Source Based Multicast Congestion Control

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

Abstract

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. This paper extends that work to handle high degrees of multiplexing and reasonably large receiver sets and demonstrates the TCP friendliness of the scheme.

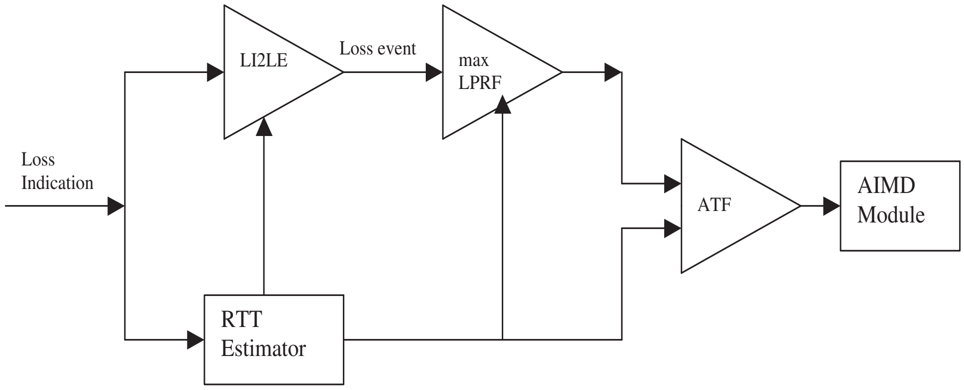

Figure 1: Illustration of the Filter Building Blocks

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.

We design a new *loss indication filter (LIF)* at the sender (source) 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. This filter called the *"max-LPRF" or max-Linear Proportional Response Filter* is an extension of Bhattacharya et al's LPRF approach. This is a *probabilistic loss filter*, which uses loss statistics as the basis for filtering and may result in multiple rate reductions (RRs) per RTT and in multiple LIs passed from a single receiver per RTT. In comparison, the canonical unicast TCP-friendly model responds to at most one loss indication per RTT (also called *"loss events" (LEs)*). To emulate this unicast 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 Indication to Loss Event Filter" (LI2LE)* 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 which feeds RTT estimates to these modules and the AIMD module. The poster will demonstrate latest simulation results with the scheme.