Low Feedback MINC Loss Tomography

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

Timur Friedman, Rémy Bellino
Laboratoire Lip6-CNRS
timur.friedman@lip6.fr, zeremz@free.fr

Nick Duffield
AT&T Labs−Research
duffield@research.att.com

***Abstract*— We present an extension to the MINC loss estimator which can be used to perform loss tomography with less feedback bits per probe. In MINC loss tomography, 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 k consecutive probes. Receivers have the option of reporting different types of feedbacks which require between 1 and $\lceil \log_2(k+1) \rceil$ bits per k probes. EMLE leverages the analysis of MINC delay distribution estimator 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 tomography tool that infers the internal characteristics of a network underlying a multicast tree based on end-to-end measurements. MINC can infer the loss rates and delay distributions of internal network links [2], [3]. To infer loss rates, the source injects probe packets into the multicast tree and each receiver reports whether it received the probe packet (`1`) or not (`0`). Based on the binary feedbacks collected from 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 end-to-end measurements are complex to deploy, [4] proposed an architecture which couples MINC measurements with RTP/RTCP [5]. In this method, RTP data packets of a multicast session act as probes and RTCP is used to report feedbacks. 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 MINC loss tomography with less feedback bits. We have designed an extension to MINC loss estimator which uses information from available probes but reports less feedback bits.

II. Extended MINC Loss Estimator (EMLE)

In EMLE, the sender injects N probe packets into the multicast tree as in MINC. Instead of providing a feedback of 0 or 1 for every probe, receivers report the number of losses observed in windows of k consecutive probes, i.e., values from $0..k$. The window size k is constant and common to all receivers. In total, receivers report N/k feedbacks which require $(N/k)\lceil \log_2(k+1) \rceil$ bits. Using these feedbacks, the passage probability of each link in the multicast tree is estimated in two steps. First, the *loss distribution* of the link is estimated. Then, its passage probability is calculated from its loss distribution. When receivers use windows of size k, the loss distribution of a link l is the set of elements

$$P_l(i|k), \quad i = 0..k-1$$

$P_l(i|k)$ is the probability of i losses on link l given that k packets were sent on link l.

The method of estimating the loss distribution of a link is similar to the method of estimating its delay distribution [3]. The loss distribution elements are estimated inductively starting from the first element $P(0|k)$; $P(i|k)$ is calculated using previously estimated elements $P(0|k)..P(i-1|k)$ and the feedback data. The number of loss distribution elements estimated vary with the number of bits used per feedback. To estimate $P(0|k)..P(i-1|k)$ for all links, receivers need to report only feedback values from $0..i$ where the largest value i is used to report i or more losses. Thus, the estimation of $P(0|k)..P(i-1|k)$ requires only $\lceil \log_2(i+1) \rceil$ bits of feedback per k probes. For instance, estimation of $P(0|k)$ alone requires 1 bit feedback per k probes. Depending on the amount of bits spent per feedback, some or all elements of loss distribution are estimated.

Using the loss distribution of a link, its passage probability is calculated as follows. If the partial list of loss distribution elements $P_l(0|k)..P_l(i-1|k)$ of a link l is estimated, its passage probability p_l is calculated by finding the root of the following polynomial

$$\sum_{j=0}^{j=i-1} \binom{k}{j} p_l^{k-j} (1-p_l)^j = \sum_{j=0}^{j=i-1} P_l(j|k) \qquad (1)$$

If all elements of loss distribution are estimated, then the passage probability of a link l is calculated as

$$p_l = \sum_{j=0}^{j=k-1} \frac{(k-j)P_l(j|k)}{k} \qquad (2)$$

The advantage of using Eq. (2) is that it does not assume the independence of probes within a window, as opposed to Eq. (1) which does.

When window size $k=1$, EMLE reduces to MINC. In this case, Eq. (1) and Eq. (2) are equivalent and as in MINC, the passage probability of link l is estimated as $p_l = P_l(0|1)$.

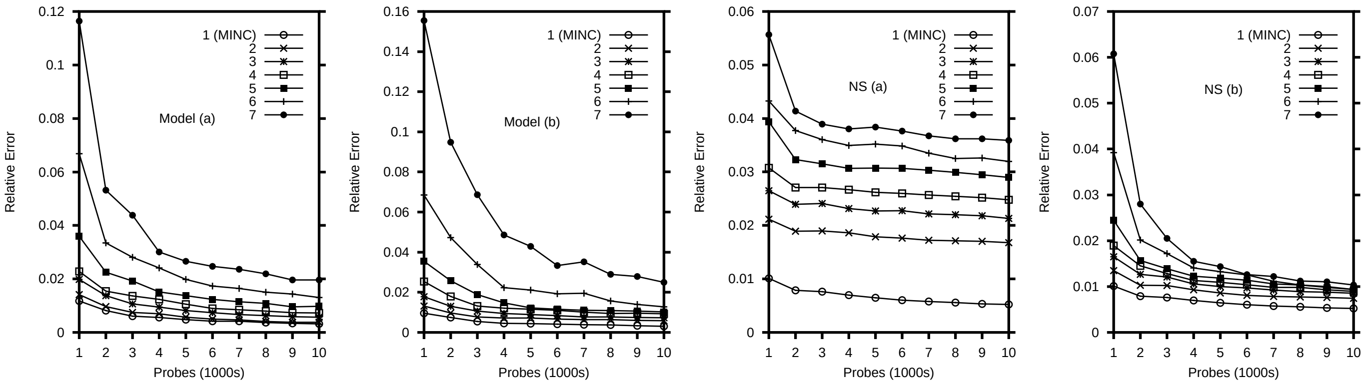

Fig. 1. Errors: (a) First element of loss distribution used (b) All elements of loss distribution used, for Model-based and NS simulations

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 [2]. 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. Fig 1 shows two cases (a) when only the first element of loss distribution is used to estimate the passage probability of each link, i.e., by using Eq. (1) as $(P(0|k))^{1/k}$ (b) when all elements of loss distribution are used to estimate the passage probability, i.e., by using Eq. (2).

Since EMLE calculates loss distribution elements inductively, the accuracy of passage probability estimated is dependent on the estimation of $P(0|k)$. Errors in $P(0|k)$ influence estimates of other loss distribution elements. As the window size k increases, $P(0|k)$ 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, since Bernoulli loss assumption holds perfectly, estimating the passage probability as $(P(0|k))^{1/k}$ is sufficient and yields low errors (Model(a)). When passage probability is estimated using Eq. (2), errors of loss distribution elements add and the passage probability has slightly larger errors (Model(b)). In NS simulations, since the Bernoulli loss assumption holds approximately, estimating the passage probability as $(P(0|k))^{1/k}$ yields high errors (NS(a)). On the other hand, since Eq. (2) does not assume the independence of probes within a window, it yields comparatively lower errors (NS(b)).

Figure 2 shows the amount of bits spent to report feedbacks for different number of probes. Window size $k = 1$ corresponds to MINC. For $k = 7$, when passage probability is estimated using the first element of loss distribution, 1 bit feedback is spent per 7 probes. When it is estimated using all elements, 3 bits of feedback are spent per 7 probes.

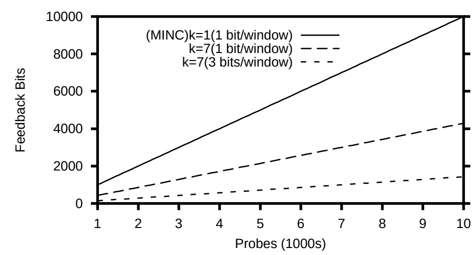

Fig. 2. Number of bits used to report feedbacks

IV. Conclusions

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 to reduce the feedback bandwidth. When MINC is used with RTCP, thinning is used, i.e., receivers report feedbacks to selective probes. EMLE like MINC can be used both with the original or thinned probes. EMLE has two limitations (a) When loss rates are high, large window sizes cannot be used since the loss distribution element $P(0|k)$ cannot be estimated accurately. (b) At present, EMLE does not handle the loss of feedbacks. In future, we shall work on these limitations 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] 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.

[3] 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.

[4] R. Caceres, N. Duffield, and T. Friedman, “Impromptu measurement infrastructures using RTP,” in *IEEE INFOCOM*, June 2002.

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