Maximum Likelihood Estimation for Multiple-Source Loss Tomography with Network Coding

Pegah Sattari[1], Athina Markopoulou[1], Christina Fragouli[2]

[1]EECS Department, University of California, Irvine, CA, USA

[2]School of Computer and Communication Sciences, EPFL, Lausanne, Switzerland

E-mail: psattari@uci.edu, athina@uci.edu, christina.fragouli@epfl.ch

Abstract—**Loss tomography aims at inferring the loss rate of links in a network from end-to-end measurements. Previous work in [1] has developed optimal maximum likelihood estimators (MLEs) for link loss rates in a single-source multicast tree. However, only sub-optimal algorithms have been developed for multiple-source loss tomography [2]–[5]. In this paper, we revisit multiple-source loss tomography in tree networks with multicast and network coding capabilities, and we provide, for the first time, low-complexity MLEs for the link loss rates. We also derive the rate of convergence of the estimators.**

Keywords—**loss tomography; network coding; maximum likelihood estimation**

I. INTRODUCTION

The successful design, control and management of networks often requires knowledge about internal network characteristics. Network monitoring can be done either by monitoring every link and node of interest or by using end-to-end measurements. The second class of monitoring problems, which aims at inferring internal network characteristics from measurements at the edge of the network, is known as network tomography [6]. In this work, we are particularly interested in loss tomography, *i.e.,* at inferring the loss rates of individual links using end-to-end probes. One of the earliest and most influential woks in loss tomography is [1], which developed MLEs of link loss rates in a single-source multicast tree. Subsequent work provided sub-optimal algorithms, but not MLEs, for loss tomography with multiple sources [2]–[5].

It has been shown that network coding can enhance tomography, for example, using active probes to infer loss [4], [5], [7]; or using passive inference to infer topology or failures [8]–[10]. The key intuition in loss tomography [5] is that network coding increases the information per probe since one coded probe observes multiple paths. At the same time, network coding reduces the bandwidth needed to cover a general graph by making exactly one probe traverse each link. Another key observation that we exploit in this paper is the symmetry between a multicast tree and a reverse multicast tree with network coding. In [5], we presented sub-optimal algorithms to infer the link loss rates in networks with network coding.

In this paper, we extend [1] and we leverage network coding to derive low complexity MLEs in trees with multiple sources. We assume that the tree topology is known and we use an active approach. Sequences of probes are sent and collected between a set of sources and a set of receivers at the edge of the network and the observations at the receivers are used for inference. Intermediate nodes perform unicast, multicast and simple coding operations, which need to be set-up once, fixed for all probes sent and be known for inference. Therefore, our approach requires more support from the network than traditional tomography, for the benefit of more accurate/efficient estimation. We develop, for the first time, efficient MLEs for all link loss rates simultaneously. We also provide their rate of convergence as the number of measurements increases.

The rest of the paper is organized as follows. Section II reviews the related work. Section III describes the problem statement. Section IV solves the likelihood equation for special cases. Section V reduces the general problem to those special cases. Section VI provides the MLEs, the rate of convergence, and an illustrative example. Section VII concludes the paper. Section VIII provides the proofs of theorems.

II. RELATED WORK

Within the large body of work on *network tomography* [6], the most relevant to this paper is the work by Caceres *et al.* [1]. It considers a single multicast tree with a known topology and infers the link loss rates from the observations at the receivers, by developing a low-complexity algorithm to compute the MLE. Follow-up work proposed a sub-optimal approach for general graphs by covering the graph with multiple multicast trees and using an EM algorithm [2]. Other approaches have been developed for unicast probes [11]–[13], and joint inference of topology and loss rates [14].

In [7], we leveraged network coding to improve *active* network monitoring. In [3], we studied sub-optimal methods for link loss estimation in multiple-source tree topologies. In [4], [5], we extended the approach to general graphs. The following papers use random network coding for *passive* network tomography. [8] passively distinguishes among failure patterns. [9], [10] provide algorithms for passive topology estimation and error localization. [15] uses subspace nesting structures to discover local bottlenecks in peer-to-peer systems.

In this paper, we follow an active approach for loss tomography. We build on the MLE for a multicast tree developed in [1], referred to as MINC, and we extend it to multiple-source trees with multicast and network coding capabilities.
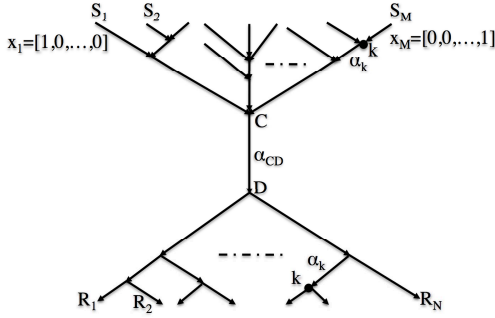
Fig. 1. A directed tree with multicast and network coding capabilities.

III. MODEL AND FRAMEWORK

A. Topology and Node Operation

Logical Tree. We consider a tree topology, like the one depicted in Fig. 1, $\mathcal{T} = (V, L)$ consisting of V nodes and L directed links. M leaf nodes, shown on top of the tree, act as sources of probe packets. The remaining N leaves, shown at the bottom of the tree, act as receivers. An intermediate node is either a coding point (with multiple incoming links and one outgoing link) or a branching point (with one incoming link and multiple outgoing links). For each node j, we denote the set of its parents (nodes with a link outgoing to j) by $f(j)$ and the set of its children (nodes with a link coming from j) by $d(j)$. The source nodes $S = \{S_1, ..., S_M\}$ have no parent and the receivers $R = \{R_1, ..., R_N\}$ have no children. $\mathcal{T}, S, R$ are considered known and fixed throughout the measurements.

We note that the tree topology we consider has the property that all coding points are located above all branching points, which is a mild assumption[1]. It is also a logical[2] tree.

Operation of Sources. Each source S_i sends a probe packet x_i, which is a vector of length M in the following form:

$$x_i = [\overbrace{0, \cdots, 0, \underbrace{1}_{i}, 0, \cdots, 0}^{M}], \quad i = 1, 2, \cdots, M$$

Operation of Intermediate Nodes. Each coding point (bit-wise) XORs all packets it receives from its parents, and forwards the result to its child[3]. Each branching point multicasts the packet it receives from its parent to all its children.

One can see that there will be a node after which $x_1 + x_2 + \cdots + x_M$ flows thought the network. We denote this node by C. Node C is the last coding point in the tree. Node C has P parents $f(C)_1, \cdots, f(C)_P$, and only one child, which we denote by node D. Node D multicasts the packet it receivers from node C to all its Q children $d(D)_1, \cdots, d(D)_Q$.

---

[1]Indeed: starting from an undirected tree, one can choose the sources so as to lead to a directed tree with this property.

[2]This is another common assumption in tomography: intermediate nodes in a logical tree have degree at least three, and in-degree and out-degree at least one. A logical link may consist of several physical links.

[3]We assume that the network is delay-free and all packet arrivals at a coding point are synchronized. Link delays only affect where the probes would meet.

We use the notation that $k < k'$, $k, k' \in V$ when k is a descendant of k', and that $k > k'$ when k is an ancestor of k'. Every node $k > C$ has multiple parents and only one child, while every node $k < D$ has one parent and multiple children. We are going to treat these two sets of nodes differently in the rest of the paper. We name any link of the tree that is above node C by its starting point, and we name any link that is below node D by its end point. In other words, link k denotes a link between nodes (k, j) if $k > C$ and $j > C$, while link k denotes a link between nodes (j, k) if $j < D$ and $k < D$.

B. Loss Model

We model the loss rate of individual links by an i.i.d. Bernoulli process, independent across links. In particular:

- A packet that traverses a link k above node C is lost with probability $\overline{\alpha}_k = 1 - \alpha_k$ and arrives at j with prob. α_k.
- A packet that traverses a link k below node D is lost with prob. $\overline{\alpha}_k = 1 - \alpha_k$ and arrives at k with prob. α_k.
- Finally, we denote the loss rate of link CD by $\overline{\alpha}_{CD}$.

We use the notation $\overline{\alpha} = 1 - \alpha$ for any quantity $0 < \alpha < 1$.

Let X_k denote the packet observed at node k, and let $X = (X_k)$, $k \in V$ denote the set of all X_k's. X_k is a binary vector of length M. Its i^{th} element, $(X_k)_i$, represents the probe packet of source i: $(X_k)_i = 1$ indicates that the probe packet of source i reaches node k, and 0 that it does not. For the sources, $X_{S_i} = x_i$, thus $(X_{S_i})_i = 1$ and $(X_{S_i})_{i'} = 0$, $\forall i' \neq i$. For any node $k \geq C$, if $(X_j)_i = 1$ for j a parent of k, $(X_k)_i = 1$ with probability α_j, and $(X_k)_i = 0$ with probability $\overline{\alpha}_j$, independently for all the parents of k. For any node $k \leq D$, if $X_k = [0, 0, \cdots, 0]$ (the all-zero vector), then $X_j = [0, 0, \cdots, 0]$, for the children j of k (and hence for all descendants of k). If $X_k \neq [0, 0, \cdots, 0]$, then for j a child of k, $X_j = X_k$ with probability α_j, and $X_j = [0, 0, \cdots, 0]$ with probability $\overline{\alpha}_j$, independently for all the children of k.

C. Data, Likelihood, and Inference

In each trial, one probe is dispatched from each source. The outcome of a single trial is a record of whether or not each source probe was received at each receiver, which is the set of vectors X_k observed at receiver $k \in R$. It is denoted by $X_{(R)} = (X_k)_{k \in R}$ and is an element of the space $\Omega = \{[\cdots, 0, 1, \cdots]\}^N$ of all such outcomes. For a given set of link loss probabilities $\alpha = (\alpha_k)_{k \in V \setminus \{C, D\}} \cup \alpha_{CD}$, the distribution of the outcomes $X_{(R)}$ on Ω will be denoted by P_α. The probability mass function for a single outcome $x \in \Omega$ is $p(x; \alpha) = P_\alpha(X_{(R)} = x)$.

We perform n trials. Let $n(x)$ denote the number of probes for which the outcome x is obtained. The probability of n independent observations $x^1, \cdots, x^n$ (each $x^t = (x_k^t)_{k \in R}$) is:

$$p(x^1, \cdots, x^n; \alpha) = \prod_{t=1}^{n} p(x^t; \alpha) = \prod_{x \in \Omega} p(x; \alpha)^{n(x)} \quad (1)$$

Our task is to estimate α using maximum likelihood, from the data $(n(x))_{x \in \Omega}$. We work with the log-likelihood function:

$$\mathcal{L}(\alpha) = \log p(x^1, \cdots, x^n; \alpha) = \sum_{x \in \Omega} n(x) \log p(x; \alpha) \quad (2)$$

The MLE of the loss rates $\breve{\alpha}$ is the α that maximizes $\mathcal{L}(\alpha)$:

$$\breve{\alpha} = \arg\max_{\alpha \in [0,1]^L} \mathcal{L}(\alpha) \quad (3)$$

IV. THE LIKELIHOOD EQUATION AND ITS SOLUTION

Candidates for the MLE are solutions $\hat{\alpha}$ of the *likelihood equation:*

$$\frac{\partial \mathcal{L}}{\partial \alpha_k}(\alpha) = 0, \quad k \in V \quad (4)$$

We define some additional variables to compute the MLEs. For each node $k \geq D$, let $\Omega^r(k)$ be the set of outcomes $x \in \Omega$ such that $(x_a)_j \neq 0$ for at least one source $j \in S$ that is an ancestor of k and for any arbitrary set of receivers $\{a\} \subset R$. Let $\gamma_k^r = \Gamma_k^r(\alpha) = P_\alpha[\Omega^r(k)]$; an estimate of γ_k^r is:

$$\hat{\gamma}_k^r = \sum_{x \in \Omega^r(k)} \hat{p}(x), \quad \text{where} \quad \hat{p}(x) = \frac{n(x)}{n} \quad (5)$$

is the observed proportion of trials with outcome x. γ_k^r shows the probability of the set of outcomes $\Omega^r(k)$ where link k has definitely worked. Link k may have also worked for some other outcomes that are not included in $\Omega^r(k)$. γ_k^r can be directly estimated from the observations at the receivers.

For each node $k \leq C$, we define $\Omega^m(k)$ to be the set of outcomes $x \in \Omega$ such that $x_j \neq [0, 0, \cdots, 0]$ for at least one receiver $j \in R$ which is a descendant of k. Let $\gamma_k^m = \Gamma_k^m(\alpha) = P_\alpha[\Omega^m(k)]$; an estimate of γ_k^m is:

$$\hat{\gamma}_k^m = \sum_{x \in \Omega^m(k)} \hat{p}(x) \quad (6)$$

γ_k^m is the probability of the outcomes $\Omega^m(k)$ in which link k has definitely worked; it can be directly estimated from the observations. Our goal is to compute $\hat{\alpha}$ from $\hat{\gamma} = (\hat{\gamma}_k^r \cup \hat{\gamma}_k^m)_{k \in V}$.

A. Special Cases

1) Multicast Tree (MINC): If $M = 1$, the general model turns into a multicast tree with a single source, which is the case considered in [1]. We represent the source node by $0 \in V$. Each node j other than the source node, has one parent $f(j)$, and a set $d(j)$ of children. We denote the link loss rates by $\overline{\alpha}_k$, where k is the end point. We simply assume that $\alpha_0 = 1$.

The outcome of each trial is $X_{(R)} = (X_k)_{k \in R}$, where each X_k is a single binary value (instead of a binary vector of length M in the general case), corresponding to whether the source probe is observed at each receiver $k \in R$ or not. The state space of the observations $X_{(R)}$ is $\Omega = \{0, 1\}^N$. We say that a link k is at level $l(k)$ if there is a chain of $l(k)$ ancestors $k < f(k) < f^2(k) \cdots < f^{l(k)}(k) = 0$ leading back to 0.

Only $\Omega^m(k)$ is used for each node k in the multicast tree; it is the set of outcomes $x \in \Omega$ where $x_j = 1$ for at least one receiver $j \in R$ that is a descendant of k. γ_k^m is like before.

The MLE for the multicast tree is computed in [1]. Let $A_k^m = \prod_{i=0}^{l(k)} \alpha_{f^i(k)}$ show the probability that the path from the source to node k works, which we denote by $P(Y_{0 \to k} = 1)$. Its estimate $\hat{A}_k^m$ is as follows. For the source node, $\hat{A}_0^m = 1$, for the leaf nodes $k \in R$, $\hat{A}_k^m = \hat{\gamma}_k^m$, and for all other nodes $k \in V \backslash \{0, R\}$, $\hat{A}_k^m$ is the unique solution in $(0, 1]$ of:

$$1 - \frac{\hat{\gamma}_k^m}{\hat{A}_k^m} = \prod_{j \in d(k)} (1 - \frac{\hat{\gamma}_j^m}{\hat{A}_k^m}) \quad (7)$$

$\hat{\alpha}_k$ is then computed from $\hat{\gamma}_k^m$, $\hat{\alpha} = \Gamma^{m-1}(\hat{\gamma^m})$, as follows:

$$\hat{\alpha}_k = \frac{\hat{A}_k^m}{\hat{A}_{f(k)}^m}, \quad k \in V \backslash \{0\} \quad (\hat{\alpha}_0 = 1) \quad (8)$$

We refer to Eq.(8) as MINC in the rest of the paper.

Note. Eq.(7) is obtained from the following relations, after some computations in [1], which we repeat here for completeness. Let $\beta_k^m = P[\Omega^m(k)|X_{f(k)} = 1]$ denote the conditional probability of $\Omega^m(k)$ given that $f(k)$ has observed something. Failure can be due to either $\overline{\alpha}_k$, or all paths towards destinations failing. The β_k^m obey the recursion:

$$\overline{\beta}_k^m = \overline{\alpha}_k + \alpha_k \prod_{j \in d(k)} \overline{\beta}_j^m, \quad k \in V \backslash R \quad (9)$$

$$\beta_k^m = \alpha_k, \quad k \in R \quad (10)$$

Eq.(7) then follows from the relation between α and γ^m:

$$\gamma_k^m = \beta_k^m \prod_{i=1}^{l(k)} \alpha_{f^i(k)} \quad (11)$$

2) Reverse Multicast Tree (RMINC): If $N = 1$, the general model turns into a reverse multicast tree with a single receiver, which we denote by $0 \in V$. Each node j other than 0 has one child $d(j)$, and a set $f(j)$ of parents. We denote link loss rates by $\overline{\alpha}_k$, where k is the starting point. We assume that $\alpha_0 = 1$.

The outcome of each trial, X_R, is a binary vector of length M. Each of its elements, $(X_R)_i$, represents whether the probe of source i is observed at the receiver or not. The state space of the observations X_R is $\Omega = \{0, 1\}^M$. We say that a link k is at level $l(k)$ if there is a chain of $l(k)$ descendants $k > d(k) > d^2(k) \cdots > d^{l(k)}(k) = 0$ leading down to the receiver.

Only $\Omega^r(k)$ is used for each node k in the reverse multicast tree; it is the set of outcomes $x \in \Omega$ where $x_j = 1$ for at least one source $j \in S$ that is an ancestor of k. γ_k^r is like before.

The MLE for the reverse multicast tree is similar to the multicast tree. Let $A_k^r = \prod_{i=0}^{l(k)} \alpha_{d^i(k)}$ show the probability that the path from node k to the receiver node works, which we denote by $P(Y_{k \to 0} = 1)$. For the receiver node, $\hat{A}_0^r = 1$, for the source nodes $k \in S$, $\hat{A}_k^r = \hat{\gamma}_k^r$, and for all other nodes $k \in V \backslash \{S, 0\}$, $\hat{A}_k^r$ is the unique solution in $(0, 1]$ of:

$$1 - \frac{\hat{\gamma}_k^r}{\hat{A}_k^r} = \prod_{j \in f(k)} (1 - \frac{\hat{\gamma}_j^r}{\hat{A}_k^r}) \quad (12)$$

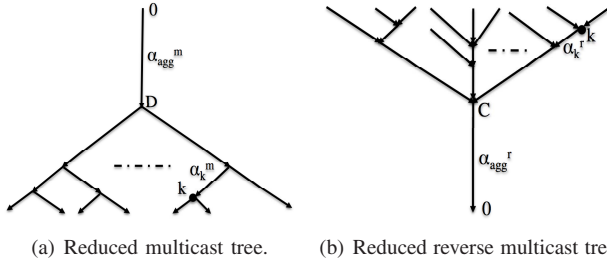(a) Reduced multicast tree.　　　(b) Reduced reverse multicast tree.

Fig. 2.　Reducing the tree in Fig. 1 to multicast and reverse multicast trees.

We can then compute $\hat{\alpha}_k$ from $\hat{\gamma}_k^r$, *i.e.,* $\hat{\alpha} = \Gamma^{r-1}(\hat{\gamma}^r)$:

$$\hat{\alpha}_k = \frac{\hat{A}_k^r}{\hat{A}_{d(k)}^r}, \quad k \in V \setminus \{0\} \quad (\hat{\alpha}_0 = 1) \tag{13}$$

We refer to Eq.(13) as RMINC in the rest of the paper.

Note. Eq.(12) results from the following relations. Let $\beta_k^r = P[\Omega^r(k)|Y_{d(k)\to 0} = 1]$ denote the conditional prob. of $\Omega^r(k)$ given that the path from $d(k)$ to the receiver works. We have:

$$\overline{\beta}_k^r = \overline{\alpha}_k + \alpha_k \prod_{j \in f(k)} \overline{\beta}_j^r, \quad k \in V \setminus S \tag{14}$$

$$\beta_k^r = \alpha_k, \quad k \in S \tag{15}$$

$$\gamma_k^r = \beta_k^r \prod_{i=1}^{l(k)} \alpha_{d^i(k)} \tag{16}$$

3) Comparison of MINC and RMINC: The reader will notice that the MLE for the multicast tree and the reverse multicast tree have the same functional form. This is a special case of the "reversibility" property, first observed in [4], [5]. Indeed, there is a 1-1 correspondence between the observable outcomes in the two cases; furthermore the corresponding outcomes have the same probability, as a function of α_k's, thus leading to the same MLE.

V. Maximum Likelihood Estimation of Loss Rates

We now present how to "reduce" the original tree to a multicast and to a reverse multicast tree, and how to estimate α_{CD}. These intermediate results are then used in the MLE algorithm in Section VI.

A. Reductions

1) Reduction to a Multicast Tree (m): If we take the upper part of the original tree in Fig. 1 and consider it as an aggregate link, we obtain the reduced multicast tree in Fig. 2(a). The aggregate link agg^m summarizes the operation of all links above C and link CD. Node D receives a packet if at least one path from the sources to C works and link CD works.

More formally, we map the outcomes $x \in \Omega$ of the original tree to the outcomes x^m of the multicast tree, as follows. Each x is a set of N binary vectors, each of length M, while each x^m is a single binary vector of length N. Any outcome x^m is obtained by taking a set of outcomes $\{x\}$, in all of

Algorithm 1 Computing the MLE of all link loss rates in the original tree topology of Fig. 1.

1: **for all** links k, where $k < D$ **do**
2:　　Reduce the original tree to a multicast tree. Use MINC [1] (Eq.(8)) to compute the MLEs $\hat{\alpha}_k^m$ and $\hat{\alpha}_{agg}^m$.
3:　　Let $\hat{\alpha}_k = \hat{\alpha}_k^m$.
4: **end for**
5: **for all** links k, where $k > C$ **do**
6:　　Reduce the original tree to a reverse multicast tree. Use RMINC (Eq.(13)) to compute the MLEs $\hat{\alpha}_k^r$ and $\hat{\alpha}_{agg}^r$.
7:　　Let $\hat{\alpha}_k = \hat{\alpha}_k^r$.
8: **end for**
9: Use Eq.(20) to compute the MLE $\hat{\alpha}_{CD}$.

---

which the same receivers have observed all-zero vectors and the same receivers have observed non-zero vectors, and by replacing each non-zero vector (that may contain any of the source probes $x_1, x_2, \cdots, x_M$) by value 1, and by replacing each all-zero vector by value 0. *I.e.:*

$$\sum_{x_{R_t} \neq [0,0,\cdots,0], x_{R_{t'}} = [0,0,\cdots,0]} n(x) = n^m(x^m),$$

$$x_{R_t}^m = 1, x_{R_{t'}}^m = 0, \ t, t' \in \{1, \cdots, N\}, \ t \neq t' \tag{17}$$

If the original tree has link loss rates α and an associated probability distribution of outcomes P_α, then the multicast tree is defined with parameters α^m and P_α^m, such that:

$$\alpha_k^m = \alpha_k, k < D, \quad \alpha_{agg}^m = \alpha_{CD}(1 - \prod_{i=1}^{P} \overline{\beta}_{f(C)_i}^r) \tag{18}$$

P_α^m can be directly calculated from P_α, since each event in P_α^m is the union of a disjoint subset of events in P_α and has probability equal to the sum of probabilities of those events.

2) Reduction to a Reverse Multicast Tree (r): Similarly, if we consider the lower part of the original tree in Fig. 1 as an aggregate link, we obtain the reduced reverse multicast tree in Fig. 2(b), with parameters α^r and P_α^r, such that:

$$\alpha_k^r = \alpha_k, k > C, \quad \alpha_{agg}^r = \alpha_{CD}(1 - \prod_{j=1}^{Q} \overline{\beta}_{d(D)_j}^m) \tag{19}$$

3) The Relation Between the Two Reduced Trees:

Lemma 5.1: We have that: $\hat{\gamma}_C^r = \hat{\gamma}_D^m = 1 - \hat{p}([0, 0, \cdots, 0])$.

The proof directly results from the definition of γ_D^m in the reduced multicast and γ_C^r in the reduced reverse multicast tree.
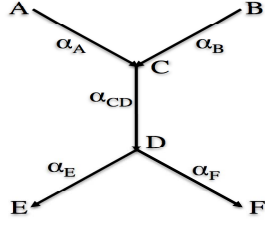
B. Estimating α_{CD}

The MLE of α_{CD} can be obtained from:

$$\hat{\alpha}_{CD} = \frac{\hat{A}_C^r \cdot \hat{A}_D^m}{\hat{\gamma}_C^r} = \frac{\hat{A}_C^r \cdot \hat{A}_D^m}{\hat{\gamma}_D^m} \tag{20}$$

The proof can be found in Section VIII.

(a) An illustrative example: the 5-link tree topology.

#	Original (5-link) tree			Reduced multicast tree			Reduced reverse multicast tree	
	E	F	Prob.	E	F	P_α^m	EF	P_α^r
1	-	-	$\hat{p}_0$	0	0	$\hat{p}_0$	[0,0]	$\hat{p}_0$
2	x_1	-	$\hat{p}_1$				[1,0]	$\hat{p}_1+\hat{p}_4+\hat{p}_7$
3	x_2	-	$\hat{p}_2$	1	0	$\hat{p}_1+\hat{p}_2+\hat{p}_3$	[0,1]	$\hat{p}_2+\hat{p}_5+\hat{p}_8$
4	$x_1 \oplus x_2$	-	$\hat{p}_3$				[1,1]	$\hat{p}_3+\hat{p}_6+\hat{p}_9$
5	-	x_1	$\hat{p}_4$				[1,0]	$\hat{p}_1+\hat{p}_4+\hat{p}_7$
6	-	x_2	$\hat{p}_5$	0	1	$\hat{p}_4+\hat{p}_5+\hat{p}_6$	[0,1]	$\hat{p}_2+\hat{p}_5+\hat{p}_8$
7	-	$x_1 \oplus x_2$	$\hat{p}_6$				[1,1]	$\hat{p}_3+\hat{p}_6+\hat{p}_9$
8	x_1	x_1	$\hat{p}_7$				[1,0]	$\hat{p}_1+\hat{p}_4+\hat{p}_7$
9	x_2	x_2	$\hat{p}_8$	1	1	$\hat{p}_7+\hat{p}_8+\hat{p}_9$	[0,1]	$\hat{p}_2+\hat{p}_5+\hat{p}_8$
10	$x_1 \oplus x_2$	$x_1 \oplus x_2$	$\hat{p}_9$				[1,1]	$\hat{p}_3+\hat{p}_6+\hat{p}_9$

(b) All possible observations and their probabilities in the 5-link tree topology as well as their corresponding observations and probabilities in the reduced multicast and reverse multicast trees.

Fig. 3. The example 5-link tree topology (a) and the table of all observations in the original tree and its reduced multicast and reverse multicast trees (b).

$$\mathcal{I}^{-1}(\alpha) = \begin{pmatrix} \frac{\alpha_A\overline{\alpha}_A}{\alpha_B\alpha_{CD}(\alpha_E+\alpha_F-\alpha_E\alpha_F)} & \frac{\overline{\alpha}_A\overline{\alpha}_B}{\alpha_{CD}(\alpha_E+\alpha_F-\alpha_E\alpha_F)} & \frac{-\overline{\alpha}_A\overline{\alpha}_B}{\alpha_B(\alpha_E+\alpha_F-\alpha_E\alpha_F)} & 0 & 0 \\ \frac{\overline{\alpha}_A\overline{\alpha}_B}{\alpha_{CD}(\alpha_E+\alpha_F-\alpha_E\alpha_F)} & \frac{\alpha_B\overline{\alpha}_B}{\alpha_A\alpha_{CD}(\alpha_E+\alpha_F-\alpha_E\alpha_F)} & \frac{-\overline{\alpha}_A\overline{\alpha}_B}{\alpha_A(\alpha_E+\alpha_F-\alpha_E\alpha_F)} & 0 & 0 \\ \frac{-\overline{\alpha}_A\overline{\alpha}_B}{\alpha_B(\alpha_E+\alpha_F-\alpha_E\alpha_F)} & \frac{-\overline{\alpha}_A\overline{\alpha}_B}{\alpha_A(\alpha_E+\alpha_F-\alpha_E\alpha_F)} & \mathcal{I}_{33}^{-1}(\alpha) & \frac{-\overline{\alpha}_E\overline{\alpha}_F}{\alpha_F(\alpha_A+\alpha_B-\alpha_A\alpha_B)} & \frac{-\overline{\alpha}_E\overline{\alpha}_F}{\alpha_E(\alpha_A+\alpha_B-\alpha_A\alpha_B)} \\ 0 & 0 & \frac{-\overline{\alpha}_E\overline{\alpha}_F}{\alpha_F(\alpha_A+\alpha_B-\alpha_A\alpha_B)} & \frac{\alpha_E\overline{\alpha}_E}{\alpha_{CD}\alpha_F(\alpha_A+\alpha_B-\alpha_A\alpha_B)} & \frac{\overline{\alpha}_E\overline{\alpha}_F}{\alpha_{CD}(\alpha_A+\alpha_B-\alpha_A\alpha_B)} \\ 0 & 0 & \frac{-\overline{\alpha}_E\overline{\alpha}_F}{\alpha_E(\alpha_A+\alpha_B-\alpha_A\alpha_B)} & \frac{\overline{\alpha}_E\overline{\alpha}_F}{\alpha_{CD}(\alpha_A+\alpha_B-\alpha_A\alpha_B)} & \frac{\alpha_F\overline{\alpha}_F}{\alpha_{CD}\alpha_E(\alpha_A+\alpha_B-\alpha_A\alpha_B)} \end{pmatrix}$$

$$\mathcal{I}_{33}^{-1}(\alpha) = \frac{1}{\alpha_A\alpha_B\alpha_E\alpha_F(-\alpha_A\overline{\alpha}_B-\alpha_B)(-\alpha_E\overline{\alpha}_F-\alpha_F)}(-\alpha_{CD}(-\alpha_B\overline{\alpha}_B\alpha_E\alpha_F-\alpha_A^2\overline{\alpha}_B\alpha_E(-1+\alpha_B(2+\alpha_{CD}(-\alpha_E\overline{\alpha}_F-\alpha_F)))\alpha_F$$
$$+ \alpha_A(-\alpha_E\alpha_F+\alpha_B^2\alpha_E\alpha_F(-3+\alpha_{CD}(\alpha_E+\alpha_F-\alpha_E\alpha_F))+\alpha_B(-\alpha_F\overline{\alpha}_F+\alpha_E(-1+7\alpha_F-3\alpha_F^2)+\alpha_E^2(1-3\alpha_F+2\alpha_F^2)))))$$

Fig. 4. The inverse of the Fisher information matrix for the confidence intervals in Eq.(21). The order of the coordinates is $\alpha_A, \alpha_B, \alpha_{CD}, \alpha_E, \alpha_F$.

VI. THE ANALYSIS OF THE MLE

A. MLE Algorithm

Algorithm 1 proceeds in the following steps: (i) it computes $\hat{\alpha}_k$ for any link k below node D from the reduced multicast tree using Eq.(8); (ii) it computes $\hat{\alpha}_k$ for any link k above node C from the reduced reverse multicast tree using Eq.(13); and (iii) it computes $\hat{\alpha}_{CD}$ from Eq.(20). These are indeed the MLEs of all link loss rates, $\hat{\alpha}$, for the tree of Fig. 1.

Theorem 6.1: The estimates computed by Algorithm 1 are the MLEs of the link loss rates in the original tree in Fig. 1.

The proof of Theorem 6.1 relies on the following lemmas, whose proofs are given in Section VIII.

Lemma 6.2: The solutions of the likelihood equations of the original tree and the reduced multicast tree are related via: (i) $\hat{\alpha}_k = \hat{\alpha}_k^m$, $k < D$; and (ii) $\hat{\alpha}_{CD} = \hat{\alpha}_{agg}^m/(1 - \prod_{i=1}^{P} \overline{\beta}_{f(C)_i}^r)$.

Lemma 6.3: The solutions of the likelihood equations of the original tree and the reverse multicast tree are related via: (i) $\hat{\alpha}_k = \hat{\alpha}_k^r$, $k > C$; and (ii) $\hat{\alpha}_{CD} = \hat{\alpha}_{agg}^r/(1 - \prod_{j=1}^{Q} \overline{\beta}_{d(D)_j}^m)$. We note that the likelihood functions of the original tree and the reduced multicast (or reverse multicast) tree are different. What the aforementioned lemmas establish is that they are maximized for the same values of their common variables.

B. Complexity

Algorithm 1 is very efficient. In the first two steps, it calls MINC and RMINC. MINC (and thus RMINC) is known to be efficient by exploiting the hierarchy of the tree topology to factorize the probability distribution and recursively compute the estimates. The computation at each node is at worst proportional to the depth of the tree [1]. The last step, $\hat{\alpha}_{CD}$, uses the estimates $\hat{A}_k, \hat{\gamma}_k$ already computed in the first 2 steps.

C. Rate of Convergence of the MLE

We now provide the rate of convergence of $\hat{\alpha}$ to the true value α. The Fisher information matrix at α based on $X_{(R)}$ is obtained from $\mathcal{I}_{jk}(\alpha) = -E\frac{\partial^2 \mathcal{L}}{\partial \alpha_j \partial \alpha_k}(\alpha)$ [1]. We have that:

Theorem 6.4: $\mathcal{I}(\alpha)$ is non-singular, and as $n \longrightarrow \infty$, $\sqrt{n}(\hat{\alpha}-\alpha)$ converges in distribution to $\mathcal{N}(0, \mathcal{I}^{-1}(\alpha))$.

The proof follows from the asymptotic properties of MLEs [1], [16]. Therefore, asymptotically for large n, with prob. $1-\delta$ (for $1-\delta$ confidence interval), $\hat{\alpha}_k$ lies between the points:

$$\alpha_k \pm z_{\delta/2}\sqrt{\frac{\mathcal{I}_{kk}^{-1}(\alpha)}{n}} \qquad (21)$$

D. Example: the 5-link tree

We illustrate our results through the 5-link tree in Fig. 3(a).

Maximum Likelihood Estimator. Sources A, B send $x_1 = [1, 0], x_2 = [0, 1]$ respectively. Ω consists of ten outcomes shown in Fig. 3(b). The table also shows the corresponding outcomes of the reduced trees. From Eq.(5), Eq.(6), we have:

$$\hat{\gamma}_A^r = \hat{p}_1 + \hat{p}_3 + \hat{p}_4 + \hat{p}_6 + \hat{p}_7 + \hat{p}_9$$

$$\hat{\gamma}_B^r = \hat{p}_2 + \hat{p}_3 + \hat{p}_5 + \hat{p}_6 + \hat{p}_8 + \hat{p}_9$$

$$\hat{\gamma}_C^r = \hat{\gamma}_D^m = \hat{p}_1 + \hat{p}_2 + \hat{p}_3 + \hat{p}_4 + \hat{p}_5 + \hat{p}_6 + \hat{p}_7 + \hat{p}_8 + \hat{p}_9 = 1 - \hat{p}_0$$

$$\hat{\gamma}_E^m = \hat{p}_1 + \hat{p}_2 + \hat{p}_3 + \hat{p}_7 + \hat{p}_8 + \hat{p}_9$$

$$\hat{\gamma}_F^m = \hat{p}_4 + \hat{p}_5 + \hat{p}_6 + \hat{p}_7 + \hat{p}_8 + \hat{p}_9$$

We then solve Eq.(7), Eq.(12) for $\hat{A}_k$, and we find $\hat{\alpha}_A, \hat{\alpha}_B$ from Eq.(13), $\hat{\alpha}_E, \hat{\alpha}_F$ from Eq.(8), and $\hat{\alpha}_{CD}$ from Eq.(20):

$$\hat{\alpha}_A = \frac{\hat{\gamma}_A^r + \hat{\gamma}_B^r - \hat{\gamma}_C^r}{\hat{\gamma}_B^r} \quad , \quad \hat{\alpha}_B = \frac{\hat{\gamma}_A^r + \hat{\gamma}_B^r - \hat{\gamma}_C^r}{\hat{\gamma}_A^r} \quad (22)$$

$$\hat{\alpha}_E = \frac{\hat{\gamma}_E^m + \hat{\gamma}_F^m - \hat{\gamma}_D^m}{\hat{\gamma}_F^m} \quad , \quad \hat{\alpha}_F = \frac{\hat{\gamma}_E^m + \hat{\gamma}_F^m - \hat{\gamma}_D^m}{\hat{\gamma}_E^m} \quad (23)$$

$$\hat{\alpha}_{CD} = \frac{\hat{\gamma}_A^r \hat{\gamma}_B^r \hat{\gamma}_E^m \hat{\gamma}_F^m}{\hat{\gamma}_D^m (\hat{\gamma}_A^r + \hat{\gamma}_B^r - \hat{\gamma}_C^r)(\hat{\gamma}_E^m + \hat{\gamma}_F^m - \hat{\gamma}_D^m)} \quad (24)$$

Confidence Intervals. Fig. 4 shows $\mathcal{I}^{-1}(\alpha)$ for the confidence intervals in Eq.(21). The confidence intervals for parameters $\hat{\alpha}$ can be obtained by inserting Eq.(22), Eq.(23), and Eq.(24) into Fig. 4.

VII. CONCLUSION

In this paper, we proposed a low complexity algorithm to compute MLEs of link loss rates in multiple-source tree networks with multicast and network coding capabilities. So far, MLEs have only been computed for single-source multicast trees [1], while sub-optimal algorithms were used for multiple-source loss tomography.

VIII. PROOFS OF THEOREMS

Estimating α_{CD}: We denote the outcomes in which link CD works by x_{CD}; the outcomes where at least one of the upstream paths to C works by x_{up}; the outcomes where at least one of the downstream paths after D works by x_{dn}. For the intersection of any two of these outcomes, *e.g.*, x_{up} and x_{dn}, we use the notation $x_{up,dn}$. The independence of link loss rates indicates that x_{up}, x_{dn}, x_{CD} are independent. Thus:

$$\hat{\alpha}_{CD} = \hat{p}(x_{CD}) = \hat{p}(x_{CD}|x_{up,dn}) = \frac{\hat{p}(x_{CD} \& x_{up,dn})}{\hat{p}(x_{up})\hat{p}(x_{dn})} \quad (25)$$

The numerator equals $1 - \hat{p}([0,...,0]) = \hat{\gamma}_C^r = \hat{\gamma}_D^m$. Also:

$$\hat{p}(x_{dn}) = \hat{p}(x_{dn}|x_{up,CD}) = \hat{p}(x_{dn}|X_D \neq [0,...,0])$$
$$= 1 - \hat{p}(x_{dn}^c|X_D \neq [0,...,0]) = 1 - \prod_{j=1}^{Q} \overline{\beta}_{d(D)_j}^m \quad (26)$$

We can derive a similar expression for $\hat{p}(x_{up})$. Therefore:

$$\hat{\alpha}_{CD} = \frac{1 - \hat{p}([0,0,\cdots,0])}{(1 - \prod_{i=1}^{P} \overline{\beta}_{f(C)_i}^r)(1 - \prod_{j=1}^{Q} \overline{\beta}_{d(D)_j}^m)} \quad (27)$$

By writing Eq.(9) for $\overline{\beta}_D^m$ in Fig. 2(a), and by writing Eq.(14) for $\overline{\beta}_C^r$ in Fig. 2(b), we conclude that:

$$1 - \prod_{j=1}^{Q} \overline{\beta}_{d(D)_j}^m = \frac{\beta_D^m}{\alpha_{agg}^m} = \frac{\gamma_D^m}{A_D^m} \;,\; 1 - \prod_{i=1}^{P} \overline{\beta}_{f(C)_i}^r = \frac{\beta_C^r}{\alpha_{agg}^r} = \frac{\gamma_C^r}{A_C^r} \quad (28)$$

Eq.(20) follows from replacing these results into Eq.(27). ∎

Proof of Lemma 6.2: In [1], it is shown that the likelihood function of the reduced multicast tree in Fig. 2(a), $\mathcal{L}^m(\alpha^m)$, can be written as the sum of three distinct parts in which the derivative $\partial \log p^m(x^m)/\partial \alpha_k^m$ is constant. These parts are $\Omega^m(k)$, the $\Omega^m(f^i(k))\backslash\Omega^m(f^{i-1}(k))$, which we represent by Ω_2^m for simplicity, for $i = 1, 2, \cdots, l(k)$, and $(\Omega^m(0))^c$. The derivative in these parts is $\frac{1}{\alpha_k^m}$, $\frac{1}{\overline{\beta}_{f^{i-1}(k)}^m} \frac{\partial \overline{\beta}_{f^{i-1}(k)}^m}{\partial \alpha_k^m}$ and $\frac{1}{\overline{\beta}_0^m} \frac{\partial \overline{\beta}_0^m}{\partial \alpha_k^m}$, respectively. Thus, the likelihood equation can be written as:

$$\frac{\partial \mathcal{L}^m}{\partial \alpha_k^m} = \frac{1}{\alpha_k^m} \sum_{x^m \in \Omega^m(k)} n^m(x^m)$$
$$+ \sum_{i=1}^{l(k)} \{ \frac{1}{\overline{\beta}_{f^{i-1}(k)}^m} \frac{\partial \overline{\beta}_{f^{i-1}(k)}^m}{\partial \alpha_k^m} \sum_{x^m \in \Omega_2^m} n^m(x^m) \} \quad (29)$$
$$+ \frac{1}{\overline{\beta}_0^m} \frac{\partial \overline{\beta}_0^m}{\partial \alpha_k^m} \sum_{x^m \in (\Omega^m(0))^c} n^m(x^m)$$

Similarly, we can split the likelihood function of the original tree, $\mathcal{L}(\alpha)$, into three parts in which $\partial \log p(x)/\partial \alpha_k$ is constant. These parts will be similar to those of a multicast tree, only with $\Omega^m(k)$ as defined for the original tree in Section IV, and with $l(k)$ representing the number of ancestors of node k up to node C (instead of 0 in the multicast tree). $\partial \log p(x)/\partial \alpha_k$ over these parts is also similar to the multicast tree, *i.e.*, $\frac{1}{\alpha_k}$, $\frac{1}{\overline{\beta}_{f^{i-1}(k)}^m} \frac{\partial \overline{\beta}_{f^{i-1}(k)}^m}{\partial \alpha_k}$ and $\frac{1}{\overline{\beta}_C^m} \frac{\partial \overline{\beta}_C^m}{\partial \alpha_k}$. Thus, we have:

$$\frac{\partial \mathcal{L}}{\partial \alpha_k} = \frac{1}{\alpha_k} \sum_{x \in \Omega^m(k)} n(x)$$
$$+ \sum_{i=1}^{l(k)} \{ \frac{1}{\overline{\beta}_{f^{i-1}(k)}^m} \frac{\partial \overline{\beta}_{f^{i-1}(k)}^m}{\partial \alpha_k} \sum_{x \in \Omega_2^m} n(x) \} \quad (30)$$
$$+ \frac{1}{\overline{\beta}_C^m} \frac{\partial \overline{\beta}_C^m}{\partial \alpha_k} \sum_{x \in (\Omega^m(C))^c} n(x)$$

(i) $\hat{\alpha}_k^m$ vs. $\hat{\alpha}_k$, $k < D$. We first compare the solutions $\hat{\alpha}_k^m$ of Eq.(29) and $\hat{\alpha}_k$ of Eq.(30) for $k < D$. From Eq.(17), we have that:

$$\sum_{x \in \Omega^m(k)} n(x) = \sum_{x^m \in \Omega^m(k)} n^m(x^m) \quad (31)$$

$$\sum_{x \in \Omega_2^m} n(x) = \sum_{x^m \in \Omega_2^m} n^m(x^m) \tag{32}$$

$$\sum_{x \in (\Omega^m(C))^c} n(x) = \sum_{x^m \in (\Omega^m(0))^c} n^m(x^m) \tag{33}$$

Therefore, for any link k located below node D, we have:

$$\frac{\partial \mathcal{L}^m}{\partial \alpha_k^m} = \frac{\partial \mathcal{L}}{\partial \alpha_k} \Longrightarrow \hat{\alpha}_k^m = \hat{\alpha}_k, \quad k < D \tag{34}$$

(ii) $\hat{\alpha}_{agg}^m$ *vs.* $\hat{\alpha}_{CD}$. For α_{agg}^m and α_{CD}, Eq.(29) and Eq.(30) consist of only the first and the last terms. We have that:

$$\frac{\partial \mathcal{L}^m}{\partial \alpha_{agg}^m} = \frac{1}{\alpha_{agg}^m} \sum_{\Omega^m(D)} n^m(x^m) + \frac{1}{\overline{\beta}_0^m} \frac{\partial \overline{\beta}_0^m}{\partial \alpha_{agg}^m} \sum_{(\Omega^m(0))^c} n^m(x^m) \tag{35}$$

$$\frac{\partial \mathcal{L}}{\partial \alpha_{CD}} = \frac{1}{\alpha_{CD}} \sum_{x \in \Omega^m(D)} n(x) + \frac{1}{\overline{\beta}_C^m} \frac{\partial \overline{\beta}_C^m}{\partial \alpha_{CD}} \sum_{x \in (\Omega^m(C))^c} n(x) \tag{36}$$

Thus, $\frac{\partial \mathcal{L}^m}{\partial \alpha_{agg}^m} \neq \frac{\partial \mathcal{L}}{\partial \alpha_{CD}}$, but the definition of $\overline{\beta}_k^m$ indicates:

$$\overline{\beta}_0^m = 1 - \alpha_{agg}^m (1 - \prod_{j=1}^Q \overline{\beta}_{d(D)_j}^m) \tag{37}$$

$$\overline{\beta}_C^m = 1 - (1 - \prod_{i=1}^P \overline{\beta}_{f(C)_i}^r) \alpha_{CD} (1 - \prod_{j=1}^Q \overline{\beta}_{d(D)_j}^m) \tag{38}$$

From Eq.(31), Eq.(33), Eq.(37), and Eq.(38), we find out that the solutions $\hat{\alpha}_{agg}^m$ of Eq.(35) and $\hat{\alpha}_{CD}$ of Eq.(36) are related via:

$$\hat{\alpha}_{CD} = \frac{\hat{\alpha}_{agg}^m}{1 - \prod_{i=1}^P \overline{\beta}_{f(C)_i}^r} \tag{39}$$

∎

Proof of Lemma 6.3: is similar to the proof above.

∎

Proof of Theorem 6.1: In [1], it is shown that $\hat{\alpha}_k^m$ in Eq.(8) are the MLE of the multicast tree. Therefore, $\hat{\alpha}_k = \hat{\alpha}_k^m$, $k < D$, are also the MLE of the corresponding links in the original tree. In addition, by following the same approach as in [1], one can show that $\hat{\alpha}_k = \hat{\alpha}_k^r$, $k > C$, are also the MLE of the corresponding links in the original tree. For $\hat{\alpha}_{CD}$, since $\hat{\alpha}_{agg}^m = \hat{A}_D^m$ and using Eq.(28), one can obtain Eq.(20) from Eq.(39). Therefore, Eq.(20) is a solution of $\frac{\partial \mathcal{L}}{\partial \alpha_{CD}} = 0$. Furthermore, from Eq.(36) and Eq.(38), we have that:

$$\frac{\partial^2 \mathcal{L}}{\partial \alpha_{CD}^2} = \frac{-1}{\alpha_{CD}^2} \sum_{x \in \Omega^m(D)} n(x) - \frac{1}{\overline{\beta}_C^{m^2}} (\frac{\partial \overline{\beta}_C^m}{\partial \alpha_{CD}})^2 \sum_{x \in (\Omega^m(C))^c} n(x) \tag{40}$$

This is always negative. Therefore, $\mathcal{L}$ is concave in α_{CD} and Eq.(20) is the unique solution of the likelihood equation.

This solution is also in the desired range $[0,1)$, because from Eq.(27) we have that:

$$\hat{\alpha}_{CD} > 0 \Longleftrightarrow \hat{p}([0,0,\cdots,0]) < 1$$

i.e., not all packets are lost, which is the default assumption in tomography: no inference can be made without data. Also:

$$\hat{\alpha}_{CD} < 1 \Longleftrightarrow 1 - \hat{p}([0,\cdots,0]) < (1 - \prod_{i=1}^P \overline{\beta}_{f(C)_i}^r)(1 - \prod_{j=1}^Q \overline{\beta}_{d(D)_j}^m)$$

This is asymptotically true for $\alpha_{CD} > 0$, because as $n \longrightarrow \infty$, the percentage of packets that are not lost approaches the probability $(1 - \prod_{i=1}^P \overline{\beta}_{f(C)_i}^r) \alpha_{CD} (1 - \prod_{j=1}^Q \overline{\beta}_{d(D)_j}^m)$.

Therefore, Eq.(20) is the MLE of α_{CD} in the original tree.

∎

REFERENCES

[1] R. Caceres, N. G. Duffield, J. Horowitz, and D. Towsley, "Multicast-based inference of network-internal loss characteristics," *IEEE Transactions on Information Theory*, vol. 45, no. 7, pp. 2462–2480, Nov. 1999.
[2] T. Bu, N. Duffield, F. Presti, and D. Towsley, "Network tomography on general topologies," *ACM SIGMETRICS Performance Evaluation Review*, vol. 30, no. 1, 2002.
[3] C. Fragouli, A. Markopoulou, R. Srinivasan, and S. Diggavi, "Network monitoring: it depends on your points of view," in *ITA Workshop*, San Diego, CA, 2007.
[4] M. Gjoka, C. Fragouli, P. Sattari, and A. Markopoulou, "Loss tomography in general topologies with network coding," in *IEEE globecom*, Washington DC, Nov. 2007.
[5] A. Markopoulou, C. Fragouli, and M. Gjoka, "A network coding approach to loss tomography," http://arxiv.org/abs/1005.4769, 2010.
[6] R. Castro, M. Coates, G. Liang, R. Nowak, and B. Yu, "Network tomography: Recent developments," *Statistical Science*, vol. 19, no. 3, pp. 499–517, 2004.
[7] C. Fragouli and A. Markopoulou, "A network coding approach to overlay network monitoring," in *Allerton Conference*, Monticello, IL, 2005.
[8] T. Ho, B. Leong, Y.-H. Chang, Y. Wen, and R. Koetter, "Network monitoring in multicast networks using network coding," in *IEEE ISIT*, Adelaide, Australia, Sept. 2005, pp. 1977–1981.
[9] G. Sharma, S. Jaggi, and B. Dey, "Network tomography via network coding," in *ITA Workshop*, San Diego, CA, Feb. 2008.
[10] H. Yao, S. Jaggi, and M. Chen, "Passive network tomography for erroneous networks: A network coding approach," http://arxiv.org/abs/0908.0711, 2010.
[11] N. Duffield, F. L. Presti, V. Paxson, and D. Towsley, "Inferring link loss using striped unicast probes," in *IEEE INFOCOM*, Anchorage, AK, Apr. 2001.
[12] M. Coates and R. Nowak, "Network loss inference using unicast end-to-end measurement," in *ITC Conference on IP Traffic, Modeling and Management*, Monterey, CA, Sep. 2000.
[13] K. Harfoush, A. Bestavros, and J. Byers, "Robust identification of shared losses using end-to-end unicast probes," in *IEEE ICNP*, Osaka, Japan, Nov. 2000.
[14] M. Rabbat, R. Nowak, and M. Coates, "Multiple source multiple destination network tomography," in *IEEE INFOCOM*, Mar. 2004.
[15] M. Jafarisiavoshani, C. Fragouli, S. Diggavi, and C. Gkantsidis, "Bottleneck discovery and overlay management in network coded peer-to-peer systems," in *ACM SIGCOMM INM Workshop*, Japan, Aug. 2007.
[16] J. A. Rice, *Mathematical Statistics and Data Analysis*. Belmont, California: Duxbury Press, 1995.