digital Object Identifier 10.1109/TSP.2003.814468

monitoring situations such as service and performance verification and detection of onset congestion. Moreover, we would like to point out that the delay point mass is implicit in canonical delay trees, which are used in discrete delay tomography, for which there is a nonzero probability that a packet traverses each link without any delay (see, e.g., [6]).

Lo Presti *et al.* [6] were the first to propose estimating internal delay distributions from end-to-end measurements. These authors uniformly discretized the internal delay distributions and derived an algorithm based on empirical histogram estimation. Their method uses multicast probes, which requires cooperation of the network to run a multicast session such as real-time transport protocol (RTP) during the probing interval. To overcome this restriction, Coates and Nowak [5] developed an internal delay histogram estimator based on an alternative unicast probing scheme in which edge sites exchange a succession of closely spaced packet pairs. Their estimator is based on a statistical inverse problem formulation and used an iterative maximum likelihood via the expectation maximization (ML-EM) approach. In related work, these authors also developed a sequential Monte Carlo method for tracking changes in nonstationary networks [11]. The principal restriction of the approaches [5], [6], [11] is the requirement of discrete-valued link delays. Overly coarse discretization, or binning, of the link delays leads to excessive model approximation error and causes bias in derived estimates such as delay mean and variance. At the opposite extreme, excessively fine discretization leads to high runtime complexity of these algorithms. Furthermore, the determination of the appropriate number and size of the bins requires tight bounds on link delay characteristics, such as maximum and minimum processing delays, which are usually unknown.

Several alternatives to the fixed and uniform binning scheme of Lo Presti *et al.* [6] have been studied. Duffield *et al.* considered a variable bin size model, where smaller bins are used to describe probability mass concentrations for small delays [12]. Tseng *et al.* [13] proposed a nonparametric algorithm where the number of bins for internal link delays is adapted to the number of measurements. (See the paper by these authors in this Special Issue [18].) They use a wavelet-based penalized maximum likelihood estimator to smooth the estimates. A nonparametric method was proposed for unicast probing by Shih and Hero [14] in which the statistical inverse problem is formulated in the cumulant generating function (CGF) domain. By sampling the CGFs arising from the least squares solution of the inverse problem, a set of continuous (un-binned) link delay density estimates can be obtained.

Herein, we propose a different method for estimation of internal delay distributions from unicast end-to-end measurements that is based on packet pair unicast probes and additive mixture models for the internal link delays. As the end-to-end delay measurement is a sum of the (assumed independent) internal link delays over the probe path, the densities of the measurements are convolutive mixtures of these additive mixture models. This makes our estimation problem more challenging than the standard mixture model estimation problem, which has received much attention in both the statistical and engineering literature [15], [27], [29], [30].

Additional issues we address are 1) the additive mixture model orders are unknown in practice, and 2) the internal link delay distributions are composed of a combination of continuous and a discrete components. We handle the convolutive mixture complication by adopting an iterative ML-EM formulation of the estimation problem using an enlarged complete data space. We handle the problem of unknown model order by adapting the unsupervised minimum-message-length (MML) approach used in Figueiredo and Jain [15]. Specifically, we add an information theoretic order selection penalty to the log-likelihood to which a penalized ML-EM (PML-EM) algorithm is applied. We handle the presence of both discrete and continuous link delay components by the following simple additive mixture model: The delay density is a (unknown) convex combination of a point mass positioned at the (unknown) transmission delay and a (unknown) number of Gaussian components with (unknown) means and variances. We adopted Gaussian continuous components to simplify the implementation, but heavy-tailed densities can also be easily accommodated in our framework.

The outline of the paper is as follows. In Section II, we give the main assumptions underlying our work. In Section III, we review the continuous and discrete delay models for unicast network delay tomography and discuss model identifiability. In Section IV, we introduce the hybrid mixture models for delays of the internal links and present the ML-EM algorithm for estimating the parameters of these models from end-to-end delay measurements. In Section IV-C, we present the PML-EM algorithm with MML penalty to control excessive order estimation of the unpenalized ML-EM algorithm. In Section V, we illustrate the performance of the ML-EM and PML-EM algorithms on simulated data using `matlab` and `ns - 2` [16] simulators.

II. NETWORK MODEL AND MAIN ASSUMPTIONS

As in Coates and Nowak [5], we adopt the back-to-back packet pair probing framework and represent the network topology as a directed logical tree $T = (V, E)$, where V is the set of nodes, e.g., routers and terminals, and E is the set of links. Let there be a total of L links in the network and number them from 1 to L . The logical tree representation has a single root node, serving as a source, several internal nodes having degree at least 2, and several leaf nodes, containing receivers. The edge of the network is the set of leaf nodes plus the root node. If there are a total of R leaf nodes, then there are R possible paths from the root to the receivers. To collect internal link information at the edge nodes, pairs of packets are transmitted from the source to pairs of leaf nodes. There are $S = C_2^R$ binary sub-trees of paths, called probe trees, on which the source can send each pair of packets to two different receivers. We define the pair of receivers at the edge of probe tree i to be W_i . Each node in the tree has two or more child nodes, except the root and leaf nodes. A network with three links is shown in Fig. 1(a) showing a single root node, a single internal node, and two leaf nodes.

In a unicast probing session, a pair of leaf nodes is (randomly) selected by the source, and two time stamped packets, called a (unicast) probe pair, are sent to them, respectively. The two packets are transmitted in rapid succession and encounter identical delays on the shared links of their paths. Each leaf

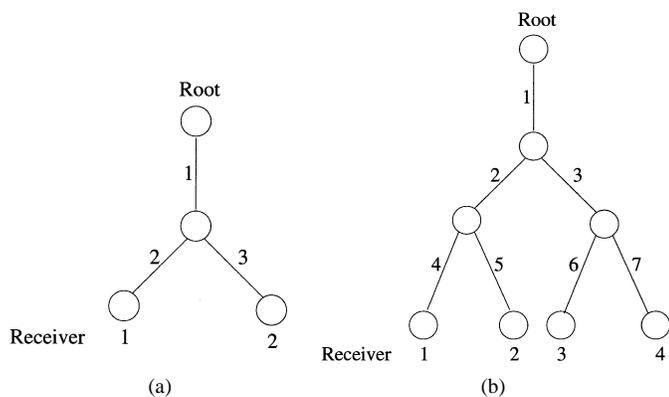


Fig. 1. (a) Two-leaf and (b) four-leaf networks, respectively. Receivers are the leaf nodes while the source is the root node. Network (b) is used for simulations in Section V.

node records the time that a packet is received. Subtracting this number from the packet's time stamp gives the end-to-end delay of the packet. End-to-end delays of the probe pairs on the same probe tree are random vectors due to the random ambient cross traffic through links along their paths. If any packet in a probe pair is dropped by the network, both packets are considered lost. Unicast probing is repeated until the session is over or enough packets are received by each leaf node to perform the next step: network delay tomography. The aim of network delay tomography is to identify the packet delay distribution for each individual internal link from the end-to-end delays observed by the receivers. Network tomography is possible since the end-to-end delay is a sum of the internal link delays encountered along the probe path and any two paths in a probe tree must cross at common links.

Let X_l be the packet delay encountered by a probe at link l , $l = 1, \dots, L$, and let Y_i be the end-to-end packet delay along the i th path $i = 1, \dots, R$. We make the following independence and stationarity assumptions.

- A1) *Spatial Independence*: Packet delays at different links are statistically independent, i.e., X_i and X_j are independent for $i \neq j$.
- A2) *Temporal Independence and Stationarity*: For a given link, the delays encountered by packets in different probe pairs at that link are statistically independent and identically distributed (i.i.d.).

For each probe pair, we make an additional consistency assumption.

- A3) The delays encountered by both packets in a probe pair on the shared links of their paths are the same (with probability 1).

It is important to point out that while A1) and A2) are normally not satisfied in practice (see, e.g., [17]), these are commonly assumed in order to permit tractable analysis. An example where spatial independence A1) is violated is when there is interaction among different data flows along the same path. As for A2), temporal independence fails when Internet traffic is bursty or the network has a long latency time, which correlates different packet pairs. Stationarity fails when the unicast probing session has a longer duration than the stationarity time of the network. However, experiments have shown that the performance of network tomography is remarkably insensitive to violations of A1)

and A2) [4]–[6], [11], [20]. In A3), the assumption of identical delays experienced by a probe pair on shared links does not hold when a small discrepancy between the two is observed from real network data (see, e.g., [18], [19]). Fortunately, this random error has mean close to 0 and can be reduced by random ordering of the two packets [5].

III. UNICAST NETWORK DELAY TOMOGRAPHY

A. Discrete Delay Model

In the widely adopted discrete link delay model [5], [6], [11], a universal bin size q is used to discretize link delays X_l at each link $l = 1, \dots, L$. The time intervals $(iq, (i+1)q]$, $i = 0, \dots, D$ are called the delay bins. Here, D is a positive integer, and $D = \infty$ can be used to account for lost probe packets or large delays that are out of range. Discretization produces the discretized delay value i when X_l falls in the i th bin. A probability mass function (pmf) or histogram $P_l = \{p_{l,d} : d = 0, \dots, D\}$ is then associated with the discretized delays over link l , where the probability $p_{l,d} = P(X_l \in (dq, (d+1)q])$ is an unknown to be estimated, and $\sum_{d=0}^D p_{l,d} = 1$. For a probe path containing j links, the discretized end-to-end packet delay varies over the range $0, \dots, j \cdot D \cdot q$.

Consider the two-leaf tree network shown in Fig. 1(a), and the associated delay pmf's $P_l = \{p_{l,d} : d = 0, \dots, D\}$ for $l = 1, 2, 3$. Probe pairs are sent from the source to receiver 1 and 2. With assumption A3), the identifiability of P_l s from end-to-end delays can be studied in a similar manner to multicast networks. More specifically, in multicast, each packet is replicated by the network at the branching points of its paths, and all the packets at the receivers again have common delays on shared links. Proof of identifiability in discrete network delay tomography with multicast probes is provided in [6], and the use of unicast probe pairs can be considered as a special case.

The discrete delay model adopted in [5], [6], [11]–[13], and [18] has two main drawbacks. First, the proper bin size needs to be carefully selected. Second, a universal bin size may not be suitable due to large variation of packet delay ranges over different links. Although in [6] it was proposed to adopt different bin sizes for different links, those bin sizes still need to be chosen in advance.

B. Continuous Delay Model

One way to avoid the pitfalls of binning is to use a flexible continuous link delay model. For example, closed-form expressions for the probability density function (pdf) of queueing delay have been derived for simple queueing models such as M/M/1. These expressions could possibly be extended to a network of queues, but the accuracy of the M/M/1 model is still in question for Internet traffic [21]. An alternative is to approximate each link delay density by a finite mixture that, with sufficiently large number of components, can describe any continuous density function [22]. Let $f_l(x)$ be the link delay pdf at link l . A finite mixture model for this pdf is

$$f_l(x) = \sum_{m=1}^{k_l} \alpha_{l,m} \phi(x; \theta_{l,m}), \quad l = 1, \dots, L \quad (1)$$

where k_l denotes the number of components, $\alpha_{l,m}$, $m = 1, \dots, k_l$ denotes the mixing parameter for the m th component ($0 < \alpha_{l,m} < 1$, $\sum_{m=1}^{k_l} \alpha_{l,m} = 1$), and $\phi(x; \theta_{l,m})$ is a density function over the x -domain parameterized by the parameter vector $\theta_{l,m}$. Many different choices for $\phi(x; \theta)$ are possible including Gaussian, generalized Gaussian, exponential, or uniform. For the case of a Gaussian mixture, $\theta_{l,m} = [\mu_{l,m}, \sigma_{l,m}^2]$ is the vector specifying the position (mean) and width [sqrt(variance)] of the m th mixture component at the l th link.

However, the use of pure continuous mixture density functions can cause serious identifiability problems. To illustrate, consider again the simple two-leaf tree of Fig. 1. Assume that all link delays are Gaussian, i.e., $k_1, k_2, k_3 = 1$ (single component mixtures) $\phi(x; \theta) = \exp(-(x - \mu)^2 / (2\sigma^2)) / (\sqrt{2\pi}\sigma)$. The end-to-end delays (Y_1, Y_2) have the following joint pdf:

$$f(Y_1, Y_2) = \frac{1}{2\pi\sqrt{\sigma_1^2\sigma_2^2 + \sigma_1^2\sigma_3^2 + \sigma_2^2\sigma_3^2}} \cdot \exp\left\{-\frac{1}{2(\sigma_1^2\sigma_2^2 + \sigma_1^2\sigma_3^2 + \sigma_2^2\sigma_3^2)} \cdot (\sigma_3^2(Y_1 - (\mu_1 + \mu_2))^2 + \sigma_2^2(Y_2 - (\mu_1 + \mu_3))^2 + \sigma_1^2((Y_1 - Y_2) - (\mu_2 - \mu_3))^2)\right\}. \quad (2)$$

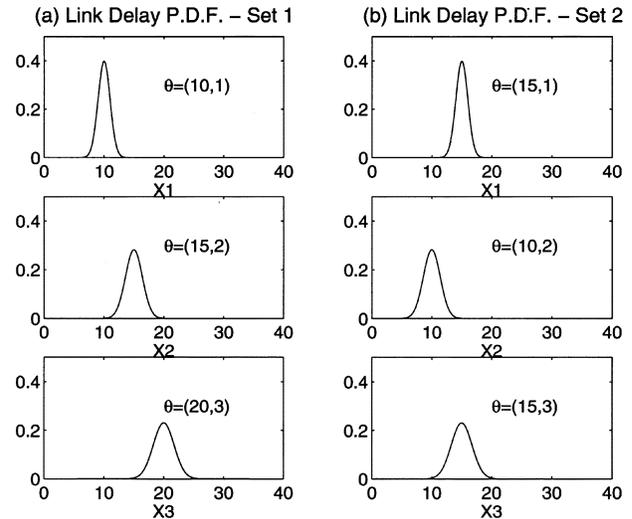
If we look at the mean parameters, they are completely described by the three parameters $\eta_1 = \mu_0 + \mu_1$, $\eta_2 = \mu_0 + \mu_2$, and $\eta_3 = \mu_1 - \mu_2$. This gives only two equations for the three unknown parameters $[\mu_1, \mu_2, \mu_3]$; therefore, the simple Gaussian model is not identifiable for any value of the mean parameters. An example is shown in Fig. 2, where (a) and (b) are two different sets of internal link delay distributions for the network in Fig. 1(a).

One can also consider the packet-stripe schemes suggested in [19], in which a “stripe” of several closely spaced unicast packets with distinct destinations are sent back-to-back from the root node. Similarly to packet pair probes, these packets are assumed to encounter virtually the same delays on shared links along their paths. As shown in [20], packet-stripe probing allows identification of higher order moments of internal link delays when the branching ratio is larger than two. However, under the Gaussian mixture link delay model, the link delay means still cannot be uniquely identified from end-to-end delays.

IV. HYBRID FINITE MIXTURE APPROACH

A. Hybrid Finite Mixture Model

In analysis of a queueing system, the utilization factor ρ is an important parameter for describing system behavior. The parameter ρ denotes the probability that the system is busy serving customers, and for a stable system, ρ must satisfy $0 \leq \rho < 1$ [23]. A lightly loaded link satisfies $\rho \ll 1$, i.e., there is a non-negligible probability that a packet encounters an empty queue, i.e., an idle router, and passes without delay. This suggests placing a point mass component with weight $1 - \rho$ in the link delay mixture model. If this point mass is included in addition to the continuous components, the link delay pdf becomes a



(c) End-to-End Delay P.D.F.

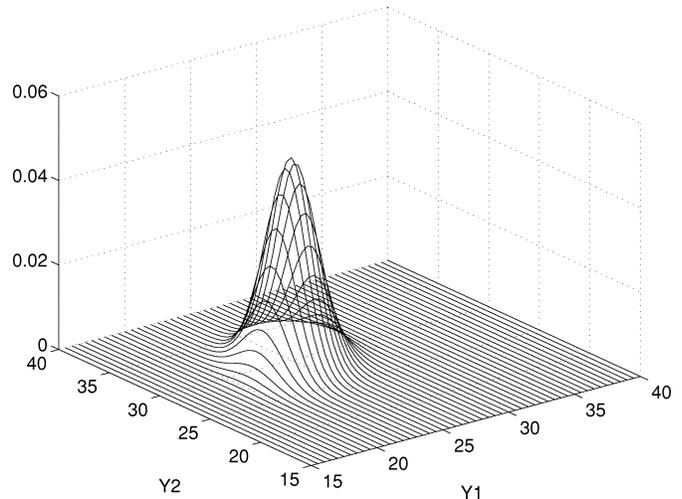


Fig. 2. Example of two sets (a) and (b) of Gaussian internal link delay densities along the two probe paths in the network in Fig. 1(a). The two end-to-end delays of each received packet pair obeys a Gaussian bivariate density shown in (c). This bivariate density is parameterized by only two location parameters, which is insufficient to recover the three location parameters in (a) and (b).

hybrid discrete/continuous finite mixture model. Hence, similar to (1), we obtain

$$f_l(x) = \alpha_{l,0}\delta(x - x_{l,0}) + \sum_{m=1}^{k_l} \alpha_{l,m}\phi(x; \theta_{l,m}). \quad (3)$$

Here, $\alpha_0 = 1 - \rho$, $\delta(x)$ is a point mass (dirac delta function) at zero, and $x_{l,0}$ is the pure (nonrandom) transmission delay experienced by the packet. All other parameters are defined as in (1), except now, the α s must satisfy $\sum_{m=0}^{k_l} \alpha_{l,m} = 1$, $\alpha_{l,m} \geq 0$. The discrete mass component $\delta(x)$ not only makes the delay distribution more precisely model the behavior of a link queue but, as shown below, also buys us identifiability of all the link delay distribution parameters.

For any probe pair, the distributions of the end-to-end probe delay densities will be the convolution of the link distributions, which are also hybrid mixtures. Now, similarly to the previous section, let us assume that the continuous mixture component

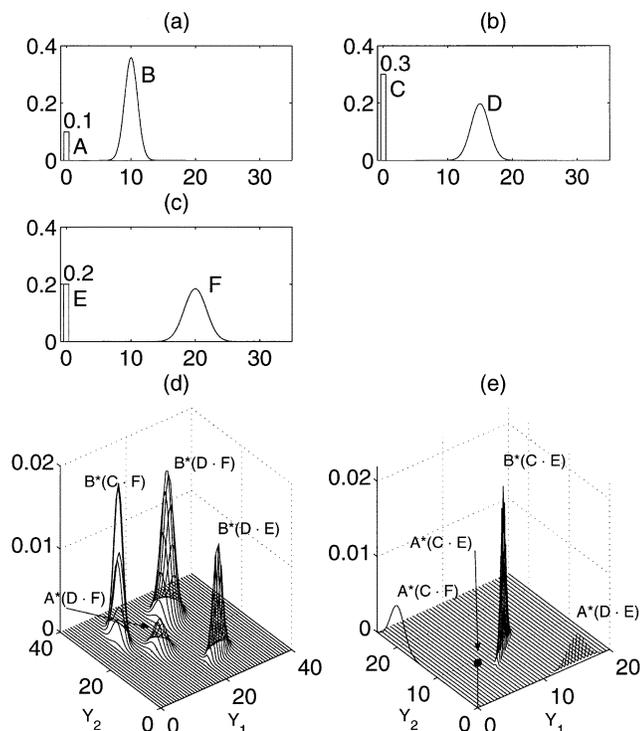


Fig. 3. Example of internal link delay mixture densities (a)–(c) for links 1, 2, 3, respectively, over the two-leaf tree of Fig. 1(a). The nonrandom minimum delays for all the links are set to 0. The end-to-end packet pair delay distribution is a hybrid mixture whose purely continuous components are shown in (d), and components associated with discrete masses are in (e). All link parameters can be identified from this two-dimensional distribution. Here, $B^*(C \cdot F)$ denotes a function of (y_1, y_2) , which is the convolution of the internal link components labeled B, C, and F in the form of $\int B(x)C(y_1 - x)F(y_2 - x) dx$.

is a single Gaussian pdf. Let the point masses $\alpha_{l,0} = \alpha_l$, and assume that they are all concentrated at zero delay, i.e., $x_{l,0} = 0$. Fig. 3 shows the end-to-end joint delay distribution in the two-leaf tree network of Fig. 1(a), whose mathematical form is

$$\begin{aligned}
 f(Y_1, Y_2) = & \beta_1\beta_2\beta_3\tilde{f}(Y_1, Y_2) + \alpha_1\alpha_2\alpha_3\delta(Y_1)\delta(Y_2) \\
 & + \alpha_1\alpha_2\beta_3\delta(Y_1)\phi(Y_2; \mu_3, \sigma_3^2) \\
 & + \alpha_1\beta_2\alpha_3\delta(Y_2)\phi(Y_1; \mu_2, \sigma_2^2) \\
 & + \beta_1\alpha_2\alpha_3\delta(Y_1 - Y_2)\phi(Y_1; \mu_1, \sigma_1^2) \\
 & + \alpha_1\beta_2\beta_3\phi(Y_1; \mu_2, \sigma_2^2)\phi(Y_2; \mu_3, \sigma_3^2) \\
 & + \beta_1\alpha_2\beta_3\phi(Y_1; \mu_1, \sigma_1^2)\phi(Y_2 - Y_1; \mu_3, \sigma_3^2) \\
 & + \beta_1\beta_2\alpha_3\phi(Y_2; \mu_1, \sigma_1^2)\phi(Y_1 - Y_2; \mu_2, \sigma_2^2)
 \end{aligned} \quad (4)$$

where $\beta_l = 1 - \alpha_l$ for $l = 1, 2, 3$, and $\tilde{f}(Y_1, Y_2)$ is the joint distribution shown in (2). Due to the point mass in (3), (4) has additional isolated Gaussian components that appear with discrete masses at locations in the (Y_1, Y_2) plane specified by $\{Y_1 = 0\}$, $\{Y_2 = 0\}$, and $\{Y_1 = Y_2\}$. It is obvious that identifiability can be achieved as long as $\alpha_l \neq 0$. It might seem strange to the reader that the addition of a point mass allows one to uniquely identify the set of parameters of the internal link components from a single probe tree. However, one still needs multiple probe trees to assign these parameters to specific links.

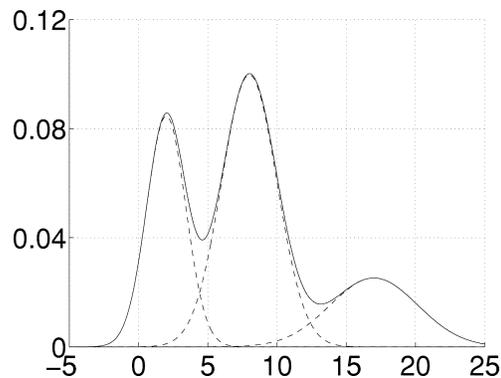


Fig. 4. Gaussian mixture example. $f_U(u) = 0.3N(u; 2, 2) + 0.5N(u; 8, 4) + 0.2N(u; 17, 10)$.

B. ML-EM Algorithm

Here, we present an ML-EM algorithm for approximating the maximum likelihood estimates of the internal link mixture model parameters from end-to-end packet pair measurements. Let U be a finite mixture random variable with k components and pdf of the form $f(U) = \sum_{m=1}^k \alpha_m \phi_m(U)$, where $\sum_{m=1}^k \alpha_m = 1$. An example of a Gaussian mixture with three components is given in Fig. 4. The solid line depicts the density function, and the dashed line shows each component. There are two different interpretations of finite mixture models, which will be useful in the sequel. The first one is simply that $f(U)$ is a multicomponent pdf for U . The second interpretation is that U is selected at random from a pool of simpler hidden random variables U_1, \dots, U_k with selection probabilities $\alpha_1, \dots, \alpha_k$, respectively. Define the binary random selection vector $\mathbf{Z} = [Z_1, \dots, Z_k]^T$, where $Z_m = 1$ if and only if the m th variable U_m is selected, and assign to this event probability α_m . U can be expressed as $U = \sum_{m=1}^k Z_m U_m$. Thus, if U_m has pdf $\phi_m(U_m)$, then this is identically the conditional pdf $f(U|Z_m = 1)$. Thus, $f(U) = \sum_{m=1}^k \alpha_m \phi_m(U)$, which is the mixture model for U with which we started out. The second interpretation is critical for development of the ML-EM algorithm, which we address below.

Assume that we have prior knowledge of all the link mixture orders $\{k_l\}_{l=1}^L$. We will relax this assumption in the next section. Let N_i be the number of packet pairs sent from the source to the receivers of probe tree i , and let M_i be the set of links along that tree. Define $X_l^{(i,n)}$ as the delay at link l encountered by the n th packet pair sent to receivers in W_i . Let $\mathbf{Z}_l^{(i,n)} = [Z_{l,0}^{(i,n)}, \dots, Z_{l,k_l}^{(i,n)}]$ be the selection vector for $X_l^{(i,n)}$.

With these definitions, maximum likelihood (ML) estimation of the set of internal link mixture densities can be formulated as a missing data problem. The expectation maximization (EM) algorithm has been extensively applied to approximate ML and penalized ML (PML) estimates for mixture models [15], [29], [30]. Let $\mathbf{X} = \{X_l^{(i,n)}\}$ and $\mathbf{Z} = \{\mathbf{Z}_l^{(i,n)}\}$ for all l, i, n . $\{\mathbf{X}, \mathbf{Z}\}$ is called *missing data* or *hidden data*. Define $\mathbf{Y}^{(i,n)} = (Y_1^{(i,n)}, Y_2^{(i,n)})$ as the pair of end-to-end delays of the n th packet pair received by two receivers in the i th probe tree. The observables $\mathbf{Y} = \{\mathbf{Y}^{(i,n)}\}_{i,n}$ are called the *incomplete data*, and the set $\{\mathbf{X}, \mathbf{Z}, \mathbf{Y}\}$ is said to be the *complete*

data. The EM algorithm generates a sequence of estimates of the unknown parameters Θ that have the property that the likelihood sequence $\mathcal{L}(\Theta) = f(\mathbf{Y}|\Theta)$ is nondecreasing.

It is easily shown that the likelihood of the *complete data* can be factorized as $\mathcal{L}_c(\Theta) \stackrel{\text{def}}{=} f(\mathbf{X}, \mathbf{Y}, \mathbf{Z}|\Theta) = f(\mathbf{Y}|\mathbf{X})f(\mathbf{X}, \mathbf{Z}|\Theta)$, and thus, maximization of $\mathcal{L}_c(\Theta)$ is equivalent to maximization of the likelihood function $\mathcal{L}(\Theta) \stackrel{\text{def}}{=} f(\mathbf{X}, \mathbf{Z}|\Theta)$. For a specific link l , X_l is a mixed random variable with density function f_l given by (3), and therefore, up to a constant, the complete data log-likelihood function is

$$\log \mathcal{L}(\Theta) = \sum_{l=1}^L \sum_{i: l \in M_i} \sum_{n=1}^{N_i} \left\{ Z_{l,0}^{(i,n)} \log \alpha_{l,0} + \sum_{m=1}^{k_l} Z_{l,m}^{(i,n)} \left(\log \alpha_{l,m} + \log \phi(X_l^{(i,n)}; \theta_{l,m}) \right) \right\}. \quad (5)$$

The EM algorithm updates parameter estimates by applying two steps at each iteration. At the t th iteration, the E-step computes conditional expectation of complete data log-likelihood given observations \mathbf{Y} and current parameter estimates $\hat{\Theta}^{(t)}$

$$Q(\Theta, \hat{\Theta}^{(t)}) = E \left[\log \mathcal{L}(\Theta) | \mathbf{Y}; \hat{\Theta}^{(t)} \right]. \quad (6)$$

The M-step maximizes the Q function computed in the E step with respect to Θ to produce

$$\hat{\Theta}^{(t+1)} = \underset{\Theta}{\operatorname{argmax}} Q(\Theta, \hat{\Theta}^{(t)}). \quad (7)$$

The E and M steps for the hybrid mixture model are similar to those for a single Gaussian mixture model [30] and are illustrated in the Appendix.

C. PML-EM Algorithm With MML Penalty

When the number of link components k_l s is unknown, the ML-EM algorithm is not guaranteed to converge. This is due to a fundamental ambiguity of unknown model order. To illustrate, consider the estimation of a k -component mixture having the form of (1) with parameters $\Theta = \{\alpha_1, \dots, \alpha_k, \theta_1, \dots, \theta_k\}$. These parameters have the same likelihood as the $k+1$ component mixture $\Theta' = \{\alpha_1, \dots, \alpha_{k-1}, (1-\beta)\alpha_k, \beta\alpha_k, \theta_1, \dots, \theta_k, \theta_k\}$ for any $0 < \beta < 1$. One of the most effective ways to eliminate this ambiguity is to add a penalty to the log-likelihood function that penalizes the addition of more components to the mixture.

Many model order penalties have been proposed including the Akaike information criterion (AIC) [24], minimum description length (MDL) [25], and minimum message length (MML) [26]. Figueiredo and Jain [15] applied the MML penalty to finite mixture models by introducing a prior to the parameters and an information-theoretic penalty, depending on quantization of parameter space. They developed an unsupervised method to si-

multaneously select model order and estimate parameters. The incomplete data penalized log-likelihood is expressed as

$$\tilde{\mathcal{L}}(\Theta) \triangleq \log f(\Theta) + \log f(\mathbf{Y}|\Theta) - \frac{1}{2} \log |\mathbf{I}(\Theta)| - \frac{c}{2} (1 + \log \kappa_c)$$

where $\mathbf{I}(\Theta)$ is the Fisher information matrix associated with the incomplete data \mathbf{Y} , $|\mathbf{A}|$ denotes the determinant of square matrix \mathbf{A} , c is the dimension of Θ , and κ_c is the so-called *optimal quantizing lattice constant* for \mathbb{R}^c .

To apply the MML algorithm [15] of Figueiredo and Jain to our network delay tomography problem, their method has to be extended to another layer of hidden data. More specifically, while in [15] the realizations from the mixture model were observed directly, in our application, only sums of these realizations (along probe paths) are observed. In other words, the end-to-end delays are themselves convolutive mixtures of the additive mixtures describing the link delays.

The standard incomplete data Fisher Information matrix $\mathbf{I}(\Theta)$ is not closed form, even for a directly observed finite mixture [27]. Therefore, similar to [15], we replace it by the complete data Fisher information matrix, which, in the network tomography setting, is $\tilde{\mathbf{I}}(\Theta) = -E[\nabla_{\Theta}^2 \log f(\mathbf{X}, \mathbf{Z}|\Theta)] = \text{block-diag}\{n_l \mathbf{I}_1(\Theta_l)\}_{l=1}^L$, where $\mathbf{I}_1(\Theta_l)$ is the Fisher information matrix associated with the complete data at link l , Θ_l denotes the parameter set of the l th link, and $n_l = \sum_{i: l \in M_i} N_i$ is the total number of packet pairs passing through the l th link. \mathbf{I}_1 itself has block-diagonal structure $\mathbf{I}_1(\Theta_l) = \text{block-diag}\{\mathbf{A}_l, \alpha_{l,1} \mathbf{I}_2(\theta_{l,1}), \dots, \alpha_{l,k_l} \mathbf{I}_2(\theta_{l,k_l})\}$, where $\mathbf{I}_2(\theta_{l,m})$ is the Fisher information matrix associated with the hidden m th component delay variable $X_{l,m}$ on link l , and $\mathbf{A}_l = \text{diag}\{\alpha_{l,m}^{-1}\}_{m=0}^{k_l}$. If any one of the $\alpha_{l,m}$ s is zero, it is removed from \mathbf{A}_l , and k_l is decreased by 1.

The prior on the parameter set was taken as $f(\Theta) = \prod_{l=1}^L \{f(\alpha_{l,0}, \dots, \alpha_{l,k_l}) \prod_{m=1}^{k_l} f(\theta_{l,m})\}$, where $f(\alpha_{l,0}, \dots, \alpha_{l,k_l})$, and $f(\theta_{l,m})$ are the noninformative Jeffreys' priors [28]: $f(\alpha_{l,0}, \dots, \alpha_{l,k_l}) \propto \sqrt{|\mathbf{A}|} = (\alpha_{l,0} \alpha_{l,1} \dots \alpha_{l,k_l})^{-1/2}$, and $f(\theta_{l,m}) \propto \sqrt{|\mathbf{I}_2(\theta_{l,m})|}$, for $\sum_{m=0}^{k_l} \alpha_{l,m} = 1$ and $0 < \alpha_{l,m} < 1$. In addition, as in [15], we make the approximation $\kappa_c = 1/12$. This yields the MML penalized likelihood function

$$\tilde{\mathcal{L}}(\Theta) = \log f(\mathbf{Y}|\Theta) - \frac{d}{2} \sum_{l=1}^L \sum_{m=1}^{k_l} \log \alpha_{l,m} - \sum_{l=1}^L \frac{k_l(d+1)+1}{2} \left(\log \frac{n_l}{12} + 1 \right) \quad (8)$$

where d is the dimension of $\theta_{l,m}$, e.g., $d = 2$ for a Gaussian component mixture.

To derive the E step of the PML-EM algorithm applied to maximizing (8), we adopt the same complete data as in the previous section. With this, it is easy to see that the E step is a modification of (6), where $Q(\Theta, \hat{\Theta}^{(t)})$ now has an additional penalty given by the second and third additive terms on the RHS of (8). The modified M-step gives the updates for the mixing parameters in (9) (see the Appendix). The M-step for the remaining

parameters depends on the specific form of the mixture density components.

$$\alpha_{l,m}^{(t+1)} = \frac{\max \left\{ \left(\sum_{i: l \in M_i} \sum_{n=1}^{N_i} \omega_{l,m}^{(i,n)} \right) - \frac{d}{2}, 0 \right\}}{\sum_{m=1}^{k_l} \max \left\{ \left(\sum_{i: l \in M_i} \sum_{n=1}^{N_i} \omega_{l,m}^{(i,n)} \right) - \frac{d}{2}, 0 \right\} + \sum_{i: l \in M_i} \sum_{n=1}^{N_i} \omega_{l,0}^{(i,n)}} \quad m = 1, \dots, k_l. \quad (9)$$

The algorithm uses the following strategy to select the number k_l of components at the l th link. It starts by setting all $k_l, l = 1, \dots, L$ to some user-specified upper bound and annihilates components as follows. If $\alpha_{l,m}^{(t+1)} = 0$, component m is removed from f_l , and its probability mass is redistributed over the other nonzero-probability components at the next iteration. Note that this procedure is myopic since it does not allow components to come back to life, and therefore, the estimate is not guaranteed to converge to the MML estimate. However, we restrict our implementation to this myopic strategy due to implementation complexity constraints.

V. EXPERIMENTAL RESULTS

A. Model Simulation: ML-EM for Known Model Order

We simulated a small network with the simple virtual tree topology shown in Fig. 1(b). Throughout this first experiment, the numbers of components $\{k_l\}_{l=1}^L$ are known to the estimator. We specialized the EM algorithm to a Gaussian continuous component mixture (see the Appendix). From two to four Gaussian components were assigned to each link in addition to a point mass. These simulations were implemented in `matlab`, and we generated 2000 i.i.d. end-to-end probe pair delays for each of the six probe tree paths. The ML-EM algorithm was applied to the estimation of the Gaussian components, their mixing parameters, and the weight of the point mass at zero. Convergence was achieved after 955 iterations, or approximately 16 iterations per parameter. Fig. 5 compares the estimated Gaussian mixture components to the true Gaussian mixture components. It also lists the number of mixture components for each link and the true/estimated probabilities $\alpha_{l,0}$ of the probe encountering empty queue on link l . The convergence curve of the log-likelihood is shown in the lowest-right graph of the figure. These results illustrate high accuracy for the case where there is no model error and the number of components is known.

B. NS Simulation: MML for Unknown Model Order

For a more realistic simulation, we used `ns-2` [16] to simulate the network shown in Fig. 1(b) with a variety of cross traffic types and router configurations. The links were assigned bandwidths and latencies listed in Table I. The `ns` parameters for each link were set to a Drop-Tail queue [first-in–first-out (FIFO) queue with finite buffer]. The queue buffer sizes were 50 packets long. Each packet in a probe pair was defined as a 40-byte user datagram protocol (UDP) packet. Probe pairs are generated independently and sent along each of the six tree paths according to a Poisson process with mean interarrival time 8 ms and rate

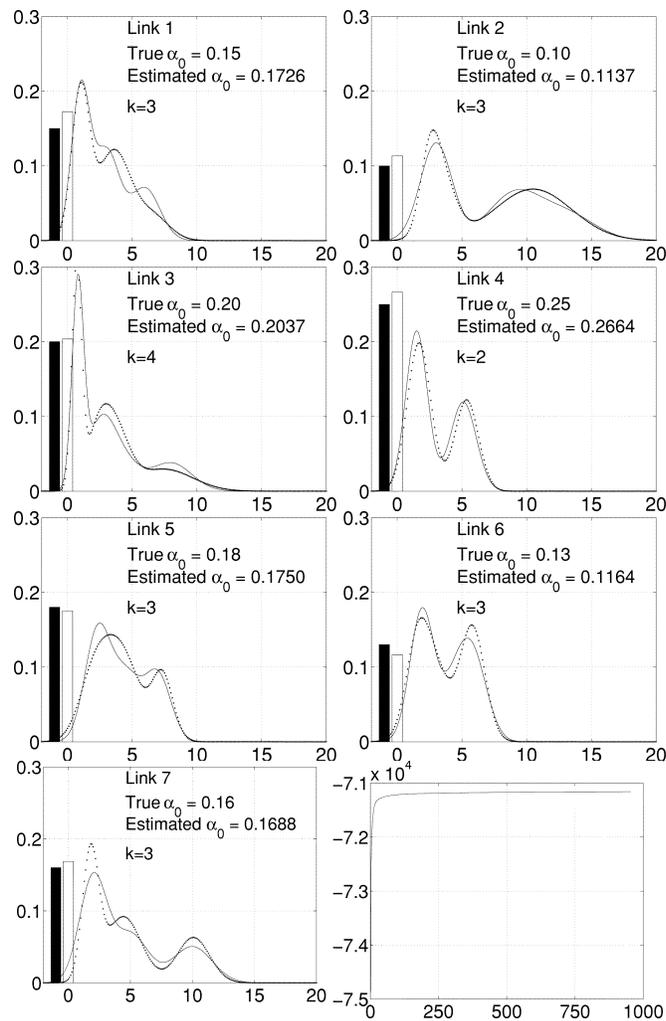


Fig. 5. True (solid curve) and estimated (dotted curve) Gaussian mixture components along with the true (black bar) and estimated (white bar) empty queue probabilities $\{\alpha_{l,0}\}$ for model simulation. The horizontal axes denote link packet delays in milliseconds. Here, the EM algorithm is used to estimate the mixed Gaussian mixture parameters for simulated measurements obeying a true mixed Gaussian mixture with known numbers of components, which are listed along with the link delay pdfs. Two thousand packet pairs are generated for each of the six probe tree paths in Fig. 1(b). The lowest-right graph shows the convergence curve of the log-likelihood function.

TABLE I
LINK BANDWIDTH AND LATENCY PARAMETERS USED IN `ns` SIMULATION

Link	1	2	3	4	5	6	7
Bandwidth (Mbps)	20	20	20	20	15	12	12
Latency (ms)	100	100	100	30	50	50	40

70 Kb/s. Cross traffic was also generated in each link by `ns` and consisted of 41 Pareto On–Off transfer control protocol (TCP) flows and 25 constant-bit-rate UDP streams with random noise introduced in the scheduled packet departure times. The design of background traffic reflects today’s IP network environment in which the UDP traffic is mainly video/audio data streams and TCP comprises the major fraction of the Internet traffic [31]. A total of $N = 500$ packet pairs for each probe tree are collected

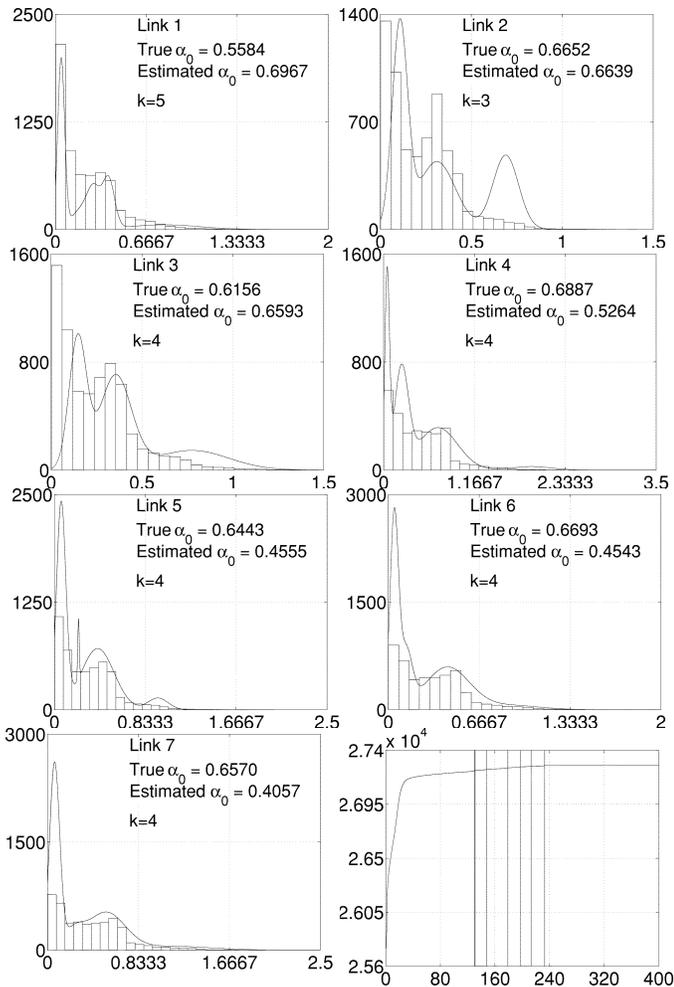


Fig. 6. Normalized ns-derived histograms for nonzero link delays and estimated Gaussian mixture density for indicated links. The horizontal axes denote link packet delays in milliseconds. Here, the unsupervised MML algorithm was initialized with five Gaussian components for each link. The lowest right graph shows the convergence curve of the MML penalized likelihood function. The vertical lines denote the iteration number where the number of Gaussian mixture components is reduced. The links affected by this reduction are in order {2, 4, 3, 2, 7, 6, 5}.

at the receiver nodes. We estimated each probe queueing delay by subtracting the minimum probe delay over the total samples for the same path.

The MML algorithm was implemented with Gaussian continuous mixtures, and the estimated number of mixture components at each link was initialized to $k_l = 5$. To accelerate the PML-EM algorithm, as in [15], we used the *component-wise EM algorithm for mixtures* (CEM²) [29]. Similarly to the SAGE algorithm of Fessler and Hero [32], the CEM² algorithm updates the parameters sequentially, instead of updating all of them simultaneously. The monotonicity property of CEM² is not affected by the order of updating. We adopted a cyclic updating procedures as follows: Update $\alpha_{1,0}$, recompute Q , update $\alpha_{1,1}$ and $\theta_{1,1}$, recompute Q , and so on, until all the parameters for link 1 are updated; then proceed in the same way for link 2, 3, and so on, until all link parameters are updated.

Link delay pdf estimates are shown in Fig. 6. To obtain ground truth, the true internal link delay distributions were estimated empirically from the ns simulated data. The mass of

the atom, which is denoted as “True $\alpha_{l,0}$ ” in the figure, is the empirically estimated probability of an empty queue at link l calculated from sample averages. The continuous portion of the true distribution is estimated by the histogram of nonzero link delay samples and normalized to have mass $1 - (\text{True } \alpha_{l,0})$. The estimated Gaussian mixtures are shown along with the normalized histogram for comparison. Note that the probability mass of the mixture is $1 - \hat{\alpha}_{l,0}$ for delays at link l . The convergence curve of the penalized likelihood function is shown in the lowest right graph in Fig. 6. The vertical lines indicate the iterations when at least one component is annihilated.

As shown in Fig. 6, the Gaussian mixture components capture the profile of the empirical continuous portion of the density for most of the links. They also provide accurate estimates to all the queueing delay ranges. Some modal mismatches occur in the estimates at, for example, links 2 and 3. This error is probably due to the limitation of the five Gaussian + 1 point mass component model. For a better fit to the internal delay histograms, it may be necessary to assign more point masses and include other density models that are flatter or more heavy tailed than Gaussian. Other sources of error might include violation of the spatial or temporal independence assumptions, insufficient number of probe samples to resolve link densities, insufficient number of iterations of the MML algorithm, existence of local maxima in the likelihood function, and burstiness (nonstationarity) of the traffic. These are topics worthy of additional investigation.

VI. CONCLUSION AND FUTURE WORK

This paper focuses on the estimation of internal link delay distributions from end-to-end unicast packet pair delay measurements when there is a positive probability of zero queueing delay, i.e., lightly to moderately loaded networks. We proposed a new hybrid discrete-continuous finite mixture model that circumvents the difficulties of link delay discretization. For the case that mixture model orders are known, we derived an EM algorithm to approximate the ML estimates. Model simulation showed that when all model assumptions hold the EM algorithm can very accurately estimate the delay distributions for each internal link. When the model orders are unknown, we implemented an MML order-selection penalty and derived an unsupervised algorithm for estimating both the number of mixture components and the continuous density parameters. Although the estimates obtained at convergence of this algorithm are not necessarily MML estimates, results of the ns-2 simulation showed that reasonably accurate estimates of internal link delay distributions are possible.

Future work includes finding ways to accelerate convergence of the ML-EM and PML-EM algorithms for real-time implementation. EM algorithms are quite slow, and the improvement made by CEM² is still limited. This makes it difficult to perform extensive comparisons. Another direction is extension of our model to include spatial dependencies of link delays among different links, especially the links along the same path. For time-varying scenarios, adaptive schemes need to be developed in order to capture possible changes in the traffic statistics and

the network environment. It may also be viable to apply these methods to detecting abnormal changes in link delay distributions. This may help early detection of possible network failures and/or malicious network activities.

APPENDIX
DERIVATION OF EM ALGORITHM

Here, we only sketch the derivation of the E- and M-step quantities required for the ML, PML, and CEM² algorithms. Throughout the derivation, we assume the point mass is located at zero delay for all the links. Let $g_i(\mathbf{y}; \Theta)$ denote the end-to-end packet pair delay joint pdf from the root node to receivers in the i th probe tree. $g_{i,(l,m)}(\mathbf{y}; \Theta)$ is defined similarly to $g_i(\mathbf{y}; \Theta)$, except in the convolution, f_l is replaced by its m th component, which is $\alpha_{l,0}$ when $m = 0$ or $\alpha_{l,m}\phi(x; \theta_{l,m})$ when $m \neq 0$. This is the likelihood of the end-to-end delays of a packet pair along the i th probe tree being equal to \mathbf{y} given the delay at link l is contributed by the m th hidden component. $h_{i,l}(\mathbf{y}; \Theta)$ is also similar to $g_i(\mathbf{y}; \Theta)$, except that the l th link is excluded in the convolution.

E-step: Compute the conditional expectation of the complete data log-likelihood in (6) of the t th iteration. Let

$$\begin{aligned}\omega_{l,m}^{(i,n)} &= E \left[Z_{l,m}^{(i,n)} | \mathbf{Y}^{(i,n)}; \hat{\Theta}^{(t)} \right] \\ &= \frac{P \left(Z_{l,m}^{(i,n)} = 1, \mathbf{Y}^{(i,n)}; \hat{\Theta}^{(t)} \right)}{P \left(\mathbf{Y}^{(i,n)}; \hat{\Theta}^{(t)} \right)} \\ &= \frac{g_{i,(l,m)} \left(\mathbf{y}^{(i,n)}; \hat{\Theta}^{(t)} \right)}{g_i \left(\mathbf{y}^{(i,n)}; \hat{\Theta}^{(t)} \right)}\end{aligned}$$

for $m = 0, \dots, k_l$ and $l = 1, \dots, L$. Define

$$\begin{aligned}Q_{l,m}^{(i,n)}(\theta_{l,m}) &= E \left[Z_{l,m}^{(i,n)} \log \phi(X_l^{(i,n)}; \theta_{l,m}) | \mathbf{Y}^{(i,n)}; \hat{\Theta}^{(t)} \right] \\ &= \int \frac{\hat{\alpha}_{l,m}^{(t)} \phi(x; \hat{\theta}_{l,m}^{(t)}) h_{i,l} \left((\mathbf{y} - \mathbf{x})^{(i,n)}; \hat{\Theta}^{(t)} \right) \log \phi(x; \theta_{l,m})}{g_i \left(\mathbf{y}^{(i,n)}; \hat{\Theta}^{(t)} \right)} \\ &\quad \cdot dx\end{aligned}$$

for $m = 1, \dots, k_l$ and $l = 1, \dots, L$, where $\mathbf{y} - \mathbf{x} = (y_1^{(i,n)} - x, y_2^{(i,n)} - x)$ if link l is shared by the two probe paths, or $(y_1^{(i,n)} - x, y_2^{(i,n)})$ if link l is on one of the branches of the probe tree, say branch 1. The conditional expectation of the log likelihood function in (6) becomes

$$Q \left(\Theta, \hat{\Theta}^{(t)} \right) = \sum_{l=1}^L \sum_{i: l \in M_i} \sum_{n=1}^{N_i} \left\{ \sum_{m=0}^{k_l} \omega_{l,m}^{(i,n)} \log \alpha_{l,m} + \sum_{m=1}^{k_l} Q_{l,m}^{(i,n)}(\theta_{l,m}) \right\}. \quad (\text{A.1})$$

Equation (A.1) holds for any choices of density ϕ . If the mixture components are Gaussian, $Q_{l,m}^{(i,n)}(\theta_{l,m})$ has the following form:

$$Q_{l,m}^{(i,n)}(\theta_{l,m}) = -\omega_{l,m}^{(i,n)} \left(\frac{\log 2\pi}{2} + \log \sigma_{l,m} + \frac{\mu_{l,m}^2}{2\sigma_{l,m}^2} \right) - \frac{\int (x^2 - 2\mu_{l,m}x) \hat{\alpha}_{l,m}^{(t)} \phi(x; \hat{\theta}_{l,m}^{(t)}) h_{i,l} \left((\mathbf{y} - \mathbf{x})^{(i,n)}; \hat{\Theta}^{(t)} \right) dx}{2\sigma_{l,m}^2 g_i \left(\mathbf{y}^{(i,n)}; \hat{\Theta}^{(t)} \right)}.$$

M-step: Update the parameter estimates by maximizing $Q(\Theta, \hat{\Theta}^{(t)})$ over Θ , as shown in (7). The estimator updates are

$$\begin{aligned}\hat{\alpha}_{l,m}^{(t+1)} &= \frac{\sum_{i: l \in M_i} \sum_{n=1}^{N_i} \omega_{l,m}^{(i,n)}}{\sum_{i: l \in M_i} N_i} \quad m = 0, \dots, k_l \\ \hat{\theta}_{l,m}^{(t+1)} &= \underset{\theta}{\operatorname{argmax}} \sum_{i: l \in M_i} \sum_{n=1}^{N_i} Q_{l,m}^{(i,n)}(\theta) \quad m = 1, \dots, k_l.\end{aligned}$$

In the case that the ϕ s are Gaussian densities, $\hat{\theta} = (\hat{\mu}, \hat{\sigma}^2)$, and

$$\begin{aligned}\hat{\mu}_{l,m}^{(t+1)} &= \frac{1}{\sum_{i: l \in M_i} \sum_{n=1}^{N_i} \omega_{l,m}^{(i,n)}} \cdot \sum_{i: l \in M_i} \sum_{n=1}^{N_i} \int x \cdot \hat{\alpha}_{l,m}^{(t)} \phi(x; \hat{\theta}_{l,m}^{(t)}) h_{i,l} \left((\mathbf{y} - \mathbf{x})^{(i,n)}; \hat{\Theta}^{(t)} \right) dx \\ &\quad \frac{g_i \left(\mathbf{y}^{(i,n)}; \hat{\Theta}^{(t)} \right)}{g_i \left(\mathbf{y}^{(i,n)}; \hat{\Theta}^{(t)} \right)} \\ \hat{\sigma}_{l,m}^{2(t+1)} &= \frac{1}{\sum_{i: l \in M_i} \sum_{n=1}^{N_i} \omega_{l,m}^{(i,n)}} \cdot \sum_{i: l \in M_i} \sum_{n=1}^{N_i} \int \left(x - \hat{\mu}_{l,m}^{(t+1)} \right)^2 \hat{\alpha}_{l,m}^{(t)} \phi(x; \hat{\theta}_{l,m}^{(t)}) h_{i,l} \left((\mathbf{y} - \mathbf{x})^{(i,n)}; \hat{\Theta}^{(t)} \right) dx \\ &\quad \frac{g_i \left(\mathbf{y}^{(i,n)}; \hat{\Theta}^{(t)} \right)}{g_i \left(\mathbf{y}^{(i,n)}; \hat{\Theta}^{(t)} \right)}.\end{aligned}$$

REFERENCES

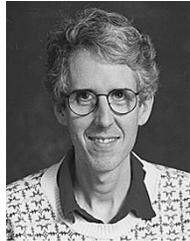
- [1] Y. Vardi, "Network tomography: Estimating source-destination traffic intensities from link data," *J. Amer. Statist. Assoc.*, vol. 91, pp. 365–377, 1996.
- [2] M. Coates, A. Hero, R. Nowak, and B. Yu, "Internet tomography," *IEEE Signal Processing Mag.*, vol. 19, pp. 47–65, May 2002.
- [3] J. Cao, D. Davis, S. V. Wiel, and B. Yu, "Time-varying network tomography: Router link data," *J. Amer. Statist. Assoc.*, vol. 95, pp. 1063–1075, 2000.
- [4] R. Caceres, N. G. Duffield, J. Horowitz, and D. F. Towsley, "Multicast-based inference of network-internal loss characteristics," *IEEE Trans. Inform. Theory*, vol. 45, pp. 2462–2480, Nov. 1999.
- [5] M. Coates and R. Nowak, "Network tomography for internal delay estimation," in *Proc. ICASSP*, Salt Lake City, UT, May 2001.
- [6] F. Lo Presti, N. G. Duffield, J. Horowitz, and D. Towsley, "Multicast-based inference of network-internal delay distributions," *IEEE/ACM Trans. Networking*, vol. 10, pp. 761–775, Dec. 2002.
- [7] National Laboratory for Applied Network Research (NLNAR). [Online]. Available: <http://moat.nlanr.net/index.html>.

- [8] The Surveyor Project. [Online]. Available: <http://www.advanced.org/surveyor>.
- [9] W. Jiang and H. Schulzrinne, "QoS measurement of Internet real-time multimedia services," in *Proc. NOSSDAV*, Chapel Hill, NC, June 2000.
- [10] K. Mochalski, J. Mischeel, and S. Donnelly, "Packet delay and loss at the Auckland internet access path," in *Proc. Passive Active Measurement Workshop*, Fort Collins, CO, Mar. 2002, pp. 46–55.
- [11] M. Coates and R. Nowak, "Sequential Monte Carlo inference of internal delays in nonstationary communication networks," *IEEE Trans. Signal Processing*, vol. 50, pp. 366–376, Feb. 2002.
- [12] N. G. Duffield, J. Horowitz, F. Lo Presti, and D. Towsley, "Network delay tomography from end-to-end unicast measurements," in *Proc. Int. Workshop Digital Commun.—Evolutionary Trends Internet*, Taormina, Italy, Sept. 2001.
- [13] Y. Tsang, M. Coates, and R. Nowak, "Nonparametric internet tomography," in *Proc. ICASSP*, Orlando, FL, May 2002.
- [14] M. Shih and A. Hero, "Unicast inference of network link delay distributions from edge measurements," in *Proc. ICASSP*, Salt Lake City, UT, May 2001.
- [15] M. Figueiredo and A. K. Jain, "Unsupervised learning of finite mixture models," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 24, pp. 381–396, Mar. 2002.
- [16] ns—Network simulator. [Online]. Available: <http://www-mash.cs.berkeley.edu/ns/ns.html>.
- [17] Y. Zhang, N. G. Duffield, V. Paxson, and S. Shenker, "On the constancy of internet path properties," in *Proc. ACM SIGCOMM Internet Measurement Workshop*, San Francisco, CA, Nov. 2001.
- [18] Y. Tsang, M. Coates, and R. Nowak, "Network delay tomography," *IEEE Trans. Signal Processing*, vol. 51, pp. 2125–2136, Aug. 2003. [Online]. Available: <http://spin.rice.edu/publications.html>.
- [19] N. G. Duffield, F. Lo Presti, V. Paxson, and D. Towsley, "Inferring link loss using striped unicast probes," in *Proc. IEEE Infocom*, Apr. 2001.
- [20] N. G. Duffield and F. Lo Presti, "Multicast inference of packet delay variance at interior network links," Preprint from MINC. [Online]. Available: <http://www-net.cs.umass.edu/minc/>.
- [21] V. Paxson and S. Floyd, "Wide area traffic: the failure of Poisson modeling," *IEEE/ACM Trans. Networking*, vol. 3, pp. 226–244, June 1994.
- [22] J. S. Marron and M. P. Wand, "Exact mean integrated squared error," *Ann. Statist.*, vol. 20, pp. 712–736, 1992.
- [23] D. Bertsekas and R. Gallager, *Data Networks*. Englewood Cliffs, NJ: Prentice-Hall, 1992.
- [24] M. Whindham and A. Cutler, "Information ratios for validating mixture analysis," *J. Amer. Statist. Assoc.*, vol. 87, pp. 1188–1192, 1992.
- [25] J. Rissanen, *Stochastic Complexity in Statistical Inquiry*. Singapore: World Scientific, 1989.
- [26] C. Wallace and P. Freeman, "Estimation and inference via compact coding," *J. R. Statist. Soc. B*, vol. 49, no. 3, pp. 241–252, 1987.
- [27] G. McLachlan and D. Peel, *Finite Mixture Models*. New York: Wiley, 2000.
- [28] J. Bernardo and A. Smith, *Bayesian Theory*. Chichester, U.K.: Wiley, 1994.
- [29] G. Celeux, S. Chrétien, F. Forbes, and A. Mkhadri, "A component-wise EM algorithm for mixtures," INRIA, Rhône-Alpes, France, Tech. Rep. 3746, 1999.
- [30] L. Xu and M. Jordon, "On convergence properties of the EM algorithm for Gaussian mixtures," *Neural Comput.*, vol. 8, pp. 129–151, 1996.
- [31] K. Claffy, G. Miller, and K. Thompson, "The nature of the beast: Recent traffic measurements from an internet backbone," in *Proc. INET*, Geneva, Switzerland, July 1998.
- [32] J. Fessler and A. Hero, "Space-alternating generalized expectation-maximization algorithm," *IEEE Trans. Signal Processing*, vol. 42, pp. 2664–2677, Oct. 1994.



Meng-Fu Shih (S'02) was born in Taipei, Taiwan, R.O.C., in 1972. He received the B.S. degree from Tatung Institute of Technology, Taipei, in 1994 and the M.S. degree from National Taiwan University, Taipei, in 1996, both in electrical engineering. Since 1999, he has been pursuing the Ph.D. degree with the Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor.

His research interests include signal processing for communications, statistical signal processing, estimation, and detection.



Alfred O. Hero, III (F'98) was born in Boston, MA, in 1955. He received the B.S. degree in electrical engineering (summa cum laude) from Boston University in 1980 and the Ph.D. from Princeton University, Princeton, NJ, in 1984, both in electrical engineering.

While at Princeton, he held the G.V.N. Lothrop Fellowship in Engineering. Since 1984, he has been a Professor with the University of Michigan, Ann Arbor, where he has appointments in the Department of Electrical Engineering and Computer Science, the Department of Biomedical Engineering, and the Department of Statistics. He has held visiting positions at I3S University of Nice, Sophia-Antipolis, France, in 2001, Ecole Normale Supérieure de Lyon, Lyon, France, in 1999, Ecole Nationale Supérieure des Télécommunications, Paris, France, in 1999, Scientific Research Labs of the Ford Motor Company, Dearborn, MI, in 1993, Ecole Nationale Supérieure des Techniques Avancées (ENSTA) and Ecole Supérieure d'Electricité, Paris, in 1990, and Lincoln Laboratory, Massachusetts Institute of Technology, Cambridge, from 1987 to 1989. His research has been supported by NIH, NSF, AFOSR, NSA, ARO, ONR, DARPA, and by private industry in the areas of estimation and detection, statistical communications, bioinformatics, signal processing, and image processing.

Dr. Hero served as Associate Editor for the IEEE TRANSACTIONS ON INFORMATION THEORY from 1995 to 1998 and again in 1999 and the IEEE TRANSACTIONS ON SIGNAL PROCESSING since 2002. He was Chairman of the Statistical Signal and Array Processing (SSAP) Technical Committee from 1997 to 1998 and Treasurer of the Conference Board of the IEEE Signal Processing Society. He was Chairman for Publicity of the 1986 IEEE International Symposium on Information Theory (Ann Arbor, MI) and General Chairman of the 1995 IEEE International Conference on Acoustics, Speech, and Signal Processing (Detroit, MI). He was co-chair of the 1999 IEEE Information Theory Workshop on Detection, Estimation, Classification, and Filtering (Santa Fe, NM) and the 1999 IEEE Workshop on Higher Order Statistics (Caesaria, Israel). He chaired the 2002 NSF Workshop on Challenges in Pattern Recognition. He co-chaired the 2002 Workshop on Genomic Signal Processing and Statistics (GENSIPS). He was Vice President (Finance) of the IEEE Signal Processing Society from 1999 to 2002. He was Chair of Commission C (Signals and Systems) of the U.S. National Commission of the International Union of Radio Science (URSI) from 1999 to 2002. He has been a member of the Signal Processing Theory and Methods (SPTM) Technical Committee of the IEEE Signal Processing Society since 1999. He is also a member of the SEDD Review Panel of the U.S. National Research Council. He will be President-Elect of the IEEE Signal Processing Society from 2004 to 2005. He is a member of Tau Beta Pi, the American Statistical Association (ASA), the Society for Industrial and Applied Mathematics (SIAM), and the U.S. National Commission (Commission C) of URSI. He received the 1998 IEEE Signal Processing Society Meritorious Service Award, the 1998 IEEE Signal Processing Society Best Paper Award, and the IEEE Third Millennium Medal. In 2002, he was appointed IEEE Signal Processing Society Distinguished Lecturer.