

Least Squares Estimates of Network Link Loss Probabilities using End-to-end Multicast Measurements

Bowei Xi, George Michailidis¹
and Vijayan N.Nair²

Department of Statistics
The University of Michigan
Ann Arbor, MI 48109-1092
email:

xbw, gmichail, vnn@umich.edu

Abstract — Estimation of quality of service parameters associated with large-scale computer and communication networks is a problem of considerable importance. In this paper, we consider estimation of link-level loss probabilities based on active tomography using multicast probing schemes. We formulate a regression framework for the problem, develop and study the properties of several types of least-squares based estimators. These include ordinary, generalized, and iteratively reweighted least squares estimators. We study the asymptotic and finite-sample properties of these estimators. The first two are simple to compute while the last two are asymptotically efficient. Computation of the variance-covariance matrix and inference using these estimators are much simpler computationally than those based on the MLE and E-M algorithm.

I. INTRODUCTION

Quality of service (QoS) parameters associated with computer and communications networks, such as link loss rates and delays, are of considerable interest to both providers and customers of network services. For example, it is now common for Internet service providers to offer a variety of service levels to customers. Service level agreements specify performance criteria that the network provider guarantees to satisfy. Such QoS criteria can include the amount of bandwidth made available to the customer and bounds on the maximum delay; the latter are important for time sensitive applications, such as Internet telephony and streaming applications. By using *active* and *passive* traffic measurement schemes, network tomography is capable of assessing the performance of modern day networks as well as localizing anomalous behavior to individual components and subnetworks [1, 4].

Large-scale network inference problems can be classified according to the type of data collected and the parameters of interest investigated. Figure 1 shows a computer network, comprised of *nodes* (computer terminals, routers or even whole subnetworks) and *links* that can be unidirectional or bidirectional, depending on the problem context and the desired level of abstraction. Messages are transmitted across the network by sending *packets* of bits from a *source* to a *destination* node along a *path* that usually passes through several other nodes.

The problem of inference in large scale networks involves estimating QoS parameters from traffic measurements at a




Fig. 1: A wired network comprised of computer terminals and routers

limited subset of nodes. Vardi [9] was among the earliest to study this problem and coined the term *network tomography*. Over the last few years, two forms of network tomography have been addressed in the literature: (1) *link-level parameter* estimation (such as packet loss rates and delay distributions) based on *end-to-end*, path-level traffic measurements and (2) *sender-receiver path-level* traffic intensity estimation based on link-level traffic measurements. The main drawback of the latter approach is that it requires gaining access to a wide range of routers and switches in an administratively diverse network, which can be a difficult task.

In link-level parameter estimation, the traffic measurements are obtained from *active probing* of the network. They consist of counts of packets transmitted and received or time delays between selected nodes, usually located on the *periphery* of the network (e.g. the numbered shaded nodes in Figure 1). The goal becomes to infer the *internal* link loss rates and delay distributions from end-to-end measurement collected from appropriately designed probing experiments. [3] advocated the use of *multicast* probing and stimulated much of the current work in the area [1, 2, 5, 6]. Unlike unicast transmissions, where each packet is sent from a source node to one and only one receiver, in multicast transmissions the sender effectively sends each packet to a group of subscribing receivers as follows. At internal routing nodes where forking occurs (see node 1 in Figure 1), the multicast packet is *replicated* and sent along each branching path [10]. The key to multicast transmissions for network tomography is that it introduces *dependencies* between end-to-end losses/delays measured by different receivers, which in turn enables inferences about (the unobserved) network internal links characteristics.

¹This work was partially supported by NSF grant IIS 9988095

²This work was partially supported by NSF grant DMS 0204247

Estimation of the link loss probabilities from multicast schemes was studied by [3] who developed a clever algorithm for computing the maximum likelihood estimator (MLE). However, computation of the Fisher information matrix (inverse of the asymptotic variance-covariance matrix of the MLE) is very involved and is difficult except in very small network topologies. Thus, confidence intervals and other inferences for the link loss probabilities are not easy to obtain. The use of the E-M algorithm for computing the MLE is discussed in [11]. While this is a conceptually simple, iterative algorithm that increases the likelihood at each iteration, it is a linear algorithm and can be slow to converge. It also suffers from the same computational problems for computing the information matrix.

In this paper, we study least-squares based estimation of internal link loss rates using multicast end-to-end measurements. By using a different parameterization, we derive a regression framework for this estimation problem. We propose various least squares (LS) estimators and derive their asymptotic properties. The ordinary least squares (OLS) and (one-step) generalized least squares (GLS) estimators are computationally easy to obtain. The GLS and iteratively reweighted least squares (IRWLS) estimators are asymptotically fully efficient. The IRWLS estimator performs very similarly to the MLE even in finite samples. Thus, we can exploit the advantages of these methods to develop computationally efficient methods of inference for very large networks.

We now introduce the modeling framework of the problem along with the necessary notation. Consider again the network shown in Figure 1. Packets are sent from a source (node $\{0\}$) to a set of destinations (nodes $\{4, 5, 6, 7\}$ following the solid paths). The end-to-end (path-level) behavior of the packets can be measured through a coordinated measurement scheme between the sender and the receivers. The sender can record whether a packet successfully reached its destination or was dropped/lost along its path. However, the sender cannot directly determine the specific link on which the packet is lost.

A physical network can be logically represented by a graph $G = (\mathcal{V}, \mathcal{E})$ consisting of nodes $v \in \mathcal{V}$ connected by edges/links $e \in \mathcal{E}$. For the problem at hand, the subset of the network over which active probing is performed is numbered according to a canonical scheme, with the sender being denoted by $\{0\}$ and the remaining nodes $1, 2, \dots$, while the links are assigned the number of the connected node below it. The graph $G = (\mathcal{V}, \mathcal{E})$ represents the *logical* topology of the monitored portion of the network. As in to [3], we investigate logical tree topologies defined as follows: let $\mathcal{T} = (\mathcal{V}, \mathcal{E})$ denote a tree with root $0 \in \mathcal{V}$, a set of nodes \mathcal{V} and a set of edges/links \mathcal{E} . The root node $\{0\} \in \mathcal{V}$ corresponds to the source of the transmitted packets. Let $\mathcal{D}(i) = \{j \in \mathcal{V} : (i, j) \in \mathcal{E}\}$ denote the set of *direct descendants* (children) of node i . The set of nodes $\mathcal{R} \subset \mathcal{V}$ such that $\mathcal{D}(i) = \emptyset$, $i \in \mathcal{V}$ (nodes without children) represents the set of *receivers*. Let $R = |\mathcal{R}|$ denote the cardinality of the receiver set. It is assumed that $|\mathcal{D}(i)| = 2$ for all $i \in \mathcal{V} - (\mathcal{R} \cup \{0\})$ (binary tree). Finally, let \mathcal{L} denote a *layer* of the tree, that is comprised of all nodes whose shortest paths from the root node $\{0\}$ are exactly L edges. It assumed that $|\mathcal{L}| = 2^{L-1}$, $L = 1, 2, \dots$ (symmetric binary tree). In this paper, we restrict attention to symmetric binary trees because their regularity makes the exposition easier. The derived results hold for any general tree topology, though (see [11]). Finally, we denote by $\mathcal{P}(i, j)$ a *path* between nodes i and j , which is

comprised of a set of connected links. An example of a 3-layer symmetric binary tree is given in Figure 2.

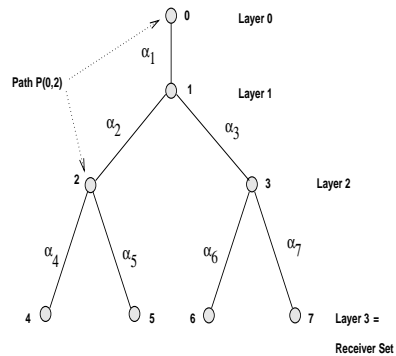


Fig. 2: A 3-layer symmetric binary tree

We study inference under the following stochastic environment for the multicast measurements. Let $\mathbf{X}_i = \{X_i(t); t = 1, 2, \dots\}_{i \in \mathcal{V}}$ be a stochastic process taking values in $\{0, 1\}$. Suppose a probe packet t was sent to node i from its *parent node*. Then $X_i(t) = 1$ indicates that the packet reached node i and $X_i(t) = 0$ that it did not. Since all packets originate at the root node, we set $X_0(t) = 1$ for all t . The link loss probabilities are given by $1 - \alpha_j$ where $\alpha_j = P(X_j(t) = 1)$. We further assume that the processes $\{\mathbf{X}_i, i \in \mathcal{V}\}$ are *mutually independent*. Note that we observe only the outcomes at receiver nodes r where $Z_r(t) = 1$ if and only if all the $X_i(t)$'s corresponding to path $\mathcal{P}(0, r)$ equal one. We can write $Z_r(t) = \prod_{j \in M_r} X_j(t)$ where M_r is the set of all nodes in the path $\mathcal{P}(0, r)$. This model is also analyzed in [3].

The paper is organized as follows. The regression formulation of the link loss problem based on end-to-end multicast measurements is presented in Section 2. The various LS estimation methods for the link loss rates are introduced and their asymptotic properties are derived in Section 3. Finite and large-sample efficiency comparisons of the various estimators are presented in Section 4. The paper ends with some concluding remarks.

II. LEAST SQUARES ESTIMATORS FOR LINK LOSS RATES AND THEIR PROPERTIES

As noted in Section 1, multicast end-to-end active probing corresponds to a multinomial experiment with 2^R possible outcomes $\mathbf{O} = \{\mathcal{O}_i\}$, $i = 1, \dots, 2^R$, whose probabilities γ_i are functions of the link loss rates α_j . We introduce next the necessary notation for being able to describe our regression model. First notice that there is a 1-1 correspondence between the canonical numbering of the receiver nodes \mathcal{R} in the logical tree topology and the set $\tilde{\mathcal{R}} = \{1, \dots, R\}$. For example, in a 3-layer tree node 4 $\in \mathcal{R}$ corresponds to the 1st receiver in $\tilde{\mathcal{R}}$, while node 7 to the 4th node in $\tilde{\mathcal{R}}$. Hence, the i^{th} multinomial outcome \mathcal{O}_i is an R -tuple taking values in $\{0, 1\}^R$, where a 1 in the r^{th} position indicates that the packet was received by the corresponding receiver in $\tilde{\mathcal{R}}$ and a 0 that it was not received. Hence, the outcome $\mathcal{O} = [1, 1, 1, 0]$ implies that the multicast packet was successfully transmitted only to receivers 4, 5 and 6, as shown in Figure 2. We denote by $\mathcal{O}_i(r)$ the r^{th} position in the tuple; e.g. in the example above we have $\mathcal{O}(1) = 1$, while $\mathcal{O}(4) = 0$.

Let \mathcal{A} describe the event that

$\mathcal{A}_S = \{at\ least\ receiver\ subset\ S\ successfully\ received\ packet\}$

, $\mathcal{S} \subseteq \tilde{\mathcal{R}}$. Formally, event $\mathcal{A}_{\mathcal{S}}$ is defined by

$$\mathcal{A}_{\mathcal{S}} = \cup_{i=1}^R \mathcal{O}_i \mathbf{1}_{\{\mathcal{O}_i(j)=1: j \in \mathcal{S}\}}, \quad \mathcal{S} \subseteq \tilde{\mathcal{R}}.$$

It is easy to see that there are $2^R - 1$ such events. For notational convenience and coherence we will describe such events by an R -tuple taking values in the set $\{1, +\}^R$, where a 1 in the r^{th} position indicates that the packet was received by the corresponding receiver in $\tilde{\mathcal{R}}$ and a + that it *may or may not* have been received. Thus, a multicast packet was certainly received by at least receivers 1 and 2 in $\tilde{\mathcal{R}}$ for the event $\mathcal{A} = [1, 1, +, +]$, while it was certainly received by at least receivers 2 and 4 for the event $\mathcal{A} = [+ , 1, +, 1]$. Let $\tilde{\gamma}_{\mathcal{S}}$ denote the probability of event $\mathcal{A}_{\mathcal{S}}$. We have that

$$\tilde{\gamma}_{\mathcal{S}} = \sum_{i=1}^R \gamma_i \prod_{k \in \mathcal{S}} \mathbf{1}_{\{\mathcal{O}_i(k)=1\}}.$$

For example the probability $\tilde{\gamma}([1, 1, +, 1]) = \gamma([1, 1, 1, 1]) + \gamma([1, 1, 0, 1])$. Notice that the $\tilde{\gamma}$'s are also functions of the underlying parameters of interest α . In general, we can write

$$\vec{\tilde{\gamma}} = Z\vec{\gamma} \quad (1)$$

for an appropriately defined binary matrix Z (an example of Z for a 3-layer tree is given in Appendix A).

Let $\vec{Y} = \log(\vec{\tilde{\gamma}})$ and let $\vec{\beta} = \log(\vec{\alpha})$, where $\vec{\tilde{\gamma}}$ are the method of moments estimates of the underlying quantities. The method of moments estimates of the multinomial probabilities are given by $\hat{\gamma}_i = N_i/N$, with N the total number of multicast probes used in the experiment and N_i the count of the event \mathcal{O}_i .

We then have that

$$\vec{Y} = X\vec{\beta} + \vec{\epsilon}, \quad (2)$$

where \vec{Y} is a $2^R - 1$ column vector, X is a $(2^R - 1) \times (2R - 1)$ binary design matrix, $\vec{\beta}$ a $2R - 1$ column vector of regression coefficients and $\vec{\epsilon}$ a column vector of unknown error terms with $E(\vec{\epsilon}) = 0$ and $E(\vec{\epsilon}\vec{\epsilon}') = V/N$. The design matrix takes values in the set $\{0, 1\}$, with $Z(i, j) = 1$ for all links j that belong to $\cup_{k \in \mathcal{R}} \mathcal{P}(0, k)$, such that $\mathcal{R}(k) = 1$ (the packet reached the k^{th} receiver). The design matrix X for the 3-layer tree shown in Figure 2 is given in Appendix B. The variance-covariance matrix V has the following form

$$V = (\text{diag}(\vec{\tilde{\gamma}}))^{-1} \Sigma (\text{diag}(\vec{\tilde{\gamma}}))^{-1},$$

where $\Sigma(i, j) = \tilde{\gamma}_k - \tilde{\gamma}_i \tilde{\gamma}_j$, $i, j = 1, \dots, 2R - 1$, with $\tilde{\gamma}_k$ corresponding to the parameter given by an *appropriate* intersection of the events i and j . For example, if $i = [1, 1, +, +]$ and $j = [+ , 1, 1, +]$ then $k = [1, 1, 1, +]$.

For example, the variance/covariance matrix V for a 2-layer tree takes the following form:

$$V = \text{diag}(\vec{\tilde{\gamma}})^{-1} \times \tilde{V} \times \text{diag}(\vec{\tilde{\gamma}})^{-1}$$

$$\tilde{V} = \begin{pmatrix} \gamma_{1,1}(1 - \gamma_{1,1}) & \gamma_{1,1}(1 - \gamma_{1,+}) & \gamma_{1,1}(1 - \gamma_{+,1}) \\ \gamma_{1,1}(1 - \gamma_{1,+}) & \gamma_{1,+}(1 - \gamma_{1,+}) & \gamma_{1,1} - \gamma_{1,+}\gamma_{+,1} \\ \gamma_{1,1}(1 - \gamma_{+,1}) & \gamma_{1,1} - \gamma_{1,+}\gamma_{+,1} & \gamma_{+,1}(1 - \gamma_{+,1}) \end{pmatrix}$$

Remark II.1 In [3] it is shown that the parameters $\vec{\alpha}$ are identifiable from the parameters of the multicast experiment $\vec{\tilde{\gamma}}$. Equation 1 establishes a 1-1 correspondence between $\vec{\tilde{\gamma}}$ and $\vec{\tilde{\alpha}}$ and therefore the link loss rates are also identifiable in the present setting.

Proposition II.1 *The design matrix of a L -layer symmetric binary tree for a multicast experiment is of full column rank. Therefore, the link loss rates are estimable by least squares methods.*

The proof of this Proposition can be found in [11].

Given model 2, the simplest method of estimation is ordinary least squares (OLS). The OLS estimator is given by

$$\vec{\beta}_{OLS} = (X'X)^{-1} X' \vec{Y}.$$

This estimator is not iterative in nature and hence is easy to compute. The only computationally demanding task is the computation of the inverse of the $X'X$ matrix. While the dimension of $X'X$ is $(2R - 1) \times (2R - 1)$, the dimension of X itself is $(2^R - 1) \times (2R - 1)$, so the number of rows grows exponentially with the number of layers in the tree. However, the matrix X has a special structure that can perhaps be exploited to develop efficient methods for computing the OLS. This and related computational issues will be studied in future work.

It is easy to establish the following results from the linear structure of the OLS estimator. Let $\vec{\beta}_0$ denotes the true value of the parameter.

Proposition II.2 *The OLS estimator is strongly consistent and asymptotically normal; i.e., as $N \rightarrow \infty$,*

$$\vec{\beta}_{OLS} \rightarrow \vec{\beta}_0 \quad \text{a.s.}$$

and

$$\sqrt{N}(\vec{\beta}_{OLS} - \vec{\beta}_0) \xrightarrow{L} W,$$

where W is normally distributed random vector with zero mean and variance-covariance matrix given by $(X'X)^{-1} X'V X (X'X)^{-1}$.

Proof: It is easily seen that $\vec{Y} \rightarrow \vec{Y}_0$ almost surely and that

$$\sqrt{N}(\vec{Y} - \vec{Y}_0) \xrightarrow{L} H$$

with H being a $N(\vec{0}, V)$ random vector since it is a linear function of the method of moments estimates of the multinomial parameters. Then, a straightforward application of Slutsky's theorem [7] establishes the result. ■

Remark II.2 *Notice that in the original scale, the limiting variance-covariance matrix of $\vec{\alpha}_{OLS}$ is given by $\text{diag}(\vec{\alpha})(X'X)^{-1} X'V X (X'X)^{-1} \text{diag}(\vec{\alpha})$.*

The observations \vec{Y} have unequal variances, so it is natural to consider the use of weighted least-squares to improve their efficiency. However, we have to estimate the unknown variance-covariance matrix. We will consider two different methods of estimating it. The first is a one-step GLS estimator which just plugs in the method of moment estimates for the unknown variance-covariance matrix. This is given by

$$\vec{\beta}_{GLS} = (X' \hat{V}^{-1} X)^{-1} X' \hat{V}^{-1} \vec{Y},$$

where \hat{V} is based on plugging in the method of moments estimates $\hat{\tilde{\gamma}}$. Note that this is also a non-iterative estimator and is easily computed. The only difficulty is the computation of the inverse of $X' \hat{V}^{-1} X$. As we see below (and is also known from the statistical literature), this one-step GLS estimator is

asymptotically as efficient as the MLE. However, the estimate of V based on the method of moments can be inefficient in finite samples. This is studied in the next section.

An alternative estimator that should behave similar to the MLE both asymptotically and in finite samples is the IRWLS estimator.

This has the same form as the GLS

$$\vec{\beta}_{IRWLS} = (X' \hat{V}^{-1} X)^{-1} X' \hat{V}^{-1} \vec{Y},$$

but it is an iterative estimator where \hat{V} is based on the previous iteration of the IRWLS estimator $\vec{\beta}$ and the iteration is repeated until convergence. While this is somewhat more demanding computationally than the MLE, the number of iterations required is much smaller. Our simulation results show that estimates based on even a single iteration have good finite sample properties.

Proposition II.3 *The GLS and IRWLS estimators are strongly consistent, asymptotically normal, and fully efficient; i.e.,*

$$\vec{\beta}_{GLS/IRWLS} \rightarrow \vec{\beta}_0 \text{ a.s.}$$

and

$$\sqrt{N}(\vec{\beta}_{GLS/IRWLS} - \vec{\beta}_0) \xrightarrow{\mathcal{L}} G,$$

where G is a normally distributed random vector with zero mean and variance-covariance matrix given by the inverse of the Fisher information matrix $I^{-1}(\vec{\beta}) = (X' V^{-1} X)^{-1}$.

Proof: Consistency and asymptotic normality follow by a similar argument as above. It is established in [11] that $X' V^{-1} X$ corresponds to the Fisher information matrix for $\log(\vec{\alpha}_{MLE})$. ■

As noted, the GLS/IRWLS estimates provide a computationally efficient alternative to the maximum likelihood estimates that are obtained in [3] and the EM-algorithm [11]. For 3 and 4-layer tree topologies, our experience suggests that the EM algorithm needs over 80 iterations to converge as opposed to 2-3 iterations required for the IRWLS estimator. Furthermore, the derivation of the Fisher information matrix is extremely cumbersome, while the calculation of $X' V^{-1} X$ is straightforward. Thus, confidence intervals and related inference problems are much more conveniently based on the LS estimation schemes.

III. FINITE AND LARGE-SAMPLE EFFICIENCY COMPARISONS

We begin with a simulation study of the finite-sample behavior of the various LS estimators. Although we have done an extensive comparison of their behavior for 3 and 4-layer trees, only selected results for the 4-layer tree are reported here due to space limitations. The data were generated under the independent Bernoulli loss model.

Figure 3 shows boxplots of the (Manhattan) distances between the true vector $\vec{\alpha}_0$ and its LS, GLS and IRWLS estimates. Here and in the rest of this section, the IRWLS estimates are based on only one iteration. All results are based on 100 replications. The left four panels (top two and bottom two) are based on 1,000 probes with all the link success probabilities set to be the same and equal, respectively, to 0.7, 0.8, 0.9 and 0.95. The right four panels show corresponding results for 10,000 probes. It can be seen that the IRWLS

estimates perform the best and are quite superior to the OLS and GLS estimates. It is somewhat surprising that the OLS method outperforms the GLS. The main reason is that the estimated covariance matrix \hat{V} based on the method of moments can be unstable which leads to instability in the estimator itself. This can be seen in more detail when we discuss Figure 4. Note also from Figure 3 that the performance of the estimates is better for higher link success probabilities and for larger probe size.




Fig. 3: Boxplots of Manhattan distances between the true link loss probability vector and its LS, GLS and IRWLS estimates for a 4-layer tree topology. Left 4 panels (top and bottom two) correspond to 1000 probe packets with all the link success probabilities set equal to .7, .8, .95 and .9 respectively. The right four panels (top and bottom two) correspond to analogous results for 10,000 probes.

Figure 4 displays boxplots of the simulation results for a representative element of $\vec{\alpha}$ from each layer of a 4-layer tree topology. The left four panels (top and bottom two) correspond to the case where all the elements of $\vec{\alpha}$ are set equal to 0.7; the right four to the case with 0.9. The results are based on 100 replications of the multicast experiment with 10000 probes each. We can see that, in general, the GLS has less variability than the OLS but that it can be biased in some cases. As noted before, this can be attributed to the instability in estimating the variance-covariance matrix. The IRWLS estimate performs well overall even with one iteration. Note also that the variability of all the estimates is smaller for the upper links (i.e. α_1 and α_2) and larger for the lower ones (i.e. α_4 and α_8). It is also smaller for larger success probabilities (.9).

Figures 5 and 6 provide some insight into the asymptotic efficiency of the LS estimator relative to the GLS/IRWLS estimators. As noted before, the GLS and IRWLS estimators have the same asymptotic variance-covariance matrix. Our experience (see [11]) suggests that similar behavior also holds in finite samples provided the number of probes sent is fairly large (e.g. larger than 5000 for a 4-layer tree). Figure 5 gives the ratio of the asymptotic variances of the OLS and GLS/IRWLS estimators for representative elements of $\vec{\alpha}$ (in log-scale). The transmission success probabilities have been set equal to one value for the top and bottom links of the tree topology (i.e. $\alpha_1 = \alpha_8 = \dots = \alpha_{15}$) and to another value those of the middle layer links (i.e. $\alpha_2 = \dots = \alpha_7$). It can be seen

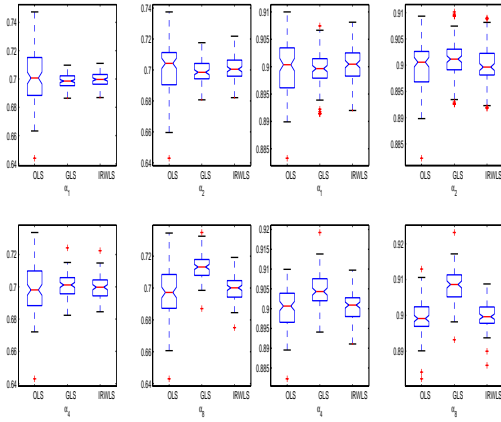


Fig. 4: Boxplots of selected elements of $\vec{\alpha}$. Left 4 panels correspond to success probabilities of .7 and right 4 panels to .9.

that for small probabilities the GLS/IRWLS estimates are orders of magnitude more efficient than the corresponding OLS estimates, with the ratio becoming almost one for probabilities very close to 1. A more detailed view is given in Figure 6, where the range of probabilities is restricted to the $[0.8, 1)$ interval. We can see from this graph (the efficiency comparisons are now plotted in the original scale) that the efficiency loss is twice as large for the OLS estimates for the upper link parameters (i.e. α_1 and α_2). Moreover, the efficiency deteriorates rapidly when the successful transmission probabilities become less than .9.




Fig. 5: Ratios of variances (in log-scale) for select elements of $\vec{\alpha}$ and for different successful transmission probabilities.

To get an overall measure of the relative efficiency, we can use the ratio of determinants of the variance-covariance matrices of the OLS and GLS estimators. This is shown in Figure 7 which demonstrates, to a large extent, a similar pattern as the one discussed above for the variances. It is of practical interest to assess the loss of efficiency as a function of computational complexity, which is a topic of current research.

IV. CONCLUSIONS AND FUTURE WORK




Fig. 6: Ratios of variances (in log-scale) for select elements of $\vec{\alpha}$ and for different successful transmission probabilities.




Fig. 7: Ratios of determinants of variance-covariance matrices (in log-scale) for different successful transmission probabilities.

This paper presents a regression framework for estimating a network's internal loss rates based on active probing and using multicast end-to-end measurements. Various least squares estimators are introduced and their finite and asymptotic properties studied. The IRWLS estimator is asymptotically efficient and inference based on the estimator is computationally simple to implement. Empirical evidence based on extensive simulations suggests that the OLS estimator has good performance with respect to bias but is fairly inefficient compared to the IRWLS estimator and is recommended only when the number of probes is very large. The one-step GLS estimator is a non-iterative scheme and is therefore also appealing. However, it can be biased when the number of probes is small to moderate.

The results thus far have assumed that the link loss probabilities are constant. In practice, however, it is more reasonable to assume that they are random. Furthermore, most of the current literature is based on the assumption of spatial

and temporal independence of the link loss behavior. Future work will focus on estimation and inference for link loss probabilities by relaxing these restrictions.

Appendix A

The transformation that leads from $\vec{\gamma}$ to $\vec{\tilde{\gamma}}$ for a 3-layer tree is given by

$$\begin{pmatrix} \tilde{\gamma}[(1, 1, 1, 1)] \\ \tilde{\gamma}[(1, 1, 1, +)] \\ \tilde{\gamma}[(1, 1, +, 1)] \\ \tilde{\gamma}[(1, +, 1, 1)] \\ \tilde{\gamma}[(+, 1, 1, 1)] \\ \tilde{\gamma}[(+, +, 1, 1)] \\ \tilde{\gamma}[(+, 1, +, 1)] \\ \tilde{\gamma}[(+, 1, 1, +)] \\ \tilde{\gamma}[(1, +, +, 1)] \\ \tilde{\gamma}[(1, +, 1, +)] \\ \tilde{\gamma}[(1, 1, +, +)] \\ \tilde{\gamma}[(+, +, +, 1)] \\ \tilde{\gamma}[(+, +, 1, +)] \\ \tilde{\gamma}[(+, 1, +, +)] \\ \tilde{\gamma}[(1, +, +, +)] \end{pmatrix} = Z \begin{pmatrix} \gamma[(1, 1, 1, 1)] \\ \gamma[(1, 1, 1, 0)] \\ \gamma[(1, 1, 0, 1)] \\ \gamma[(1, 0, 1, 1)] \\ \gamma[(0, 1, 1, 1)] \\ \gamma[(0, 0, 1, 1)] \\ \gamma[(0, 1, 0, 1)] \\ \gamma[(0, 1, 1, 0)] \\ \gamma[(1, 0, 0, 1)] \\ \gamma[(1, 0, 1, 0)] \\ \gamma[(1, 1, 0, 0)] \\ \gamma[(0, 0, 0, 1)] \\ \gamma[(0, 0, 1, 0)] \\ \gamma[(0, 1, 0, 0)] \\ \gamma[(1, 0, 0, 0)] \end{pmatrix}$$

with Z given by

$$Z = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 1 & 1 & 1 & 1 & 1 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 0 \\ 1 & 1 & 0 & 1 & 1 & 1 & 0 & 1 & 0 & 1 & 0 & 0 & 1 & 0 & 0 \\ 1 & 1 & 1 & 0 & 1 & 0 & 1 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 \\ 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 0 & 0 & 0 & 1 \end{pmatrix}$$

Appendix B

The regression model for a 3-layer symmetric binary tree with multicast measurements is given next:

$$\begin{pmatrix} \log(\hat{\gamma}(1, 1, 1, 1)) \\ \log(\hat{\gamma}(1, 1, 1, +)) \\ \log(\hat{\gamma}(1, 1, +, 1)) \\ \log(\hat{\gamma}(1, +, 1, 1)) \\ \log(\hat{\gamma}(+, 1, 1, 1)) \\ \log(\hat{\gamma}(+, +, 1, 1)) \\ \log(\hat{\gamma}(+, 1, +, 1)) \\ \log(\hat{\gamma}(+, 1, 1, +)) \\ \log(\hat{\gamma}(1, +, +, 1)) \\ \log(\hat{\gamma}(1, +, 1, +)) \\ \log(\hat{\gamma}(1, 1, +, +)) \\ \log(\hat{\gamma}(+, +, +, 1)) \\ \log(\hat{\gamma}(+, +, 1, +)) \\ \log(\hat{\gamma}(+, 1, +, +)) \\ \log(\hat{\gamma}(1, +, +, +)) \end{pmatrix} = X \times \begin{pmatrix} \log(\alpha_1) \\ \log(\alpha_2) \\ \log(\alpha_3) \\ \log(\alpha_4) \\ \log(\alpha_5) \\ \log(\alpha_6) \\ \log(\alpha_7) \end{pmatrix} + \vec{\epsilon}$$

$$X = \begin{pmatrix} 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 & 1 \\ 1 & 1 & 1 & 1 & 0 & 1 & 1 \\ 1 & 1 & 1 & 0 & 1 & 1 & 1 \\ 1 & 0 & 1 & 0 & 0 & 1 & 1 \\ 1 & 1 & 1 & 0 & 1 & 0 & 1 \\ 1 & 1 & 1 & 0 & 1 & 1 & 0 \\ 1 & 1 & 1 & 1 & 0 & 0 & 1 \\ 1 & 1 & 1 & 1 & 0 & 1 & 0 \\ 1 & 1 & 0 & 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 1 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 & 1 & 0 & 0 \\ 1 & 1 & 0 & 0 & 1 & 0 & 0 \\ 1 & 1 & 0 & 1 & 0 & 0 & 0 \end{pmatrix}$$

REFERENCES

- [1] Adams, A., Bu, T., Caceres, R., Duffield, N.G., Friedman, T., Horowitz, J., Presti, F.L., Moon, S.B., Paxson, V., and Towsley, D. (2000), The use of end-to-end multicast measurements for characterizing internal network behavior, *IEEE Communications Magazine*, 38, 152-159.
- [2] Bu, T., Duffield, N., Presti, F.L., and Towsley, D. (2002), Network tomography on general topologies, *Proceedings of ACM SIGMETRICS*.
- [3] Caceres, R., Duffield, N.G., Horowitz, J. and Towsley, D. (1999), Multicast based inference of network internal loss characteristics, *IEEE Transactions on Information theory*, 45, 2462-2480
- [4] Coates, M., Hero, A., Nowak, M. and Yu, B. (2002), Internet tomography, *IEEE Signal Processing Magazine*, 19, 47-65.
- [5] Duffield, N.G., Horowitz, J., Presti, F.L., and Towsley, D. (2002), Multicast topology inference from measured end-to-end loss, *IEEE Transactions on Information theory*, 48, 26-45
- [6] Duffield, N.G., Horowitz, J., Towsley, D., Wei, W., and Friedman T., Multicast-based loss inference with missing data, to appear in *IEEE Journal on Selected Areas in Communications*.
- [7] Ferguson, T.S. (1996), *A Course in Large Sample Theory*, Chapman & Hall, London
- [8] Shih, Meng-Fu, Hero, Alfred O, Unicast-based Inference of Network Link Delay Distributions Using Mixed Finite Mixture Models, submitted to *IEEE Transactions on Signal Processing*.
- [9] Vardi, Y. (1996), Network tomography: estimating source-destination traffic intensities from link data, *Journal of the American Statistical Association*, 91, 365-377
- [10] Walrand, J., and Varaiya, P. (2000), High-performance communication networks, *Morgan Kaufmann*.
- [11] Xi, B., Michailidis, G. and Nair, V.N. (2002), Least squares and maximum likelihood estimators for network link loss rates based on active end-to-end multicast measurements, Technical Report # 392, Department of Statistics, The University of Michigan