

The following paper posted here is not the official IEEE published version. The final published version of this paper can be found in the Proceedings of the IEEE Consumer Communications and Networking Conference : Consumer networking : closing the digital divide : Las Vegas, Nevada USA, 5-8 January 2004 / Robert S. Fish (general chair):pp.164-168

Copyright © 2004 IEEE.

Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

Analysis on Binary Loss Tree Classification with Hop Count for Multicast Topology Discovery

Hui Tian and Hong Shen
Graduate School of Information Science
Japan Advanced Institute of Science and Technology
1-1 Asahidai, Tatsunokuchi, Ishikawa, 923-1292, Japan

Abstract—The use of multicast inference on end-to-end measurement has recently been proposed as a means to obtain underlying multicast topology. In this paper we analyze the algorithm of the binary loss tree classification with hop count (HBLT). We compare it with the algorithm of binary loss tree classification (BLT) and show that the probability of misclassification of HBLT decreases more quickly than BLT as the number of probing packets increases. The inference accuracy of HBLT is always 1 — the inferred tree is identical to the physical tree — in the case of correct classification, whereas that of BLT is dependent on the shape of the physical tree and inversely proportional to the number of internal nodes with single child. Our analytical result shows that HBLT is superior to BLT not only on time complexity but also on misclassification probability and inference accuracy.

Key words: Multicast, topology inference, hop count, misclassification probability, inference accuracy.

I. INTRODUCTION

For the efficient use of network resources, multicast has become one of the most popular forms of communication. The potential benefits obtained from topological information on multicast distribution trees becomes increasingly important. Knowledge of multicast topology can significantly facilitate resource management, and can be applied to build schemes for loss recovery and congestion control in the context of multicast sessions supporting heterogeneous receivers [1], [4].

There has been much research on multicast topology discovery. Existing approaches can be classified into two types: those based on end-to-end measurements [3], [2], [6], [5], [7], [11] and those requiring the help of intervening network nodes [8], [9].

The key idea underlying the first approach is that receivers sharing common paths on the multicast tree associated with a given source will see correlations in their packet losses. The multicast tree can thus be inferred based on the shared loss statistics for transmitted probe packets. The key advantage of the approach lies in its applicability to inferring multicast trees requiring no support from internal nodes. It permits however, identification of the logical multicast topology rather than the actual physical topology. As discussed in the sequel, a long path with no branches would be identified as a single logical link. With this approach, all single-child nodes are deleted in the inferred logical topology. In practice this may not be an appropriate inference of the actual topology because there may exist many nodes with only one child. The scheme of *binary loss tree classification* (BLT) is a representative of

this approach which provides a good combination of inference accuracy and computational simplicity [7], [10].

The second approach to multicast topology discovery is based on the use of the MTRACE feature currently implemented in the IGMP protocol. The main advantage of this approach is that it provides full information on the multicast topology based on currently available IGMP features. The physical topology including interface addresses of routers can therefore be obtained. This however, shall rely heavily on special services at routers.

Based on algorithm proposed in [5], [7], [10] from end-to-end loss measurement, an improved algorithm, namely *binary loss tree classification with hop count* (HBLT), has recently been proposed in [12] which can identify nodes with both two children and single child, and thus produce the inferred topology much closer to the physical topology than that of BLT. In HBLT, an additional parameter of hop count is incorporated. It doesn't require support from internal nodes which is necessary in the second approach, and can infer multicast topologies which BLT is unable to infer. This paper is concerned with the analysis of misclassification probability and inference accuracy of HBLT.

The paper is organized as follows. In Section II mathematical models of multicast network are introduced. The procedures of BLT and HBLT are described in Section III. Misclassification analysis is given in Section IV. Section V compares the degree of inference accuracy for both algorithms. Section VI concludes the paper.

II. MATHEMATICAL MODELS OF MULTICAST NETWORK

We begin with reviewing the tree and loss models used to formulate the BLT Algorithm and the improved algorithm HBLT. The physical multicast tree is modelled as a tree comprising actual network elements (the nodes) and the communication links that join them.

• Tree model

Let $T = (V, L)$ denote a multicast tree with node set V and link set L . The root node 0 is the source of probes, and $R \subset V$ denotes the set of leaf nodes representing the receivers. A link is said to be internal if neither of its endpoints is the root or a leaf node. Let W denote $V \setminus (\{0, 1\} \cup R)$, where 1 is the child node of 0. Each non-leaf node k has a set of children node $d(k) = \{d_i(k) \mid 1 \leq i \leq n_k\}$, and non-root node k has a parent $P(k)$. The link $(p(k), k) \in L$ can be

simply denoted by link k . If j is descended from k , it could be written $j \prec k$ or $k = p^r(j)$ for r being a positive integer. Let $a(U)$ mean the nearest common ancestor of a subset $U \subset V$, Nodes U are said to be siblings if they have the same parent, i.e., if $f(k) = a(U), \forall k \in U$. The subtree of T rooted at k is denoted by $T(k) = (V(k), L(k))$, and the receiver set $R(k) = R \cap V(k)$. Figure 1 shows an example of multicast tree model.




Fig. 1. A multicast tree model

• Loss model

For each link an independent Bernoulli loss model is assumed: each probe is successfully transmitted across link k with probability α_k . Thus the progress of each probe down the tree is described by an independent copy of a stochastic process $X = (X_k)_{k \in V}$ as follows. $X_0 = 1, X_k = 1$ if the probe packet reaches node $k \in V$ and 0 otherwise. If $X_k = 0, X_j = 0, \forall j \in d(k)$. Otherwise, $P[X_j = 1 | X_k = 1] = \alpha_j$ and $P[X_j = 0 | X_k = 1] = 1 - \alpha_j$. Define $\alpha_0 = 1, \alpha = (\alpha_i)_{i \in V}$. The pair (T, α) is called a loss tree. $P_{T, \alpha}$ denotes the distribution of X on the loss tree (T, α) . If $0 < \alpha_k < 1, \forall k \in V \setminus \{0\}$, the loss tree is said to be a canonical tree. Any tree (T, α) in non-canonical form can be reduced to a canonical tree (T', α') [5]. Henceforth only canonical loss trees are discussed in this paper.

III. ALGORITHMS FOR MULTICAST TOPOLOGY DISCOVERY

In this section, we summarize the procedures of BLT and HBLT. For simplicity only inference algorithms for binary tree inference are concerned here.

A. Binary Loss Tree Classification (BLT)

In BLT, When a probe packet is sent on the path from the root to the nearest ancestor of two receivers, the same link condition is provided. Thus the pair of receivers share the same loss rate. Since the loss rate cannot decrease as the path is extended, the pair of receivers whose shared loss rate is greatest will be siblings. This pair is identified as a pair of siblings and replaced by a composite node that represents their parent. Iterating this procedure shall then reconstruct the binary tree.

For identifying the siblings, the loss distribution function $B(U)$ [5] is applied. When U is a set of siblings, $B(U)$ is

minimized, which implies the sharing loss rate of the set is maximized. BLT works as follows:

- Step1: Find the set of receivers which minimize $B(U)$;
- Step2: Mark the set siblings and deletes them from the set requiring inferred;
- Step3: Substitute a composite node that represents their parent for the siblings and adds it to non-inferred set;
- Step4: Iterate the above procedure in the rest of the set requiring inferred.

In this way the logical multicast tree is reconstructed by observing the probability distribution of the receivers in bottom-up fashion.

In practice, $B(U)$ is estimated by $\hat{B}(U)$.

$$\hat{B}(U) = \frac{\sum_{i=1}^n \hat{X}_{u_1}^{(i)} \sum_{i=1}^n \hat{X}_{u_2}^{(i)}}{n \sum_{i=1}^n \hat{X}_{u_1}^{(i)} \hat{X}_{u_2}^{(i)}} \quad (1)$$

Here $U = \{u_1, u_2\}$, n is the number of probe packets, $(X_k^i)_{k \in R}^{i=1, \dots, n}$ denotes the measured outcomes observed at the receiver arising from each of n probes.

When u_1, u_2 are siblings, $\hat{B}(U)$ will be minimized.

Note that the inferred tree by BLT is a logical tree which is followed by such rules. The logical topology induced by a physical topology T is formed from T after all internal nodes with only one child have been collapsed into their parent recursively (as defined in [10]). Thus a physical tree can be converted to a logical tree. The logical multicast tree has the property that every node has at least two descendants, apart from the root node (which have only one) and the leaf-nodes (which have none). BLT can infer such kind of logical tree in correctly inferring case. But it can never infer the actual physical tree if single-child node exists.

B. Binary Loss Tree Classification with Hop Count (HBLT)

In order to identify those single-child nodes to obtain a multicast topology which is closer to the actual physical tree, HBLT takes hop count into account. It can infer all the internal nodes with different values of hop count, including those nodes with single child. Links terminating at single-child nodes can also be identified.

Though one more parameter is considered, HBLT doesn't add more burden to the multicast network because hop count can be computed by simply reading the TTL values of the probe packets at receivers. For internal nodes, the values of hop count can be easily computed by degression. Therefore application of hop count to topology inference is very efficient.

HBLT infers the multicast topology from the node with the maximum hop count and proceeds in a bottom up fashion. Since real siblings must have the same value of hop count, the set of all the receivers is classified to different sets according to the value of hop count. In the procedure of identifying siblings, only comparison between the sets with the same value of hop count is necessary. It works as follows:

- Step 1: Classify all the receivers nodes to different node sets according to hop count;

Step 2: Identify siblings in the set of nodes with the maximum value of hop count;

Step 3: Substitute them with their parent node, the value of hop count is reduced by 1, and loss distribution is obtained by “OR” operation of the siblings;

Step 4: After the siblings are identified in the node set with certain hop count, repeat the same procedure in the node set with hop count value decreased by 1 until the set with hop count value 1.

The multicast topology is reconstructed in this way. HBLT provides higher performance in inference speed than BLT. As analyzed in [12], the time complexity of HBLT in the worst case is $O(l^2)$, while that of BLT is $O(l^3)$. The expected time complexity of HBLT is estimated at $O(l * \log_2 l)$, which is much better than that of BLT $O(l^3)$.

IV. MISCLASSIFICATION PROBABILITY ANALYSIS

In this section, we describe the model of failure of HBLT, and analyze the probability of misclassification. Since HBLT proceeds by recursively classifying receivers, topology misclassification can be analyzed by looking at how sets of receivers can be misclassified in the estimated topology \hat{T} .

Definition 1: Let (T, α) be a loss tree with $T = (V, L)$, and $(\hat{T}, \hat{\alpha})$ be an inferred loss tree with $\hat{T} = (\hat{V}, \hat{L})$. Denote the receiver set of i by $R_T(i)$ and let $R^T = \cup_i R^T(i)$. If $\forall i \in W = V \setminus (\{0, 1\} \cup R)$, $R_T(i) = R_{\hat{T}}(i)$, the node i is said to be classified correctly.

For a binary tree, the topology is correctly classified if and only if $\forall i \in W$ is correctly classified. So we can study topology misclassification by looking at the misclassification of receivers for each i . For both BLT and HBLT, computing loss distribution of internal nodes is necessary when estimating $B(U)$. Therefore we use the following general formula of loss distribution function, as presented in [5].

$$\begin{aligned} \hat{B}(S_1, S_2) &= \frac{\sum_{i=1}^n \vee_{j \in S_1} Y_j^{(i)} \sum_{i=1}^n \vee_{j \in S_2} Y_j^{(i)}}{n \sum_{i=1}^n (\vee_{j \in S_1} Y_j^{(i)}) \cdot (\vee_{j \in S_2} Y_j^{(i)})} \\ &= \frac{\hat{\gamma}(S_1) \hat{\gamma}(S_2)}{\hat{\gamma}(S_1) + \hat{\gamma}(S_2) - \hat{\gamma}(S_1 \cup S_2)} \end{aligned} \quad (2)$$

Where

$$Y_k^{(i)} = X_k^{(i)} = \vee_{j \in d(k)} Y_j^{(i)} = \vee_{j \in R_T(k)} Y_j^{(i)},$$

S_1, S_2 are two non-empty disjoint subsets of R_T . $Y_k^{(i)}$ is the loss distribution of node k for the i th probe packet, which is the union of all its descendants' loss distribution. For $U \subset V$, $\hat{\gamma}(U) = n^{-1} \sum_{i=1}^n \vee_{j \in U} Y_j^{(i)}$ is the fraction of probe packets that reach some receiver descended from some node in U .

As the number of probe packets n increases, $\hat{B}(S_1, S_2)$ asymptotically approaches to $B(S_1, S_2)$, $\lim_{n \rightarrow \infty} \hat{B}(S_1, S_2) = B(S_1, S_2)$. $B(U)$ is the function of loss distribution at a set of nodes U that is minimized when U is a set of siblings. $\gamma(U) = P[\vee_{k \in U} \vee_{j \in R(k)} X_j = 1]$, which is approximated by $\hat{\gamma}(U)$ in practice.

$$\begin{aligned} B(S_1, S_2) &= \frac{P[\vee_{j \in S_1} X_j = 1] P[\vee_{j \in S_2} X_j = 1]}{P[\vee_{j \in S_1} X_j \cdot \vee_{j \in S_2} X_j = 1]} \\ &= \frac{\gamma(S_1) \gamma(S_2)}{\gamma(S_1) + \gamma(S_2) - \gamma(S_1 \cup S_2)} \end{aligned} \quad (3)$$

Assume that the height of the multicast tree is h , so the maximum value of hop count is h . Each node set with the same hop count value k is denoted by $H_k, 1 \leq k \leq h$. The coherence coefficient of H_k , denoted by ε_k , is defined to be the correlation among the nodes within $H_k, 1 \leq k \leq h$. Coherence coefficients $\varepsilon_1, \varepsilon_2, \dots, \varepsilon_h$ are used to decide whether two nodes are siblings. Let the event G_i describe that HBLT correctly classifies nodes in $R_T(i)$ for some $i \in W$. This happens if every descendent of i finds its sibling in HBLT inference correctly.

Let $D(S_1, S_2, S_3)$ be the difference between $B(S_1, S_2)$ and $B(S_1, S_3)$, and $\hat{D}(S_1, S_2, S_3)$ the difference between $\hat{B}(S_1, S_2)$ and $\hat{B}(S_1, S_3)$. Define a set $S(i)$ as

$$\begin{aligned} S(i) &= \{(S_1, S_2, S_3) : S_1, S_2, S_3 \subset H_k, 1 < k < h, \\ &\quad S_1, S_2 \subset R_T(i), S_3 \subseteq R_T \setminus R_T(i), \\ &\quad S_p \neq \Phi, p = 1, 2, 3, S_p \neq S_q, p \neq q\}. \end{aligned}$$

Lemma 1: A sufficient condition for correct classification of i is that

$$\hat{D}(S_1, S_2, S_3) = \hat{B}(S_1, S_2) - \hat{B}(S_1, S_3) > 0 \quad (4)$$

and

$$\hat{B}(S_1, S_2) < \varepsilon_k \quad (5)$$

for all $(S_1, S_2, S_3) \in S(i)$.

Let $Q(S_1, S_2, S_3)$ be the event that equation (4) and (5) hold, Q_i be the event that $Q(S_1, S_2, S_3)$ holds for all $S_1, S_2, S_3 \in S(i)$. Thus, $Q_i = \bigcap_{(S_1, S_2, S_3) \in S(i)} Q(S_1, S_2, S_3)$. Let \bar{G}_i be the event that i is classified incorrectly, P_i^f be the misclassification probability on node i . Then

$$G_i \supseteq Q_i \quad (6)$$

$$P_i^f = P[\bar{G}_i] \leq \sum_{(S_1, S_2, S_3) \in S(i)} P[\bar{Q}(S_1, S_2, S_3)] \quad (7)$$

Denote by G the multicast tree that is correctly classified, then we have that

$$G \supseteq \bigcap_{i \in W} Q_i = \bigcap_{i \in W} \bigcap_{(S_1, S_2, S_3) \in S(i)} Q(S_1, S_2, S_3) \quad (8)$$

Consequently the misclassification probability of the inferred tree by HBLT is as follows.

$$\begin{aligned} P_{HBLT}^f &= P[\bar{G}] \\ &\leq \sum_{i \in W} \sum_{(S_1, S_2, S_3) \in S(i)} P[\bar{Q}(S_1, S_2, S_3)] \end{aligned} \quad (9)$$

Theorem 1: For each $i \in W$ and all $((S_1, S_2, S_3) \in S(i))$, $\sqrt{n} \cdot (\hat{D}(S_1, S_2, S_3) - D(S_1, S_2, S_3))$ converges in distribution to a Gaussian random variable with mean 0 and

variance $\frac{\sigma^2(S_1, S_2, S_3)}{\xi}$ as the number of probes $n \rightarrow \infty$, where $D(S_1, S_3, S_3) = B(S_1, S_3) - B(S_1, S_2)$, $\xi \geq 2$.

Proof:

From Theorem 10 of [5], we know that by use of BLT, for each $i \in W$ and all $((S_1, S_2, S_3) \in S(i))$, $\sqrt{n} \cdot (\widehat{D}(S_1, S_2, S_3) - D(S_1, S_2, S_3))$ can be described as a standard normal distribution $N(0, \sigma^2(S_1, S_2, S_3))$ as $n \rightarrow \infty$. For every classification of siblings descended from i , S_1, S_2 are selected from subsets of all nodes which are descendent of i , and S_3 is selected from the subsets that don't belong to the subtree rooted at i . In HBLT, only S_1, S_2, S_3 within the subsets with the same value of hop count are sampled, ε_k acts as imaginative node in S_3 , which makes equation (4) hold after replacing S_3 by ε_k . All these (S_1, S_2, S_3) can be known as a sampling from the normal total sets. Assume the number of sampling is ξ , then the mean of sampling is still 0, and the variance is the original variance divided by ξ . ■

By Theorem 1, we can approximate $P[\widehat{Q}(S_1, S_2, S_3)]$ by $\Psi(-\sqrt{n} \cdot \sqrt{\xi} \cdot \frac{D(S_1, S_2, S_3)}{\sigma(S_1, S_2, S_3)})$, Ψ is the cdf of the standard normal distribution. For large n , $P[\widehat{Q}(S_1, S_2, S_3)]$ can be approximated as follows.

$$P[\widehat{Q}(S_1, S_2, S_3)] \approx e^{-(n/2) \cdot \xi \cdot \frac{D^2(S_1, S_2, S_3)}{\sigma^2(S_1, S_2, S_3)}} \quad (10)$$

The logarithmic asymptotic of $P[\widehat{Q}(S_1, S_2, S_3)]$ is the following:

$$\log P[\widehat{Q}(S_1, S_2, S_3)] \sim -\frac{n \cdot \xi \cdot D^2(S_1, S_2, S_3)}{2\sigma^2(S_1, S_2, S_3)} \quad (11)$$

From equation (7) we know P_i^f is dominated by $P_{max}(\widehat{Q}(S_1, S_2, S_3))$. $P_i^f \approx P_{max}(\widehat{Q}(S_1, S_2, S_3))$ when $\min_{(S_1, S_2, S_3) \in S(i)} \frac{D^2(S_1, S_2, S_3)}{\sigma^2(S_1, S_2, S_3)}$. After substituting $D^2(S_1, S_2, S_3)$ and $\sigma^2(S_1, S_2, S_3)$, it was shown from Theorem 10 in [5] that

$$\min_{(S_1, S_2, S_3) \in S(i)} \frac{D^2(S_1, S_2, S_3)}{\sigma^2(S_1, S_2, S_3)} = \hat{\alpha}_i + O(\|\alpha\|^2), \quad (12)$$

where α_i is the loss rate of link $(f(i), i)$, $\hat{\alpha}_i = 1 - \alpha$, $\|\alpha\| = \max_{k \in V} \hat{\alpha}_k \rightarrow 0$. P_i^f can be approximated as:

$$P_i^f \approx e^{-\hat{\alpha}_i \frac{n}{2} \xi} \quad (13)$$

From equation (9) we know P_{HBLT}^f is dominated by $\max_{i \in W} P_i^f$ likewise, that is, when

$$\hat{\alpha}^f = \min_{i \in W} \hat{\alpha}_i \quad (14)$$

holds. Under this condition, logarithm on P_{HBLT}^f is expected to be asymptotically linear on n with negative slope:

$$\frac{\xi}{2} \cdot \hat{\alpha}^f$$

Compared with the asymptotical slope of BLT, $\hat{\alpha}^f/2$, HBLT shows better performance. Therefore, misclassification probability of HBLT decreases more quickly than that of BLT as the number of probe packet increases, as depicted in Figure 2.

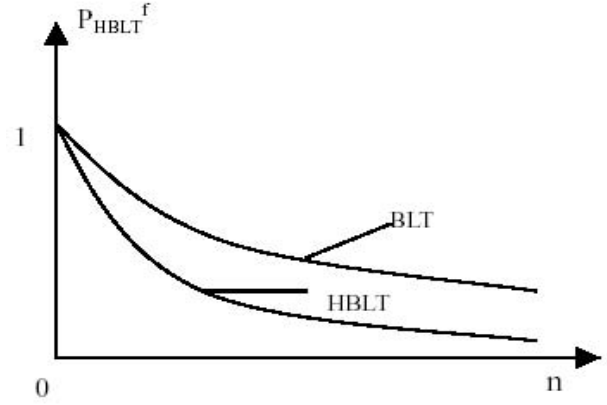


Fig. 2. Misclassification probability comparison between BLT and HBLT

V. ANALYSIS ON INFERENCE ACCURACY

In this section, we compare the inferred trees by BLT and HBLT to analyze the inference accuracies. It is clear that inferred trees for BLT contains only those nodes with two children in the physical tree, whereas that by HBLT contains nodes with both two children and one child in the physical tree. Considering that both algorithms classify correctly, the objected trees are much different from each other. From this view inference accuracy is discussed in this section.




Fig. 3. Logical trees inferred by BLT and HBLT when misclassification probability is 0

Assume that both algorithms classify correctly, the inferred trees are shown in Figure 3, where the left tree is inferred tree by HBLT which is the same as the original physical tree and the right tree is the inferred tree by BLT. It is obvious that the correctly inferred tree by BLT changes a lot in the structure from the original physical tree, whereas the correctly inferred tree by HBLT is the same as the actual physical tree. We analyze their difference in inference accuracy.

Assume that the actual physical tree has n internal nodes. Every internal node has 1 child at the probability of α , and has 2 children at the probability of $1 - \alpha$. Let n_1 be the number of nodes that have 1 child, and n_2 be the number of nodes that have 2 children. Then we have $n = n_1 + n_2$. Denote by

c_i the number of children of internal node i .

$$\begin{cases} P\{c_i = 1\} = \alpha \\ P\{c_i = 2\} = 1 - \alpha \end{cases} \quad (15)$$

$$\begin{aligned} E(n_1) &= \sum_{x=1}^2 P\{c_i = x\} \cdot E[n_1 | c_i = x] \\ &= P\{c_i = 1\} \cdot E[n_1 | c_i = 1] \\ &+ P\{c_i = 2\} \cdot E[n_1 | c_i = 2] \\ &= n \cdot \alpha \end{aligned} \quad (16)$$

When all internal nodes have 1 child, $E[n_1 | c_i = 1]$ is n ; if all internal nodes have 2 children, $E[n_1 | c_i = 2]$ equals 0.

Definition 2: For a multicast tree that has n internal nodes, the inference accuracy of an algorithm is defined as the number of inferred internal nodes divided by the actual total number of internal nodes in the case that the algorithm classifies correctly.

Clearly, the expected number of internal nodes that have single child is $n\alpha$. In BLT all these single-child nodes are deleted. It only infers all those internal nodes that have two children. Henceforth, the BLT's accuracy degree is

$$\frac{n - E(n_1)}{n} = 1 - \alpha \quad (17)$$

Suppose that $\alpha = 1/2$, a half of internal nodes will fail to be identified. The inference accuracy of BLT is $1/2$. As α increases, i.e., the probability of every internal node having single child increases, the inference accuracy of BLT will decrease. If with a large probability every internal node has single child, the inferred tree by BLT will be very different from the original physical tree. The inference accuracy of BLT is thus very low in such case.

However, for HBLT, the inference accuracy is irrelevant to α , so it can identify all the internal nodes regardless of the number of children they have. The inference accuracy is therefore always equal to 1.

VI. CONCLUSION

In this paper, we have analyzed misclassification probability and inference accuracy of the recently proposed algorithm HBLT. Comparison between HBLT and BLT is discussed, which shows that HBLT outperforms BLT significantly. The misclassification probability of HBLT decreases more quickly than BLT as the number of probe packets increases. The inference accuracy of HBLT is always 1 and the inferred tree by HBLT is thus identical to the actual physical tree, in the case of correct classification. However, the inference accuracy of BLT is inversely proportional to the number of internal nodes with single child in the multicast tree. The inferred logical tree may differ significantly from the actual physical tree. Our analysis shows that incorporating hop count in the process of topology inference is a promising technique and may benefit various applications on network topology discovery.

ACKNOWLEDGEMENT

This work was supported by Japan Society for the Promotion of Science (JSPS) Grant-in-Aid for Scientific Research under its General Research Scheme (**B**) Grant No. 14380139.

REFERENCES

- [1] B.N.Levine, S.Paul, and J.J.GARCIA-Luna-Aceves. Organizing multicast receivers deterministically according to packet-loss correlation. In *Proceedings of ACM Multimedia'98*, 1998.
- [2] R. Caceres, N. G. Duffield, J. Horowitz, F. Lo Presti, and D. Towsley. Loss based inference of multicast network topology. In *IEEE Conference on Decision and Control, Phoenix, AZ*, 1999.
- [3] Rui Castro, Mark Coates, and Robert Nowak. Maximum likelihood identification from end-to-end measurements. Technical report, DIMACS Workshop on Internet Measurement, Mapping and Modeling, DIMACS Center, Rutgers University, Piscataway, New Jersey, February 2002.
- [4] D.Li and D.R.Cherton. OtersOn-Tree Efficient Recovery using Sub-casting:a reliable multicas protocol. In *Proceedings of IEEE ICNP'98*, 1998.
- [5] N. G. Duffield, J. Horowitz, F. Lo Presti, and D. Towsley. Multicast topology inference from measured end-to-end loss. *IEEE Trans. Information Theory*, 2002.
- [6] N.G. Duffield, J. Horowitz, and F. Lo Presti. Adaptive multicast topology inference. In *IEEE Infocom 2001, Anchorage, Alaska*, April 22-26, 2001.
- [7] N.G. Duffield, J. Jorowitz, F.Lo Presti, and D. Towsley. Multicast topology inference from end-to-end measurements. Technical report, ITC Seminar on IP Traffic Measurement, Modeling and Management, Monterey, CA, Sept. 2000.
- [8] M. Fenner and S. Casner. A traceroute facility for ip multicast. Technical report, IETF draft, draft-ietf-idmr-traceroute-ipm-06.tex, March, 2000.
- [9] Jangwon Lee and Gustavo de Veciana. Resource and topology discovery for ip multicast using a fan-out decrement mechanism. In *Proc. IEEE Inforcom'2001*, 2001.
- [10] F. Lo Presti, N. G. Duffield, J. Horowitz, and D. Towsley. Multicast-based inference of network-internal delay distributions. *IEEE/ACM Transactions on Networking*, Vol. 10, No. 6, Dec. 2002.
- [11] S. Ratanasamy and S. McCanne. Inference of multicast routing tree topologies and bottleneck bandwidths using end-to-end measurements. In *Proc. IEEE Inforcom'99, New York, NY*, 1999.
- [12] Hui Tian and Hong Shen. An improved algorithm of multicast topology inference from end-to-end measurements. In *Proceedings of the 5th Int'l Symp. on High Performance Computing ISHPC-V*, Oct. 2003.