Random Early Detection Gateways for Congestion Avoidance

Sally Floyd and Van Jacobson
Paper review by **Ehsun Daroodi**
810181050

In the absence of explicit feedback from the gateway, transport-layer protocols could infer congestion from the estimated bottleneck service time, from changes in throughput, from changes in end-to-end delay, as well as from packet drops or other methods. The most effective detection of congestion can occur in the gateway itself. Random Early Detection (RED) gateway do congestion avoidance by controlling average queue size and marking packets. Marking packets means notifying the end source of congestion either by explicit messages or by dropping its packets.

The RED algorithm: when the average queue size (*avg*) is less than the minimum threshold (MINth), no packets are marked. When avg is greater than the maximum threshold (MAXth), all arriving packets are marked. When avg is between the MINth and MAXth, each arriving packet is marked with probability Pa, where Pa is a function of average queue size. The RED gateway functions most effectively when |MAXth - MINth| is larger than the typical increase in the calculated average queue size in one RTT.

The initial packet-marking probability, Pb, is calculated as a linear function of the average queue size. Pa can be calculated by either geometric random variables or uniform random variables. The second method is superior as marks packets less clustered and avoids global synchronization.

Simulation results: RED gateway can control the average queue size while accommodating transient congestion and network power is higher with RED gateway than with Drop Tail gateways.

RED gateways' properties: appropriate time scales, no global synchronization, simplicity, maximizing global power and appropriate for a wide range of environments. RED gateways don't attempt to ensure that each connection receives the same fraction of total throughput and don't explicitly control misbehaving users.

Because RED gateways randomly choose packets to be marked during congestion, they could easily identify which connections have received a significant fraction of the recently marked packet. When the number of marked packets is sufficiently large, a connection that has received a large share of marked packets is also likely to be a connection that has received a large share of the bandwidth. This information could be used by higher policy layer to restrict the bandwidth of those connections during congestion.

RED gateway algorithm can be implemented efficiently, with only a small number of add and shift instructions for each packet arrival.