Low Feedback MINC Loss Tomography

Vijay Arya
Advisor: Thierry Turletti
INRIA, Sophia Antipolis, France
{vijay.arya, thierry.turletti}@sophia.inria.fr

Timur Friedman, Rémy Bellino
Laboratoire Lip6-CNRS
Paris, France
timur.friedman@upmc.fr, zeremz@free.fr

Nick Duffield
AT&T Labs–Research
NJ USA
duffield@research.att.com

***Abstract*— We present a simple extension to the MINC loss estimator which can be used to perform loss tomography with less feedback bits per probe. In MINC loss inference, each receiver in the multicast tree reports one bit of feedback per probe. This poses constraints when MINC is used with RTCP and feedback bandwidth must not exceed 5% of data bandwidth. In the Extended MINC loss estimator (EMLE), receivers report feedbacks for groups of w consecutive probes. Each feedback requires only 1 bit per w probes. EMLE leverages the analysis of MINC loss estimator itself and results in the reduction of feedback bits without substantial loss of accuracy.**

I. Introduction

Multicast-based Inference of Network internal Characteristics (MINC) [1] is a method of performing network tomography in which internal characteristics of a network are inferred from end-to-end multicast measurements. MINC can infer internal characteristics of a network that lies under a multicast tree. MINC can infer loss rates and delay distributions of internal network links [2], [3]. To infer loss rates, the source injects a stream of probe packets into the multicast tree. Corresponding to each probe, each receiver reports whether it received the probe packet (`1`) or not (`0`). Using the binary feedbacks collected from all receivers, per link loss rates in the multicast tree are inferred. In this way, one bit of feedback is needed per probe in MINC.

Since dedicated infrastructures to perform large scale measurements are generally complex to deploy, authors of [4] proposed an architecture which couples the process of performing end-to-end multicast measurements with RTP/RTCP [5]. In this method, RTP data packets of a multicast session act as probes and the feedbacks to perform MINC loss inference are piggybacked on RTCP packets. One of the constraints here is that in large multicast groups, receivers are unable to provide one bit feedback per probe since this can cause the feedback bandwidth to exceed 5% of data bandwidth [4]. In this work, we consider the problem of performing loss tomography with less feedback bits. We have designed a simple extension to the MINC loss estimator which uses information from available probes but reports less feedback bits.

II. Extended MINC Loss Estimator (EMLE)

In EMLE, N probe packets are injected from the source of the multicast tree as in MINC. Instead of providing a feedback corresponding to every probe, receivers report feedbacks for windows of w consecutive probes. Each receiver reports only two values - whether it observed 0 losses or more than 0 losses in windows of w consecutive probes. The window size w is constant and common to all receivers. In total, receivers report N/w feedbacks which require N/w bits.

Using these feedbacks, the passage probability of each link in the multicast tree is estimated in the following manner. With usual notation, let V denote the set of all nodes in the logical multicast tree, let $R \subset V$ denote the set of receiver nodes, and let $S \in V$ denote the source node. Let $f(k)$ denote the father of node k and let $d(k)$ denote the set of children of node k. Instead of modeling the passage of each probe packet as in MINC, we model the passage of w consecutive probe packets through the multicast tree. For each node $k \in V$, we estimate the quantity $A_k(0|w)$ that is the probability of observing 0 losses on the path from source S to k given that w probe packets are sent from S. From this, the passage probability of the path from S to k, denoted by p_k is calculated as $p_k = (A_k(0|w))^{1/w}$. Then, the passage probability of the link terminating at k, denoted by α_k is calculated as $\alpha_k = p_k/p_{f(k)}$.

For each node k, $A_k(0|w)$ is estimated as follows. Let $\gamma_k(0|w)$ denote the probability that at least one receiver in the subtree rooted at k observes 0 losses given that w probes are sent from S. For each node k, $\gamma_k(0|w)$ can be calculated directly from the feedback data. Then,

$$A_S(0|w) = 1$$

$$A_k(0|w) = \gamma_k(0|w), \forall k \in R$$

For all $k \in V - R - S$, $A_k(0|w)$ is estimated by finding the root of the following polynomial

$$\gamma_k(0|w) = A_k(0|w)\left\{1 - \prod_{d \in d(k)} (1 - \frac{\gamma_d(0|w)}{A_k(0|w)})\right\}$$

When window size $w = 1$, EMLE reduces to MINC and the passage probability of the path from S to k is estimated as $p_k = A_k(0|1)$.

EMLE does not require any explicit implementation. The quantities defined above are analogous to those used in MINC loss estimator, but have extended meanings. Any implementation of MINC loss estimator can be used for estimating passage probabilities, if receivers report binary feedbacks in the following manner

$$\text{feedback} = \begin{cases} \mathbf{1} & \text{if } 0 \text{ losses among } w \text{ probes,} \\ \mathbf{0} & \text{if } 1 \text{ or more losses among } w \text{ probes.} \end{cases}$$

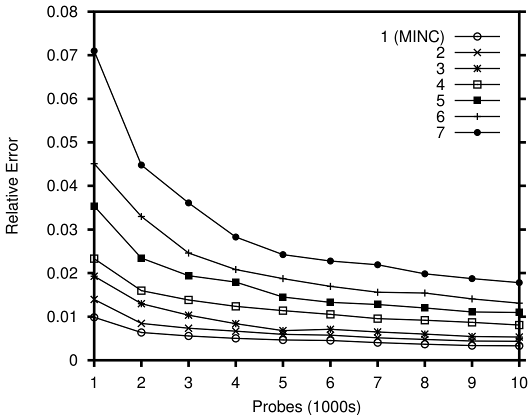

(a) Model-based Simulation

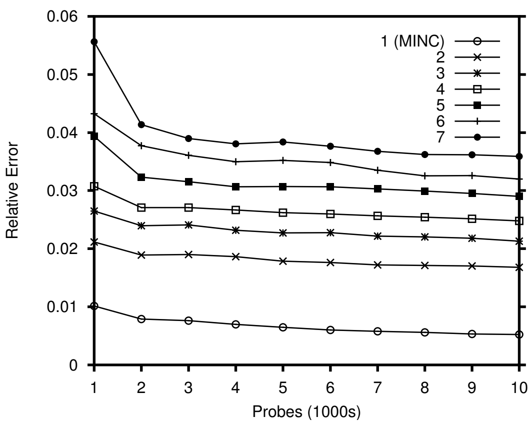

(b) NS Simulation

Fig. 1. Relative Error for windows of different sizes

Giving these "new" feedbacks to an implementation of MINC loss estimator will result in the estimation of quantity $(\alpha_k)^w$ for each link terminating at k. Quantities α_k are then obtained in a straight forward manner.

III. Experiments

Fig 1 shows the behavior of EMLE for Model-based and NS simulations. For these experiments, we simulated an 8-receiver complete binary tree. In Model-based simulations, losses on links are created using Bernoulli losses. In NS simulations, losses on links occur due to buffer overflows on nodes as the probe packet competes with background TCP and exponential on-off UDP traffic. For NS simulations, the parameters were set as in [3]. For both cases, passage rates of links varied from 85% to 95% and the simulations were run 100 times. Fig 1 plots the average absolute relative error for one of the links in the tree for window sizes 1 to 7. Window size 1 corresponds to MINC.

The accuracy of EMLE depends on the estimation of $A(0|w)$ for each node. As the window size w increases, $A(0|w)$ estimated is less accurate since the number of feedbacks used for its estimation decrease (feedbacks corresponding to all probes in a window having been received). In Model-based simulations, the Bernoulli loss assumption holds perfectly. Thus, estimating the passage probability as $(A(0|w))^{1/w}$ yields low errors as compared to NS simulations, where the Bernoulli loss assumption holds only approximately. With a window of size 7, EMLE spends $1/7$ bits of feedback per probe. In MINC, 1 bit of feedback is spent per probe.

IV. Conclusions and Future Work

In this abstract, we introduced the Extended MINC Loss Estimator which can be used to perform loss tomography with less feedback bits per probe. We showed its behavior for Model-based and NS simulations. When used with RTCP, it can help in the reduction of feedback bandwidth. When MINC is used with RTCP, thinning is used, i.e., receivers report feedbacks corresponding to selective probes. EMLE like MINC can be used both on the original or thinned probes.

At present, EMLE estimates only the first element of *loss distribution* i.e., $A(0|w)$. If receivers report the number of losses in each window (i.e., values $0..w$), we would like to know if all elements of loss distribution can be estimated, i.e. elements $A(i|w)$, $0 \leq i \leq w-1$. If all elements of loss distribution are available, the passage probability can be estimated in the following manner.

$$p_k = \sum_{j=0}^{j=w-1} \frac{(w-j)A_k(j|w)}{w}$$

If receivers report the number of losses observed in each window, then $\lceil \log_2(w+1) \rceil$ bits are needed per feedback.

EMLE has two limitations: (a) when loss rates are high, large window sizes cannot be used since the loss distribution element $A(0|w)$ cannot be estimated accurately, (b) at present, EMLE does not handle the loss of feedbacks. If the feedback loss process is MAR (missing at random), we would like to know if ideas of MINC loss estimator with missing data [6] apply to EMLE as well. In future, we shall work on these problems and on ways of estimating loss rates of links with less feedback bits.

References

[1] A. Adams, T. Bu, T. Caceres, N. Duffield, T. Friedman, J.Horowitz, F. Presti, S. Moon, V. Paxson, and D. Towsley, "The use of end-to-end multicast measurements for characterising internal network behavior," *IEEE Communications Magazine*, May 2000.

[2] F.L.Presti, N. Duffield, J. Horowitz, and D. Towsley, "Multicast-based inference of network-internal delay distributions," *IEEE/ACM Trans. Netw.*, vol. 10, no. 6, pp. 761–775, 2002.

[3] R. Caceres, N. 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, 1999.

[4] R. Caceres, N. Duffield, and T. Friedman, "Impromptu measurement infrastructures using RTP," in *IEEE INFOCOM*, June 2002, pp. 1490–1499.

[5] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson, "RFC 1889: RTP: A transport protocol for real-time applications," Jan 1996.

[6] N. G. Duffield, J. Horowitz, D. Towsley, W. Wei, and T. Friedman, "Multicast-based loss inference with missing data," *IEEE Journal of Selected Areas of Communications*, vol. 20, no. 4, pp. 700–713, 2002.