MAXIMUM LIKELIHOOD IDENTIFICATION OF NETWORK TOPOLOGY FROM END-TO-END MEASUREMENTS

Rui Castro, Mark Coates, and Robert Nowak

ECE Dept., Rice University, Houston, TX 77005

1. INTRODUCTION

One of the predominant schools of thought in networking today is that monitoring and control of large scale networks is only practical at the edge. With intelligent and adaptive elements at the edge of the network, core devices can function as simple, robust routers. However, the effectiveness of edge-based control can be significantly enhanced by information about the internal network state. If the core is endowed with minimal monitoring and data collection capabilities, then methods for inferring state information from edge-based traffic measurements are of great interest. One of the most fundamental components of the state is the routing topology. The focus of this paper is a new Maximum Likelihood approach to topology identification that makes use only of measurements performed between host computers and requires no special support (e.g., ICMP responses) from internal routers.

2. PROBLEM STATEMENT

Consider a sender transmitting information through a network to a set of receivers denoted by $\mathcal{R}$. Assume that the routes from the sender to the receivers are fixed. The problem we address is the identification of the network topology based on end-to-end measurements that measure the degree of correlation between receivers [1, 2, 3, 4, 5]. With this limited information, it is only possible to identify the so-called "logical topology" defined by the branching points between paths to different receivers. This corresponds to a tree-structured topology with the sender at the root and the receivers at the leaves as depicted in Figure 1(a).

Associate a metric $\gamma_{i,j}$ with each pair of receivers $i, j \in \mathcal{R}$. The value of $\gamma_{i,j}$ is related to the extent of the shared portion of the paths to i and j. More specifically, the metrics must have the following

Monotonicity Property: Let i, j and k be any three receivers and let p_i, p_j, and p_k denote the paths from the sender to each. If p_i shares more links with p_j than with p_k, then $\gamma_{i,j} > \gamma_{i,k}$.

The property can be used to identify the underlying topology. For example, referring to Figure 1, the metric $\gamma_{18,19}$ will be strictly greater than $\gamma_{i,19}$ for all $i \in \mathcal{R}/\{18, 19\}$, revealing that receivers 18 and 19 have a common parent in the logical tree. The property can be exploited in this manner to devise a simple and effective bottom-up merging algorithms that identify the full, logical topology [1, 2, 3, 4].

Metrics possessing the Monotonicity Property can be estimated from a number of different end-to-end measurements including counts of losses, counts of zero delay events (utilization), and delay correlations [1, 2, 3, 4]. These estimated metrics, denoted $\{x_{i,j}\}$, can be interpreted as statistics derived from repeated measurements. Randomness in network conditions can lead to variability in the measurements and hence the estimated metrics. Most of the previous work in this area does not incorporate the variability of the estimated metrics (which can also be assessed directly from the measurements) into the identification process. We claim that this variability can have a major impact on the performance of topology identification algorithms.

Several methods have been proposed for topology identification in both unicast and multicast settings [1, 2, 3, 4], but all have a very similar structure. The DBT algorithm proposed in [2] is a representative example. The algorithm is a recursive selection and merging/aggregation process that generates a binary tree from the bottom-up (receivers to sender). In this paper, we describe a new algorithm that specifically addresses the uncertainty in estimated metrics, providing substantial performance improvements in certain cases.

3. NEW CONTRIBUTIONS

To address the issue of metric variability and uncertainty we pose topology identification as a maximum likelihood estimation (MLE) problem. The MLE approach is selected for its well known asymptotic optimality properties (under mild conditions which are applicable in this problem the MLE is the best unbiased estimator as the number of measurements tends to infinity). There are three main contributions.

1. Likelihood formulation: The estimated metrics $\mathbf{x} \equiv \{x_{i,j}\}$ can be interpreted as observations of the true metric values $\boldsymbol{\gamma} \equiv \{\gamma_{i,j}\}$ contaminated by some randomness or noise. We model this contamination probabilistically. The estimated metrics are randomly distributed according to a density (whose precise form depends on the contamination model) that is parameterized by the underlying topology $\mathcal{T}$ and the set of true metric values, written as $p(\mathbf{x}|\boldsymbol{\gamma}, \mathcal{T})$. The $\mathbf{x}$ are observed and henced fixed, and when $p(\mathbf{x}|\boldsymbol{\gamma}, \mathcal{T})$ is viewed as a function of $\mathcal{T}$ and $\boldsymbol{\gamma}$ it is called the likelihood of $\mathcal{T}$ and $\boldsymbol{\gamma}$. The maximum likelihood tree is given by

$$\mathcal{T}^* = \max_{\mathcal{T} \in \mathcal{F}} \max_{\boldsymbol{\gamma} \in \mathcal{G}} p(\mathbf{x}|\boldsymbol{\gamma}, \mathcal{T}), \qquad (1)$$

where $\mathcal{F}$ denotes the *forest* of all possible tree topologies connecting the sender to the receivers and $\mathcal{G}$ denotes the set of all metrics satisfying the Monotonicity Property.

To illustrate our approach we will focus on one type of metric. In earlier work we proposed a metric based on delay differences [5]. Each estimated metric is modeled as

$$x_{i,j} \sim \mathcal{N}(\gamma_{i,j}, \sigma^2_{i,j}), \qquad (2)$$

where $\sigma^2_{i,j}$ is measured variability of the $x_{i,j}$ and $\mathcal{N}(\gamma, \sigma^2)$ denotes the Gaussian density with mean γ and variance σ^2. The motivation for the model above is that the average of several independent measurements tends to a Gaussian distribution according to the Central Limit Theorem. The likelihood function in this case is a product of Gaussian densities of this form, one factor for each pair of receivers.

2. Characterization of the Maximum Likelihood Tree:
The maximizations involved in (1) are quite formidable. We are not aware of any method for computing the global maximum except

Supported by the National Science Foundation, grant no. MIP–9701692, the Office of Naval Research, grant no. N00014-00-1-0390, and the Army Research Office, grant no. DAAD19-99-1-0290.

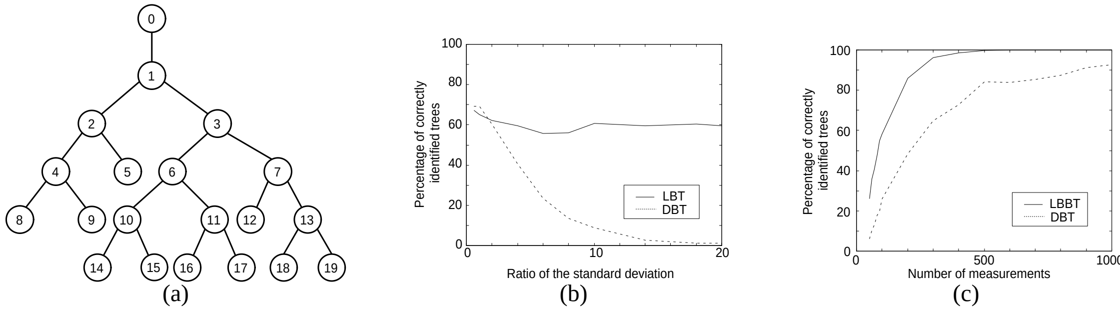

Fig. 1. (a) A binary logical tree topology. (b) Comparison of LBT and DBT algorithms based on simulated delay difference measurements [5] using the topology in (a). The plots depict the percentage of correctly identified trees versus the ratio of largest and smallest standard deviations of the estimated metrics. 1000 independent trials were computed for each ratio. (c) Percentage of correctly identified trees versus the number of measurements when the ratio of the largest to the smallest standard deviations is equal to 6.

by a brute force examination of each tree in the forest. Consider a network with N receivers. A very loose lower bound on the size of the forest $\mathcal{F}$ is $N!/2$. For example, if $N = 10$ then there are more than 1.8×10^6 trees in the forest. This explosion of the search space precludes the brute force approach in all but very small (logical) networks. Moreover, the inner maximization is non-trivial because it involves a constrained optimization over $\mathcal{G}$. The following theorem establishes a key property of the MLE solution that leads to an important simplification.

Theorem: *Let $\mathcal{T}^*$ denote the solution to (1). Then*

$$\arg\max_{\boldsymbol{\gamma}\in\mathbb{R}^N} p(\mathbf{x}|\boldsymbol{\gamma},\mathcal{T}^*) \;=\; \arg\max_{\boldsymbol{\gamma}\in\mathcal{G}} p(\mathbf{x}|\boldsymbol{\gamma},\mathcal{T}^*).$$

The theorem is proved by contradiction. Suppose that some tree $\mathcal{T}$ is the MLE solution. If the argument of the unconstrained maximization over all real-valued $\boldsymbol{\gamma}$ does not belong to the monotonic set $\mathcal{G}$, then this can be shown to contradict the fact that $\mathcal{T}$ is the MLE. We omit the full proof here due to space limitations. The theorem shows that it is unnecessary to perform the constrained optimization. For each tree, we can compute the unconstrained optimization, which simply involves calculating a weighted sums of metrics, and check if the resulting maximizer lies in the set $\mathcal{G}$. If not, we know that the tree cannot be the MLE solution we seek. Avoiding the constrained optimization can significantly reduce the complexity of searches through the forest.

3. Likelihood-based Binary Tree (LBT) Algorithm:

While determining the globally optimal tree is prohibitive in most cases, the theorem above motivates a new, improved (but suboptimal) bottom-up algorithm based on our likelihood formulation of the problem. The new approach, called the LBT algorithm, is based on two properties: (A) the identified tree should satisfy the condition stated in the Theorem above; (B) the bottom-up process should preserve the likelihood structure.

Before describing the algorithm, we point out an important issue arising when considering the variability associated with the estimated metrics. The estimated values $x_{i,j}$ and $x_{j,i}$ may be based on different measurements (e.g., as in [5]) and consequently they may not be equal nor have equal variances (even though they both are contaminated versions of the same underlying true metric). To apply the DBT algorithm directly in such cases, one may be tempted to simply average $x_{i,j}$ and $x_{j,i}$ to form a symmetric quantity. This, however, can lead to violations of Property (A).

The LBT algorithm follows a similar strategy to the DBT algorithm with the following key modifications. The receiver pairs are selected by finding i, j such that

$$\left(\frac{x_{i,j}}{\sigma_{i,j}^2} + \frac{x_{j,i}}{\sigma_{j,i}^2}\right) / \left(\frac{1}{\sigma_{i,j}^2} + \frac{1}{\sigma_{j,i}^2}\right)$$

is maximized. This guarantees that the resulting tree will satisfy Property (A). This selection criterion is employed in subsequent steps of the algorithm with aggregated metrics. The aggregation step requires another critical modification. To ensure Property (B) the aggregation also depends on the variances associated with the estimated metrics. This is accomplished as follows. Suppose that nodes i, j are selected for merging/aggregation and let k denote the new parent node. The new (aggregated) metric value relating k to any other node l is given by

$$x_{k,l} = \left(\frac{x_{l,i}}{\sigma_{l,i}^2} + \frac{x_{l,j}}{\sigma_{l,j}^2}\right) / \left(\frac{1}{\sigma_{l,i}^2} + \frac{1}{\sigma_{l,j}^2}\right),$$

and similarly for $x_{l,k}$. Also the variances for the new metrics must be updated in a similar fashion

$$\sigma_{k,l}^2 = \left(\sigma_{k,i}^2\sigma_{k,j}^2\right) / \left(\sigma_{l,i}^2 + \sigma_{l,j}^2\right).$$

This guarantees that the maximum likelihood tree derived from aggregated metrics is a subtree of full maximum likelihood tree. Aside from these two simple, yet crucial modifications, the algorithm operates in the same manner as the DBT algorithm. The performance improvements provided by these modifications are examined in the simulation experiment described in Figure 1(b)-(c). The results show that as the disparity between the variances of the estimated metrics increases, the performance is significantly improved by the proposed modifications.

4. REFERENCES

[1] S. Ratnasamy and S. McCanne, "Inference of multicast routing trees and bottleneck bandwidths using end-to-end measurements," in *Proceedings of IEEE INFOCOM 1999*, New York, NY, March 1999.

[2] N.G. Duffield, J. Horowitz, F. Lo Presti, and D. Towsley, "Multicast topology inference from end-to-end measurements," in *ITC Seminar on IP Traffic, Measurement and Modelling*, Monterey, CA, Sep. 2000.

[3] N.G. Duffield, J. Horowitz, F. Lo Presti, and D. Towsley, "Multicast topology inference from measured end-to-end loss," to appear in *IEEE Trans. Information Theory*, 2002.

[4] A. Bestavros, J. Byers, and K. Harfoush, "Inference and labeling of metric-induced network topologies," Tech. Rep. BUCS-2001-010, Computer Science Department, Boston University, June 2001.

[5] R. Castro, M. Coates, M. Gadhiok, R. King, R. Nowak, E. Rombokas, and Y. Tsang, "Maximum likelihood network topology identification from edge-based unicast measurements," Tech. Rep. TREE-0107, Rice University, Aug. 2001.