Explicit Loss Inference in Multicast Tomography

N.G. Duffield
AT&T Labs–Research
180 Park Avenue
Florham Park, NJ 07932, USA
duffield@research.att.com

J. Horowitz
Dept. Math. & Statistics
University of Massachusetts
Amherst, MA 01003, USA
joeh@math.umass.edu

F. Lo Presti
Dipartimento di Informatica
Universita dell Aquila
67010 Coppito (AQ), Italy
lopresti@di.univaq.it

D. Towsley
Dept. of Computer Science
University of Massachusetts
Amherst, MA 01003, USA
towsley@cs.umass.edu

Abstract— **Network performance tomography involves correlating end-to-end performance measures over different network paths to infer the performance characteristics on their intersection. Multicast based inference of link loss rates is the first paradigm for the approach. Existing algorithms generally require numerical solution of polynomial equations for a maximum likelihood estimator (MLE), or iteration when applying the expectation maximization (EM) algorithm. The purpose of this note is to demonstrate a new estimator for link loss rates that is computationally simple, being an explicit function of the measurements, and that has the same asymptotic variance as the MLE, to first order in the link loss rates.**

Keywords— **Statistical Inference, Link Loss Rates, End-to-end measurement**

I. Introduction

A. Summary

Network tomography is becoming a rapidly established discipline. One branch of this focuses on the development of statistical techniques for inferring internal network properties, such as link loss rates [2], [4], [1], link delay statistics [7], [12], and topology [6], [9], based on end-to-end packet measurements. In this note, we focus on the loss inference problem where loss observations are taken at the leaves of a tree over which packets are multicast. The current solution to this inference problem relies on obtaining maximum likelihood estimates (MLEs) of link-level loss rates. In general, this requires finding the roots of polynomial equations associated with the internal nodes in the tree where the polynomial degree corresponds to the branching factor of the associated node. Iterative solutions via the EM algorithm can also be obtained.

In this note we derive a simple *explicit formula* for the link loss rate estimates. Although they do not correspond to the MLEs, the estimators are consistent, i.e., they converge to the true loss rates as the number of measurements grows. Furthermore, the asymptotic variance of the explicit estimator equals that of the MLE to at least first order in the loss rates.

In the remainder of this section we describe some related work. Section II defines the underlying model for multicast loss inference. The existing MLE and the explicit estimator are described in Section III. Last, various aspects, including computational complexity and efficiency, of the two estimators are compared in Section IV.

B. Related Work

In addition to the particular problem motivating our work, namely the inference of link-level characteristics based on end-to-end observations, two other problems are of interest. The first of these is concerned with identifying source-destination traffic matrices based on link-level packet counts made within a network, [13], [3], [14]. The second is concerned with identifying *path-level* characteristics based on end-to-end observations made over a single path. Much of the path-level characterization work has focussed on estimating available bandwidth [8], [10].

II. Model

We describe the model for the multicast trees and the propagation of probe packets down the trees.

A. Model of Multicast Trees

Let $\mathcal{T} = (V, L)$ denote the logical multicast tree from a given source, consisting of the set of nodes V, including the source and receivers, and the set of links L. A link is an ordered pair $(j, k) \in V \times V$ denoting a link from node j to node k. The set of children of a node j is denoted by d_j (i.e. $d_j = \{k \in V : (j, k) \in L\}$). $U = V \setminus \{0\}$ denotes the set of all non-root nodes. For each node $j \in U$ there is a unique node $k = f(j)$, the parent of j, such that $(k, j) \in L$. We shall define $f^n(k)$ recursively by $f^n(k) = f(f^{n-1}(k))$, with $f^0(k) = k$. We say that j is a descendant of k, and write $j \prec k$, if $k = f^n(j)$ for some integer $n > 0$. $j \vee k$ denotes the closest common ancestor of nodes j and k, i.e., the unique node ℓ such that $j, k \prec \ell$, with no other node $\ell' \prec \ell$ with this property. The set of leaves of $\mathcal{T}$ is denoted by R. $\mathcal{T}_k = (V_k, L_k)$ denotes the subtree of $\mathcal{T}$ descended from k, and $R_k = V_k \cap R$.

B. Model of Probe Process

We now describe the loss model in more detail. With each node $k \in V$ we associate a probability $\alpha_k \in [0, 1]$ that a given probe packet is not lost on the link terminating at k. We model the passage of probes down the tree by a stochastic process $X = (X_k)_{k \in V}$ where each X_k takes a value in $\{0, 1\}$; $X_k = 1$ signifies that a probe packet reaches node k, and 0 that it does not. The packets are generated at the source, so $X_0 = 1$. For all other $k \in V$, the value of X_k is determined as follows. If $X_k = 0$ then $X_j = 0$ for the children j of k (and hence for all descendants of k). Conditioned on $X_k = 1$, then $\{X_j : j \in d_k\}$ are independent random variables, with the probability of the probe to reach k from $f(k)$ being $\mathsf{P}[X_k = 1 \mid X_{f(k)} = 1] = \alpha_k$. Although there is no link terminating at 0, we shall adopt the convention that $\alpha_0 = 1$, in order to avoid excluding the root link from expressions concerning the α_k. We will denote by $\overline{x}$ the quantity $1 - x$. $A_k = \prod_{j \succeq k} \alpha_j$ is the probability that a probe reaches k from the root 0.

III. Measurement and Inference

We describe the measurements obtained from a set of multicast probes. The existing MLE is reviewed in Section III-C. Our new explicit estimator is described in Section III-D.

A. Measurements and Data

Consider an experiment in which n probes are dispatched from the root 0. Each probe $i = 1, \dots, n$ gives rise to an independent realization $X^{(i)}$ of the probe process X. The leaf measurements $(X_k^{(i)})_{k\in R}^{i=1,\dots,n}$ comprise the data for inference. For each outcome $x \in \Omega \equiv \{0,1\}^{\#R}$, let $n(x)$ denote the number of probes $i = 1, \dots n$ for which $X_k^{(i)} = x_k$ for all $k \in R$.

B. Likelihood Function

Write $\alpha = \{\alpha_k : k \in U\}$ and let $p_\alpha(x) = \mathsf{P}_\alpha[X_k = x_k, \ \forall k \in R]$ denote the probability of an outcome $x \in \Omega$. The log-likelihood to obtain the data $(X^{(1)}, \dots, X^{(n)})$ can be written in terms of the $n(x)$ as

$$\mathcal{L}(\alpha) = \log \mathsf{P}_\alpha[(X^{(1)}, \dots, X^{(n)})] = \sum_{x\in\Omega} n(x) \log p_\alpha(x). \quad (1)$$

C. Maximum Likelihood Estimator

We characterize the Maximum Likelihood Estimator (MLE) of α, namely, $\arg\max_\alpha \mathcal{L}(\alpha)$, as follows. For each $k \in U$ set

$$Y_k = \max_{j\in R(k)} X_j \quad \text{and} \quad \gamma_k = \mathsf{P}[Y_k = 1]. \quad (2)$$

Thus $Y_k = 1$ if the probe reaches at least one receiver descended from node k. Denote by $\widehat{\gamma}_k$ the corresponding empirical quantity, i.e., the proportion of the n probes that reach at least one leaf descended from k. Hence, with

$$Y_k^{(i)} = \max_{j\in R(k)} X_j^{(i)} \quad \text{then} \quad \widehat{\gamma}_k = n^{-1} \sum_{i=1}^{n} Y_k^{(i)}. \quad (3)$$

In what follows we consider α to lie in the open parameter set $\mathcal{A} = \{\alpha \mid \alpha_k \in (0,1), k \in U\}$. Some of the results of the following theorem also hold on subsets of the boundary of $\mathcal{A}$.

Theorem 1 ([2], [5]) Assume $\alpha \in \mathcal{A}$.
(i) For each $k \in U$, $(1 - \gamma_k/A_k) = \prod_{j\in d_k}(1 - \gamma_j/A_k)$, where the empty product for $k \in R$ is taken as 0.
(ii) Let $\mathcal{G} = \{(\gamma_k)_{k\in U} : \gamma_k > 0 \,\forall k;\ \gamma_k < \sum_{j\in d_k} \gamma_j \,\forall k \in U \setminus R\}$. For each $\gamma \in \mathcal{G}$ and $k \in U$, the equation in (i) has a unique solution, denoted $\mathcal{H}_k(\gamma)$, in the interval $(\gamma_k, 1)$. Moreover, it is the largest real solution of that equation.
(iii) If $\widehat{\gamma} \in \mathcal{G}$, the likelihood equation $\{\partial\mathcal{L}_c/\partial\alpha_k = 0 : k \in U\}$ has unique solution $\widehat{\alpha}_k = \widehat{A}_k/\widehat{A}_{f(k)}$ where $\widehat{A}_k = \mathcal{H}_k(\widehat{\gamma})$. Moreover, $\widehat{\alpha}_k$ and $\widehat{A}_k$ are the MLEs for α_k and A_k respectively.

D. Explicit Estimator

For $k \notin R$, define $Z_k = \min_{j\in d_k} Y_j$, thus $Z_k = 1$ iff for each child j of k, a probe reaches some leaf descended through j. Similarly, define $Z_k^{(i)} = \min_{j\in d_k} Y_j^{(i)}$ Let $B_k = \mathsf{P}_\alpha[Z_k = 1]$, with the corresponding empirical empirical probability $\widehat{B}_k = n^{-1} \sum_{i=1}^{n} Z_k^{(i)}$.

Theorem 2: Let $\alpha \in \mathcal{A}$.
(i) For $k \notin R$.

$$A_k = \Phi_k(B_k, \gamma) := \left(\frac{\prod_{j\in d_k} \gamma_j}{B_k} \right)^{1/(\#d_k - 1)} \quad (4)$$

(ii) Define $\check{A}_k = \widehat{\gamma}_k$ for $k \in R$, and $\check{A}_k = \Phi_k(\widehat{B}_k, \widehat{\gamma})$ otherwise. Then $\check{A}_k$ is a consistent estimator of A_k, and hence $\check{\alpha}_k = \check{A}_k/\check{A}_{f(k)}$ is a consistent estimator of α_k.

Proof: (i) $Z_k = 1$ implies $X_k = 1$. Hence

$$\begin{aligned} B_k &= \mathsf{P}_\alpha[Z_k = 1,\ X_k = 1] = A_k \mathsf{P}_\alpha[Z_k = 1 \mid X_k = 1] \\ &= A_k \prod_{j\in d_k} \mathsf{P}_\alpha[Y_j = 1 \mid X_k = 1] = A_k \prod_{j\in d_k} (\gamma_j/A_k) \quad (5) \end{aligned}$$

(ii) By the law of large numbers, $\widehat{\gamma} \to \gamma$ and $\widehat{B}_k \to B$ almost surely as $n \to \infty$. Thus the results follows from (i) and the evident continuity of Φ_k on the set of values (B_k, γ) generated as α ranges over $\mathcal{A}$. ∎

IV. Comparison of Estimators

We compare the forms of the ML and explicit estimators for simple examples in Section IV-A. Computational complexities are compared in Section IV-B. The variance of the explicit estimator is calculated in Section IV-C, and compared with that of the MLE to first order in loss rates in Section IV-D. Higher order behavior in a simple topologies is compared in Section IV-E.

A. Examples

Binary Node: The ML and explicit estimators are equal. Let k have children j and j'. The MLE for A_k is $\widehat{\gamma}_j\widehat{\gamma}_{j'}/(\widehat{\gamma}_j + \widehat{\gamma}_{j'} - \widehat{\gamma}_k)$. Now $\widehat{B}_k = \widehat{\gamma}_j + \widehat{\gamma}_{j'} - \widehat{\gamma}_k$ and one sees from Theorem 2 that the explicit estimator is the same as the MLE.

Tertiary Node: Let k have children j, j', j''. By Theorem 1 the MLE $\widehat{A}_k = \mathcal{H}_k(\widehat{\gamma}_k, \widehat{\gamma}_j, \widehat{\gamma}_{j'}, \widehat{\gamma}_{j''})$ is the largest root of the quadratic equation $\widehat{A}_k^2(\widehat{\gamma}_j + \widehat{\gamma}_{j'} + \widehat{\gamma}_{j''} - \widehat{\gamma}_k) - \widehat{A}_k(\widehat{\gamma}_j\widehat{\gamma}_{j'} + \widehat{\gamma}_{j'}\widehat{\gamma}_{j''} + \widehat{\gamma}_{j''}\widehat{\gamma}_j) + \widehat{\gamma}_j\widehat{\gamma}_{j'}\widehat{\gamma}_{j''} = 0$. The explicit estimator is $\check{A}_k = (\widehat{\gamma}_j\widehat{\gamma}_{j'}\widehat{\gamma}_{j''}/\widehat{B}_k)^{1/2}$.

B. Computational Complexity

Determining $\widehat{A}_k$ requires finding the roots of a polynomial equation of degree $\#d_k - 1$. Thus, by a well known result in Galois theory, $\widehat{A}_k$ must in general be determined numerically when $\#d_k > 5$. On the other hand, the explicit estimate $\check{A}_k$ is determined explicitly from the measured data.

In calculation of the empirical probabilities $\widehat{\gamma}_k$, both methods require computation of the quantities $Y_j^{(i)}$. This takes $O(n\#U)$ operations for calculating over all probes and nodes. In addition, the explicit estimator $\check{A}_k$ uses the empirical probability $\widehat{B}_k$ which requires the computation of $Z_k^{(i)}$ for each node. This takes another $O(n\#U)$ operations.

Summarizing, in computing the empirical probabilities, the explicit estimator requires twice as many operations as the MLE, but thereafter the computation of the explicit estimator is far simpler, needing an explicit computation on the empirical probabilities, while the MLE requires finding a root in general.

C. Variance

As an MLE, $\widehat{A}_k$ inherits the property of asymptotic efficiency common to all MLEs, i.e., its asymptotic variance as the number of probes $n \to \infty$ is the same as that of the minimum variance unbiased estimator of A_k. Since the explicit estimator is asymptotically unbiased, it is expected to have variance no less than the MLE. For both estimators, even the asymptotic variances have a relative complex form. We now show that the asymptotic variances of $\widehat{A}_k$ and $\check{A}_k$ are equal, to first order in the link loss rates $\overline{\alpha}_k$. Let $s_k = \sum_{j \succeq k} \overline{\alpha}_k$, $t_k = \sum_{j \in d_k} \overline{\alpha}_k$, and $\|\overline{\alpha}\| = \max_{k \in U} |\overline{\alpha}_k|$.

Theorem 3: (i) As $n \to \infty$, $n^{1/2}(\widehat{A}_k - A_k)$ has an asymptotically Gaussian distribution of mean 0 and variance v_k^{MLE}, where $v_k^{\mathrm{MLE}} = s_k + O(\|\overline{\alpha}\|^2)$.
(ii) As $n \to \infty$, $n^{1/2}(\check{A}_k - A_k)$ has an asymptotically Gaussian distribution of mean 0 and variance v_k^{E}, where $v_k^{\mathrm{E}} = s_k + O(\|\overline{\alpha}\|^2)$.

For the proof of Theorem 2 and other results, it is useful to collect together some results established during the proof of Theorem 5 in [2].

Lemma 1: (i) $\mathsf{Cov}(Y_j, Y_k) = \gamma_k(1 - \gamma_j)$ when $j \succeq k$ and $\gamma_j \gamma_k (1/A_{j \vee k} - 1)$ otherwise.
(ii) $A_k = 1 - s_k + O(\|\overline{\alpha}\|^2)$, $\gamma_k = 1 - s_k + O(\|\overline{\alpha}\|^2)$ and $\mathsf{Cov}(Y_j, Y_k) = s_{j \vee k} + O(\|\overline{\alpha}\|^2)$

Proof of Theorem 3: (i) is proved in Theorem 5 of [2].

(ii) *Components of the asymptotic variance:* by the central limit theorem, $n^{1/2}(\widehat{B}_k - B_k)$ and $n^{1/2}(\widehat{\gamma}_j - \gamma_j)$ converge in distribution to zero mean Gaussian random variables as $n \to \infty$. Hence, by the Delta method (see Chapter 7 of [11]), $n^{1/2}(\check{A}_k - A_k)$ converges in distribution to a zero mean Gaussian random variable with variance

$$v_k^{\mathrm{E}} = \nabla \Phi_k \cdot C^{\mathrm{E}} \, \nabla \Phi_k \tag{6}$$

where $\nabla \Phi_k$ is the $\# d_k + 1$ dimensional derivative

$$\nabla \Phi_k = \left(\frac{\partial \Phi_k(B_k, \gamma)}{\partial B_k}, \left\{ \frac{\partial \Phi_k(B_k, \gamma)}{\partial \gamma_j} \, : \, j \in d_k \right\} \right)^T \tag{7}$$

and C^{E} is the $\# d_k + 1$ dimensional square covariance matrix

$$\begin{pmatrix} \mathsf{Var}(Z_k) & \{\mathsf{Cov}(Z_k, Y_j) \, : \, j \in d_k\} \\ \{\mathsf{Cov}(Z_k, Y_j) \, : \, j \in d_k\}^T & \{\mathsf{Cov}(Y_j, Y_{j'}) \, : \, j, j' \in d_k\} \end{pmatrix} \tag{8}$$

It remains to calculate $\nabla \Phi_k$ and C^{E}, and derive their first order behavior in $\overline{\alpha}$. For the derivative

$$\nabla \Phi_k(B_k, \gamma) = (-1/B_k, \{1/\gamma_j \, : \, j \in d_k\}) \Phi_k / (\# d_k - 1) \tag{9}$$

For the covariance, $\mathsf{Var} Z_k = B_k(1 - B_k)$ while $\mathsf{Cov}(Y_j, Y_{j'})$ is described in Lemma 1(i). Finally, note that $Z_k = 1$ requires $Y_j = 1$ for all $j \in d_k$, and hence $\mathsf{Cov}(Z_k, Y_j) = B_k(1 - \gamma_j)$.

D. Leading Order Variance in $\overline{\alpha}$:

We derive the leading order behavior of v_k^{E} in $\overline{\alpha}$. Combining Lemma 1(ii) with (5), one finds $B_k = 1 - s_k - t_k + O(\|\overline{\alpha}\|^2)$. Inserting into the previous expressions for the components $\nabla \Phi_k$ and C^{E}:

$$\begin{aligned} \nabla \Phi_k &= \frac{1}{\# d_k - 1} (-1, 1, \ldots, 1)^T + O(\|\overline{\alpha}\|) \\ C^{\mathrm{E}} &= s_k + \begin{pmatrix} t_k & \{\overline{\alpha}_j \, : \, j \in d_k\} \\ \{\overline{\alpha}_j \, : \, j \in d_k\}^T & \mathrm{Diag}\{\overline{\alpha}_j \, : \, j \in d_k\} \end{pmatrix} \\ & \quad + O(\|\overline{\alpha}\|^2) \end{aligned}$$

where $\mathrm{Diag}(S)$ denotes the diagonal matrix with diagonal entries in S. The stated first order form for v_k^{E} then follows. ∎

E. Example: Estimator Variance at a Tertiary Node

We consider the tree with three leaves 1, 2 and 3 having common parent c, whose own parent is the root node 0, all links having the same transmission probability α. For both the ML and explicit estimators, we use the Delta-method to calculate the asymptotic variance. We find that at the second order in $\overline{\alpha}$, the MLE has lower asymptotic variance than the explicit estimator.

ML Estimator: The asymptotic variance is $\nabla \mathcal{H}_k \cdot C^{\mathrm{MLE}} \nabla \mathcal{H}_k$ where $\nabla \mathcal{H}_k = \nabla \mathcal{H}_k(\gamma_c, \gamma_1, \gamma_2, \gamma_3)$ from Section IV-A, and C^{MLE} is the covariance matrix with elements $C_{\eta\zeta}^{\mathrm{MLE}} = \mathsf{Cov}(\widehat{Y}_\eta, \widehat{Y}_\zeta)$ where η, ζ run over the node set $\{c, 1, 2, 3\}$. Using Lemma 1(ii) and $\gamma_1 = \gamma_2 = \gamma_3 = \alpha^2$ and $\gamma_c = \alpha(1 - \overline{\alpha}^3)$, one obtains the variance $v_c^{\mathrm{MLE}} = \overline{\alpha} - \overline{\alpha}^2 + O(\|\overline{\alpha}\|^3)$.

Explicit Estimator: We combine (6), (7) and (8), using Lemma 1(i), $B_c = \alpha^4$, $\gamma_1 = \gamma_2 = \gamma_3 = \alpha^2$ and $\gamma_c = \alpha(1 - \overline{\alpha}^3)$ to obtain the variance $v_c^{\mathrm{E}} = \overline{\alpha} - \overline{\alpha}^2/4 + O(\|\overline{\alpha}\|^3)$.

References

[1] T. Bu, N.G. Duffield, F. LoPresti, D. Towsley. "Network Tomography on General Topology", in *Proc. ACM SIGMETRICS 2002*, June 2002.

[2] R. Cáceres, N.G. Duffield, J.Horowitz D. Towsley. "Multicast-Based Inference of Network Internal Loss Characteristics," *IEEE Trans. on Information Theory*, **45**(7), 2462-2480, 1999.

[3] J. Cao, D. Davis, S. Vander Wiel, B. Yu "Time-varying network tomography: router link data", *J. Amer. Statist. Assoc.* **95**, 1063–1075, 2000.

[4] M. Coates, R. Nowak. "Network loss inference using unicast end-to-end measurement", *Proc. ITC Conf. IP Traffic, Modeling and Management*, Sept. 2000, 28.1-28.9.

[5] N.G. Duffield, J. Horowitz, D. Towsley, W. Wei, T. Friedman, "Multicast-based loss inference with missing data", *IEEE Journal on Selected Areas in Communications, 2002*, **20**(4), 700-713, 2002.

[6] N.G. Duffield, J. Horowitz, F. LoPresti, D. Towsley. "Multicast Topology Inference from Measured End-to-End Loss," *IEEE Transactions on Information Theory*, **48**(1), 26-45, 2002.

[7] F. Lo Presti, N.G. Duffield, J. Horowitz, D. Towsley. "Multicast -Based Inference of Network-Internal Delay Distributions," *IEEE/ACM Transactions on Networking*, **10**(6), 761-775, December 2002.

[8] R. Prasad, C. Dovrolis, M. Murray, kc claffy. "Bandwidth estimation: metrics, measurement techniques, and tools," *IEEE Network*, **17**(6), 27-35, 2003.

[9] M. Rabbat, R. Nowak, M. J. Coates. "Multiple Source, Multiple Destination Network Tomography," *Proc. IEEE Infocom'04*, Hong Kong, March 2004.

[10] V.J. Ribeiro, R.H. Riedi, J. Navratil, L. Cottrell, R.G. Baraniuk. "pathChirp: Efficient Available Bandwidth Estimation for Network Paths," *Proc. Workshop on Passive and Active Measurement PAM2003*, March 2003.

[11] M.J. Schervish, "Theory of Statistics", Springer, New York, 1995.

[12] Y. Tsang, M. J. Coates, R. Nowak. "Network Delay Tomography", *IEEE Trans. Signal Processing*, **51**(8), 2125-2136, Aug. 2003.

[13] Y. Vardi, "Network Tomography: estimating source-destination traffic intensities from link data," *J. Am. Statist. Assoc.*, 91: 365–377, 1996.

[14] Y. Zhang, M.Roughan, N.G. Duffield, A. Greenberg, "Fast Accurate Computation of Large-Scale IP Traffic Matrices from Link Loads", in *Proc. ACM Sigmetrics 2003*, June 2003.