Statistical Inference for Internal Link Parameters in a Network[1]

R. Cáceres, N.G. Duffield — AT&T Labs-Research
J. Horowitz — University of Massachusetts
F. LoPresti — AT&T Labs-Research
D. Towsley[2] — University of Massachusetts

Abstract — **In this paper we describe the use of end-to-end loss measurements of multicast traffic to infer the network-internal loss characteristics. For the case that the tree spanning a multicast source and its receivers is known, we present statistically rigorous techniques for estimating loss rates on the paths between branch points based on losses observed by multicast receivers. Last, we show how these techniques canalso be used to identify the structure of the tree spanning a multicast source and its receivers.**

I. Introduction

In this paper, we describe an approach for characterizing link-level loss behavior in a network based on end-to-end *multicast* probe traffic, that is, traffic between a source and several receivers. Multicast introduces dependence in the losses measured at the receivers which can be used to infer the loss rates at links in the routing tree spanning the sender and receivers.

The results presented here are as follows. For the case that the topology is known, we present maximum likelihood estimators (MLEs) of the link loss rates based on a sequence of n independent probes, under the assumption that losses are independent on different links. These estimators have been shown to be strongly consistent as the number of probes tends to infinity. In the second part of the paper, we show how these estimators can be adapted to infer the topology of the multicast tree.

II. Description of the Basic Model

Let $\mathcal{T} = (V, L)$ denote the multicast tree, consisting of the set of nodes V, including the source and receivers, and the set of links L, which are ordered pairs (j, k) of nodes, indicating a (directed) link from j to k. The set of *children* of node j is denoted by $d(j)$. Let $0 \in V$ denote the root of the tree, the source of the probes and $R \subset V$ the set of receivers (leaf nodes in the tree). For each node j, other than the root 0, there is a unique node $f(j)$, the *parent* of j, such that $j \in d(f(j))$. Last, we assume that the root node has one child, the receivers have no children, and all other nodes have two or more children.

A probe packet is sent down the tree, starting at the root. If it reaches a node j, copies of the packet are produced and sent down the links toward each child of j. As a packet traverses a link to k, it is lost with probability $\overline{\alpha}_k = 1 - \alpha_k$ and arrives at k with probability α_k. We shall use the notation $\overline{\alpha} = 1 - \alpha$ for any probability α The losses on different links are assumed to be independent and to occur with the probabilities $\overline{\alpha}_k$ as described.

The passage of probes down the treeis described by a stochastic process $X = (X_k)_{k\in V}$ where each X_k equals 0 or 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). If $X_k = 1$, then for j a child of k, $X_j = 1$ with probability α_j, and $X_j = 0$ with probability $\overline{\alpha}_j$, independently for all the children of k. We write $\alpha_0 = 1$ to simplify expressions concerning the α_k.

III. Results on Maximum Likelihood Estimators

The outcome of a probe is a record of whether or not a copy of it was received at each receiver. Expressed in terms of the process X, the outcome is a configuration $X_{(R)} = (X_k)_{k\in R}$ of zeroes and ones at the receivers (1 = received, 0 = lost). Only the values of X at the receivers are observable; the values at the internal nodes are invisible. The state space of the observations $X_{(R)}$ is thus the set of all such configurations, $\Omega = \{0,1\}^R$. For a given set of link probabilities $\alpha = (\alpha_k)_{k\in V}$, the distribution of $X_{(R)}$ on Ω will be denoted by $\mathbf{P}_\alpha$. The probability mass function for a single outcome $x \in \Omega$ is $p(x;\alpha) = \mathbf{P}_\alpha(X_{(R)} = x)$.

Let us dispatch n probes, and, for each $x \in \Omega$, let $n(x)$ denote the number of probes for which the outcome x is obtained. The probability of n independent observations $x^1, \dots, x^n$ (with each $x^m = (x_k^m)_{k\in R}$) is then

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

We work with the log-likelihood function

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

In this notation we suppress the dependence of $\mathcal{L}$ on n and $x^1, \dots, x^n$. For each node k, let $\Omega(k)$ be the set of outcomes $x \in \Omega$ such that $x_j = 1$ for at least one receiver $j \in R$ which is a descendant of k, and let $\gamma_k = \gamma_k(\alpha) = \mathbf{P}_\alpha[\Omega(k)]$. An estimate of γ_k is

$$\widehat{\gamma}_k = \sum_{x\in\Omega(k)} \widehat{p}(x), \quad (3)$$

where $\widehat{p}(x) := n(x)/n$ is the observed proportion of trials with outcome x. We will show that α can be calculated from $\gamma = (\gamma_k)_{k\in V}$, and that the MLE

$$\breve{\alpha} = \arg\max_{\alpha\in[0,1]^{\#L}} \mathcal{L}(\alpha) \quad (4)$$

can be calculated in the same manner from the estimates $\widehat{\gamma}$. The relation between α and γ is as follows. We use U to denote the set of nodes other than the root.

Theorem 1 *Let $\mathcal{A} = \{(\alpha_k)_{k\in U} : \alpha_k > 0\}$, and $\mathcal{G} = \{(\gamma_k)_{k\in U} : \gamma_k > 0\,\forall k;\ \gamma_k < \sum_{j\in d(k)} \gamma_j\,\forall k \in U \setminus R\}$. There is a bijection Γ from $\mathcal{A}$ to $\mathcal{G}$ which is differentiable and has a differentiable inverse.*

[1] This work was supported in part by DARPA and the Air Force Research Laboratory under agreement F30602-98-2-0238.

[2] Correponding author: Dept. of Computer Science, Univ. of Massachusetts, Amherst, MA 01003-4610, towsley@cs.umass.edu

0-7803-5268-8/99/$10.00 ©1999 IEEE

The proof of this and other theorems in this section can be found in [1]. Candidates for the MLE are solutions of the *likelihood equation*:

$$\frac{\partial \mathcal{L}}{\partial \alpha_k}(\alpha) = 0, \quad k \in U. \tag{5}$$

Theorem 2 *When $\widehat{\gamma} \in \mathcal{G}$, the likelihood equation has the unique solution $\widehat{\alpha} := \Gamma^{-1}(\widehat{\gamma})$.*

The proof depends on a detailed analysis of $\mathcal{L}$ on the sets $\Omega(k)$.

Let us now restrict our attention to link probabilities $\alpha \in \mathcal{B} = (0,1)^{\#R}$. We have the following result.

Theorem 3

(i) The model is identifiable in $\mathcal{B}$, i.e., $\alpha, \alpha' \in \mathcal{B}$ and $\mathbf{P}_\alpha = \mathbf{P}_{\alpha'}$ together imply $\alpha = \alpha'$. Thus, distinct link probabilities α produce distinct statistical behavior of the $\widehat{\gamma}$ as $n \to \infty$.

(ii) As $n \to \infty$, $\breve{\alpha} \to \alpha$, with $\mathbf{P}_\alpha$- probability 1,

(iii) With probability 1, for sufficiently large n, $\breve{\alpha} = \widehat{\alpha}$)

IV. Topology Inference for Binary Trees

The use of estimate of losses on the common portion of the path between receivers in binary trees has been proposed recently in order to group multicast receivers that share the same set of bottlenecks on the path from the source [5]. In this and the succeeding section we shall investigate the analytic and experimental properties of this algorithm and some generalization that infer general trees (i.e. those which are not necessarily binary).

Before we introduce the algorithm for identifying a binary multicast tree, we generalize the terminology introduced in the preceding section. Define the family of random variables $(Y(S))_{S \in \mathcal{W}}$ recursively by $Y(S) = \max_{S' \in S} Y(S')$ with $Y(j) = X_j$ for $j \in R$. We can relate the internal link probabilities and the distribution of the visible $(X_k)_{k \in R}$ as follows. For $S \in \mathcal{W}$ set $\gamma(S) = \mathbf{E}_\alpha[Y(S)] = \mathbf{P}_\alpha[Y(S) = 1]$, the probability that a probe reaches any receiver descended from S. For $j \in V$ set $A(j) = \alpha_j \alpha_{f(j)} \ldots \alpha_0$, the probability that a probe reaches the node j. It can be shown that these are related by

$$(1 - \gamma(k)/A(k)) = \prod_{j \in d(k)} (1 - \gamma(j)/A(k)). \tag{6}$$

It was shown in Lemma 1 of [1] that $A(k)$ is the unique solution to (6) in $(\gamma(k), 1]$ provided that $\gamma(k) < \sum_{j \in d(k)} \gamma(j)$. (This allows us to extend (6) to give a *definition* of $A(S)$ when S is an arbitrary element of $\mathcal{W}$). Thus the collection of $A(k)$, and hence also the $\alpha_k = A(k)/A(f(k))$, can be determined from the $\gamma(k)$. At a node k with two offspring $\{j, j'\}$, (6) reduces to

$$A(k) = \frac{\gamma(j)\gamma(j)}{\gamma(j) + \gamma(j') - \gamma(k)} = \frac{\mathbf{P}_\alpha[Y(j) = 1]\mathbf{P}_\alpha[Y(j') = 1]}{\mathbf{P}_\alpha[Y(j) = Y(j') = 1]} \tag{7}$$

In order to estimate the $A(k)$ from measurements, n probes are multicast from the source, giving rise to outcomes $\widehat{X}^{(1)}, \ldots, \widehat{X}^{(n)}$. We define subsidiary quantities $\widehat{Y}^{(i)}(S)$ in terms of $\widehat{X}^{(i)}$ for $i = 1, \ldots n$ analogously to the definition of Y in terms of X, and empirical probabilities corresponding to the $\gamma(S)$ by

$$\widehat{\gamma}(S) = n^{-1}\#\{i : \widehat{Y}^{(i)}(S) = 1\}. \tag{8}$$

Input: The set of receivers R and the set of receiver traces $(X_k^i)_{k \in R, i=1,2,\ldots N}$;
$R' := R$, $V := \emptyset$; $L = \emptyset$;
while $|R'| > 1$ **do**
  select $S = \{S_1, S_2\} \subset R'$ that minimizes $\widehat{A}(\cdot)$;
  $V := V \cup \{S\}$; $R' := (R' \setminus S) \cup \{S\}$;
  foreach $S' \in S$ **do**
    $\alpha(S') := A(S)/A(S')$; $L := L \cup \{(S, S')\}$;
  enddo
enddo
Output: loss tree $((V, L), \alpha)$

Figure 1: Binary Loss Tree Classification Algorithm (BLT).

We estimate $A(k)$ by $\widehat{A}(k)$, the solution to (6) that results from using the $\widehat{\gamma}$ in place of γ.

Lemma 1

(i) The model is identifiable, i.e. $P_\alpha = P_{\alpha'}$ implies $\alpha = \alpha'$.

(ii) with probability 1, for sufficiently large n, $\widehat{A}$ is the Maximum Likelihood Estimator of A.

As a consequence of the MLE property, $\widehat{A}$ is (strongly) consistent ($\widehat{A} \overset{n\to\infty}{\longrightarrow} A$ with probability 1), and asymptotically normal ($\sqrt{n}(\widehat{A} - A) \overset{n\to\infty}{\longrightarrow} G$ for some multivariate Gaussian random variable G); see [6].

In what follows we shall work exclusively with **canonical loss trees**. First a loss tree consists of a tree-loss rate combination $(\mathcal{T}, \alpha)$. A loss tree is said to be in canonical form if $\alpha_k < 1, \forall k \in V$ except for $k = 0$. Any tree $(\mathcal{T}, \alpha)$ not in canonical form can be reduced to a loss tree, $(\mathcal{T}', \alpha')$, in canonical form such that the distribution of $(X_k)_{k \in R}$ is the same under the corresponding probabilities $\mathbf{P}_{\mathcal{T},\alpha}$ and $\mathbf{P}_{\mathcal{T}',\alpha'}$. Henceforth, we will only consider canonical loss trees.

In the **binary loss tree** (BLT) classification algorithm, the largest set of nodes with highest common loss is grouped to form a composite node that is identified as their common parent. Grouping proceeds iteratively until all a single node (the root) is formed. A formal specification is given in Figure 1.

Let $P_A^f(\mathcal{T}, \alpha)$ denote the probability that algorithm A fails to correctly identify the tree $\mathcal{T}$.

Theorem 4

(i) A canonical loss tree is identifiable, i.e., $\mathbf{P}_{\mathcal{T},\alpha} = \mathbf{P}_{\mathcal{T}',\alpha'}$ implies $(\mathcal{T}, \alpha) = (\mathcal{T}', \alpha')$.

(ii) For any binary canonical loss tree $(\mathcal{T}, \alpha)$, $P_{BLT}^f \to 0$ as $n \to \infty$.

The proof of this theorem can be found in [3]

V. Inference of General Trees

In this section we describe how the binary loss tree classification algorithm can be extended to treat general trees. The new algorithm uses a parameter $\varepsilon > 0$ in the following manner. For each $\varepsilon > 0$ we define BLTP(ε) (BLT with Pruning) by a two step process: (i) applying BLT to the receiver traces resulting in a binary loss tree $((V, L), \alpha)$; then (ii) for each node k with loss probability $1 - \alpha_k < \varepsilon$, remove the link $(f(k), k)$ from L and identify the endpoints in V.

Theorem 5 *Given a canonical loss tree* $(\mathcal{T}, \alpha)$, *then* $P^f_{BLTP} \to 0$ *as* $n \to \infty$ *provided that* $\varepsilon < \min\{\alpha_k\}$.

VI. Discussion

We have introduced the use of end-to-end measurements of multicast traffic to infer internal link loss probabilities, and have shown that maximum likelihood estimation is feasible when the losses are independent. Simlation results and measurements show this to be accurate, even when the independence assumptions are violated, see [2] and [4] for further details of the simulation and measurement studies.

References

[1] R. Cáceres, N.G. Duffield, J. Horowitz, D. Towsley, "Multicast-based inference of network-internal loss characteristics", Comp. Sci. Tech. Rep. 98-17, University of Massachusetts at Amherst, February 1998. `ftp://gaia.cs.umass.edu/pub/CDHT98:MINC.ps.Z`

[2] R. Cáceres, N.G. Duffield, J. Horowitz, D. Towsley, T. Bu," Multicast-Based Inference of Network-Internal Characteristics: Accuracy of Packet Loss Estimation", Proceedings IEEE Infocom'99, New York, (1999).

[3] R. Cáceres, N. Duffield, J. Horowitz, F. LoPresti, D. Towsley. "Loss-based Inference of Multicast Network Topology", submitted for publication.

[4] R. Cáceres, N. Duffield, S.B. Moon, D. Towsley. "Inferring Link-level Performance from End-to-End Multicast Measurements," submitted for publication.

[5] S. Ratnasamy & S. McCanne, "Inference of Multicast Routing Tree Topologies and Bottleneck Bandwidths using End-to-end Measurements", Proceedings IEEE Infocom'99, New York, (1999).

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