POSTER SUMMARY: Source Based Multicast Congestion Control [1]

Puneet Thapiyal, Sidhartha, Shivkumar Kalyanaraman
Department of ECSE, Rensselaer Polytechnic Institute
Troy, NY 12180-3590
Contact Email: {*thaps, sidha*}@cs.rpi.edu, *shivkuma*@ecse.rpi.edu

1 Introduction

Congestion control is a key requirement for reliable multicast transport (RMT) protocols. We have recently developed a purely source-based scheme (GSC) which requires no receiver/network support[3]. This paper extends that work to handle high degrees of multiplexing and reasonably large receiver sets.

2 Congestion Scheme Approach

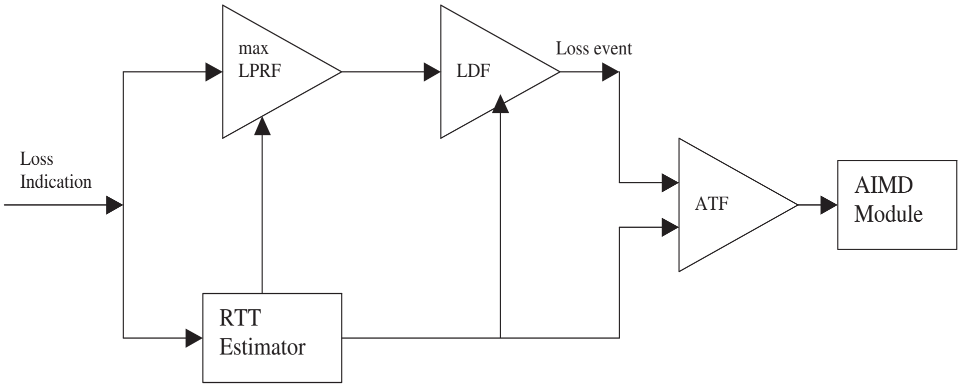

Figure 1: Illustration of the Filter Building Blocks

The two key problems a multicast congestion control (MCC) scheme needs to handle are:

Drop-to-Zero: The problem of reacting to more LIs (loss indications) than necessary. This occurs due to multiple independent congestion/loss paths and high degrees of multiplexing.

TCP-unfriendliness: The problem of reacting to much less LIs than necessary than an equivalent hypothetical TCP connection competing on the worst loss path.

The canonical sender AIMD-based approach involves designing a *loss indication filter (LIF)* which takes as input all the LIs (ΣX_i, where X_i is the number of LIs from receiver i) and passes only $\max X_i$ LIs. We propose a new filter *"max-LPRF" or max-Linear Proportional Response Filter*, an extension of Bhattacharya et al's [1] LPRF approach. Observe that this is a *probabilistic loss filter*, which uses loss statistics as the basis for filtering and may result in

[1] The final poster will be a set of 6-9 slides

multiple rate reductions (RRs) per RTT and in multiple LIs passed from a single receiver per RTT.

The canonical unicast TCP-friendly model responds to at most one loss indication per RTT (also called "loss events" (LEs) [2]). To emulate this behavior, we design two filters, which together converts LIs to LEs on a *per-receiver* basis, and ensures that at most one rate reduction (RR) is effected in an RTT, *across all receivers*. We call the first filter as a *"Loss Delay Filter" (LDF)* and the second as an *"Adaptive Time Filter" (ATF)*. Observe that these are *time filters*, which use time as the basis for filtering. These three filters are cascaded in series before the AIMD module as shown in Figure 1. There is an additional RTT estimation module (described in [3]) which feeds RTT estimates to these modules and the AIMD module.

3 Max-LPRF

Max-LPRF filters out loss indications (LIs) using the following uniform probability (Figure 2):

$$P(Filtering\ LI) = 1 - \frac{\max X_i}{\Sigma X_i} \tag{1}$$

where, X_i is the number of loss indications received from receiver i. $\max X_i$ is the LI count of the receiver sending maximum LIs. ΣX_i is the total number of LIs received from all receivers.

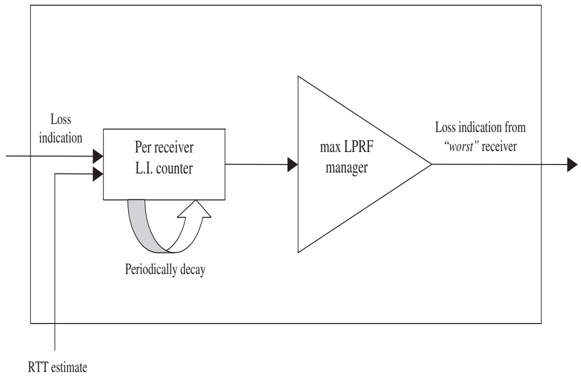

Figure 2: max-Loss Proportional Rate Filter

4 Loss Delay Filter (LDF)

The first step in LDF is to maintain a *per-receiver* moving average of inter-LI arrival times, where the LIs refer to the filtered LIs coming out of the max-LPRF filter. A LI is passed only if its inter-LI arrival time is within an order of the RTT. This filters out "stray" LIs (Figure 3).

5 Adaptive Time Filter (ATF)

The ATF simply filters out LIs (and suppresses rate increases) for a period of RTT after the most recent rate-reduction (RTT), thus enforcing the condition of at most one RR per RTT

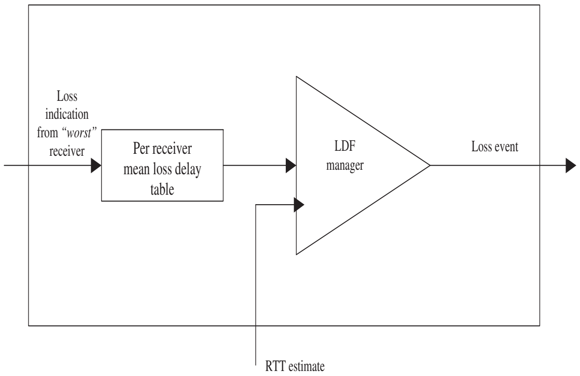

Figure 3: Loss Delay Filter (LDF)

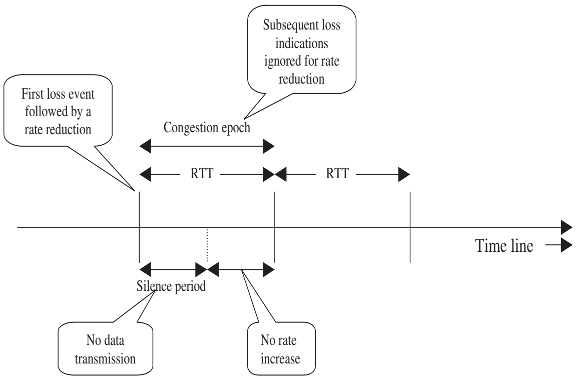

Figure 4: Adaptive Time Filter (ATF)

(Figure 4) [3]. All the components (including the AIMD policy module) are clocked by a carefully designed RTT estimator, and are relatively robust to upto an order of magnitude RTT estimation errors.

6 Simulation Results

The sample simulation results in Figure 6 (based upon the configuration in Figure 5) indicate drop-to-zero resistance. Figure 6(a) shows the scheme with just the ATF (without max-LPRF and LDF) suffering from drop-to-zero when one RM flow participates in four multiplexed bottleneck paths. Figure 6(b) shows the RM flow with the new scheme displaying no drop-to-zero effects.

7 Summary and Future Work

The focus of this work is pure source-based MCC implementation, while being compatible with a simple AIMD model for MCC. Current work involves analytical justification, experimental activity and study of issues including aggregation.

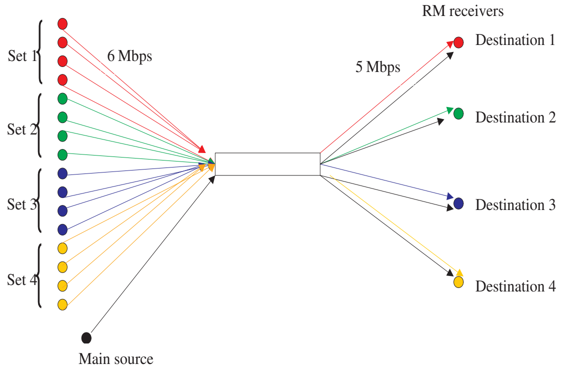

Figure 5: Configuration to test Drop-to-Zero Resistance

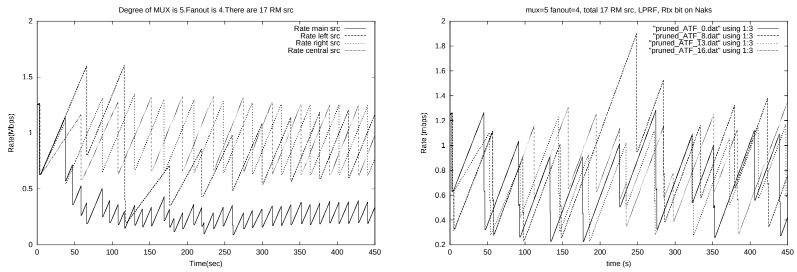

(a) Rates (old scheme)

(b) Rates (new scheme)

Figure 6: Fixing Drop-to-Zero Problems with the Revised Filter Design

References

[1] S. Bhattacharya et al, "A Novel Loss Indication Filtering Approach for Multicast Congestion Control," to appear in *Journal of Computer Communications*, 2000. Available from: http://www.sprintlabs.com/People/supratik/

[2] S. Floyd, M. Handley, J. Padhye, and J. Widmer, "Equation-Based Congestion Control for Unicast Applications," *SIGCOMM'2000*, August 2000.

[3] N. Natu, P. Rajagopal, S. Kalyanaraman, "GSC: A Generic Source-based Congestion Control Algorithm for Reliable Multicast," to appear in *Journal of Computer Communications*, 2000. Available from:http://www.ecse.rpi.edu/Homepages/shivkuma/research/papers-rpi.html#mult