A Comparison of RED's Byte and Packet Modes*

Wesley M. Eddy, Ohio University, weddy@irg.cs.ohiou.edu
Mark Allman, BBN Technologies/NASA GRC, mallman@bbn.com

Longer version: http://irg.cs.ohiou.edu/~weddy/papers/redmodes.ps

The Random Early Detection (RED) active queue management scheme [FJ93], recommended for deployment by the IETF [B98], is used by routers to control their congestion levels by informing end hosts of incipit congestion. Provisions in the original RED research give routers the option of measuring their queue length in terms of the number of packets buffered or in terms of the number of bytes of buffer space consumed by the queued packets. In addition, a packet's size may factor into the computation of its *marking probability* (the probability a given packet will be dropped or the end host notified via ECN to reduce the sending rate). Given a network with heterogenous packet sizes (e.g., the Internet) the choice of using packets or bytes as the units in RED's calculations can have implications on performance and fairness. As an example, when queueing and marking in terms of packets a RED queue may mark small packets at an inappropriately high rate based on the small packet's use of the bandwidth. Through simulations we quantify the performance implications of selecting amongst combinations of queue measurement and marking modes from the standpoints of link utilization and fairness. Additionally we have investigated the impact of setting RED's *mean packet size* parameter that the byte modes depend upon.

Simulations of a link congested with several bulk FTP transfers show that selecting packet modes provides greater link utilization, while byte modes enhance fairness among competing flows. We also observe that the decision of whether or not to adjust the marking probability of a packet based on its size has more effect on performance than the choice of units for queue measurement. These observations hold constant across different mixes of packet sizes and traffic levels, however the results are most apparent when there is a high load and wide range of packet sizes.

Our next set of simulations involves HTTP traffic congesting the RED gateway. We first note that in lightly loaded networks all RED variants performed roughly the same and therefore choosing one particular variant does not seem important. In networks with moderate to heavy congestion levels we note that marking based on a packet's size yields higher fairness. Also, measuring the queue in terms of bytes causes better link utilization in our simulations, while at the same time lowering the average queueing delay and reducing the degree of jitter in queueing delay. Since the majority of real-world Internet traffic today consists of such HTTP transfers, this makes a strong argument for consideration of byte-based RED modes for deployment.

RED's byte modes use a constant called the *mean packet size* (MPS) to calculate a given packet's probability of being marked. Next, we investigate how difficult it may be to set the MPS and the sensitivity of our results to the MPS setting. We examine packet traces collected at two different points in the Internet. Our analysis indicates that the degree of mixture and rate of change in packet sizes varies across both time and network. We conclude from this that the MPS must remain a site-selectable parameter – or possibly something that RED could tune on-the-fly. Our simulations of both FTP and HTTP traffic show that the MPS parameter should be *approximately accurate* for the mix of traffic traversing a link, but does not need to be exact. Specifically, we note that settings within several hundred bytes of the measured mean packet size are roughly equivalent in performance.

Finally, we introduce a rating system for comparing the performance of various RED modes based on link utilization and fairness, and favoring some degree of balance between the two metrics. We show that the combination of a queue measured in bytes and marking based on the number of bytes in a packet generally offers the best ranking among other options tested across our simulation scenarios. However, a concrete suggestion on which RED variant to use remains elusive, as in some scenarios byte modes rank lowest. Therefore, we conclude that some testing of different modes may be necessary on a site-by-site basis in order to achieve desired performance properties.

*The authors thank NASA's Glenn Research Center for funding this work.

References

[B98] Robert Braden et. el. Recommendations on Queue Management and Congestion Avoidance in the Internet, April 1998. RFC 2309.

[FJ93] Sally Floyd and Van Jacobson. Random Early Detection Gateways for Congestion Avoidance. *IEEE/ACM Transactions on Networking*, 1(4):397–413, August 1993.