

# Random Early Detection Gateways for Congestion Avoidance

Reviewer: Sayandeep Sen

February 20, 2008

This paper presents a scheme to implement congestion control at the gateway nodes. The superiority of such a scheme (congestion control at gateways) over end-to-end congestion control comes from the fact that, a gateway has a global view of current link utilization (can differentiate between queuing and propagation delays). Hence they can trigger congestion avoidance mechanisms when the delays are congestion induced while letting normal traffic flow to carry on if high latency is the main culprit for end-to-end delays. Also, congestion control at gateways could potentially be part of an architecture to enforce differential treatment to flows based on certain metrics (to maintain fairness between flows).

The method proposed by the authors involves maintaining a max and a min threshold of permissible average queue size at the router. Once the queue grows beyond the high mark, the algorithm starts to drop packets. When the queue length is between the two thresholds the algorithm probabilistically marks packets to inform the sending hosts of impending congestion. The intent being, the sender would throttle its own traffic and hence a congestion would be avoided.

The advantage of the algorithm over droptail strategy are

a) it is able to react earlier than droptail strategy and hence would lead to potentially better throughputs for all the senders (droptail would only kickoff when the queue size is above the high mark),

b) because it drops packets probabilistically it doesn't lead to global synchronization of the backoffs at all the senders. c) also because of being probabilistic it is not biased against bursty traffic.

However there are a few weak points of the work,

a) the algorithm gives the desired performance only when all the senders implement their own congestion control strategy. In case a sender doesn't implement congestion control algorithm it might end up getting unfair advantage over a congestion controlled source. For example, in a scenario when the average queue size is between the two thresholds (max, min), as the controlled source would throttle its sending rate on getting marked packets while the other source would keep on pumping packets. The authors do not do any simulations to model this scenario.

b) The success of the algorithm depends of correct estimation of the max and min thresholds, which in turn depends on the traffic patterns observed over the links. While the traffic patterns at the core of the Internet is stable enough to accurately set the threshold values, the traffic at edge gateways (which are nearest to the senders and hence most effective traffic controllers) is generally bursty in nature. Hence, determination of the two thresholds is a fairly non-trivial problem.