A Comparison of RED’s Byte and Packet Modes

Wesley M. Eddy
Ohio University

Mark Allman
NASA GRC/BBN

Baseline Simulations

- Used homogenous packet sizes to gauge inherent differences independent of packet size
- Throughput
  - For packet sizes < MPS, byte modes generate higher throughput
  - For packet sizes > MPS, packet modes generate higher throughput
  - This difference is due purely to the MPS parameter
- Fairness
  - All modes performed comparable well in terms of fairness
  - Typically saw over 0.99 fairness index

Extreme Mix of Packet Sizes

- Bulk transfer packet sizes of 296, 576, 1500, 4352, and 16k bytes
- Marking mode more important than queue measurement mode
- Observed inverse relationship between utilization and fairness

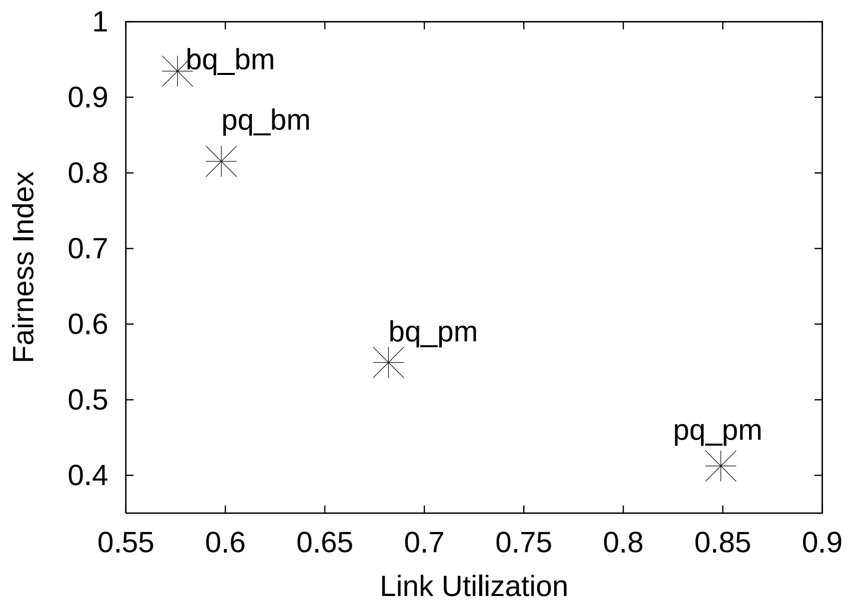

HTTP Traffic with Realistic MTUs

- Queue lengths when RED measures the queue in packets and in bytes:

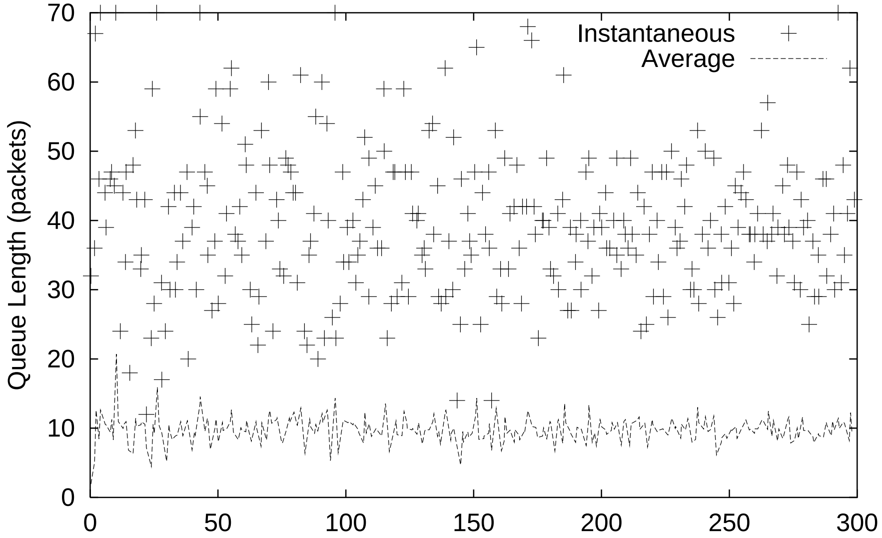

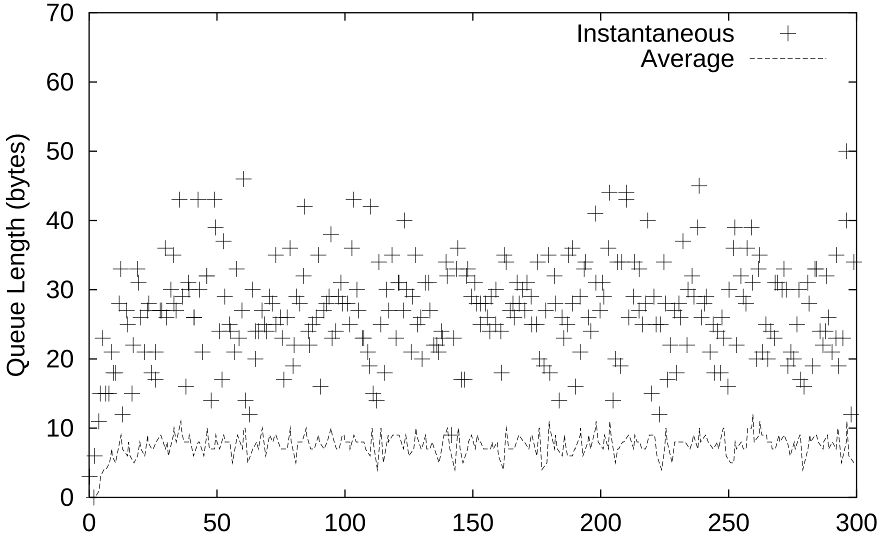

Mean Packet Size Dependence

- Byte-marking robust to poor MPS settings
- MPS setting can be used to tweak performance of byte modes
- Utilization and fairness as functions of MPS in WWW simulations:

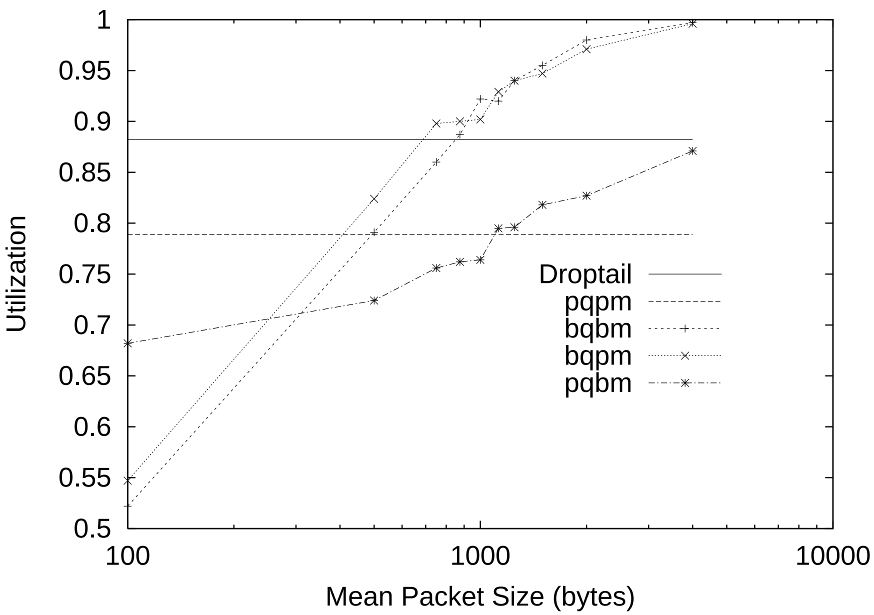

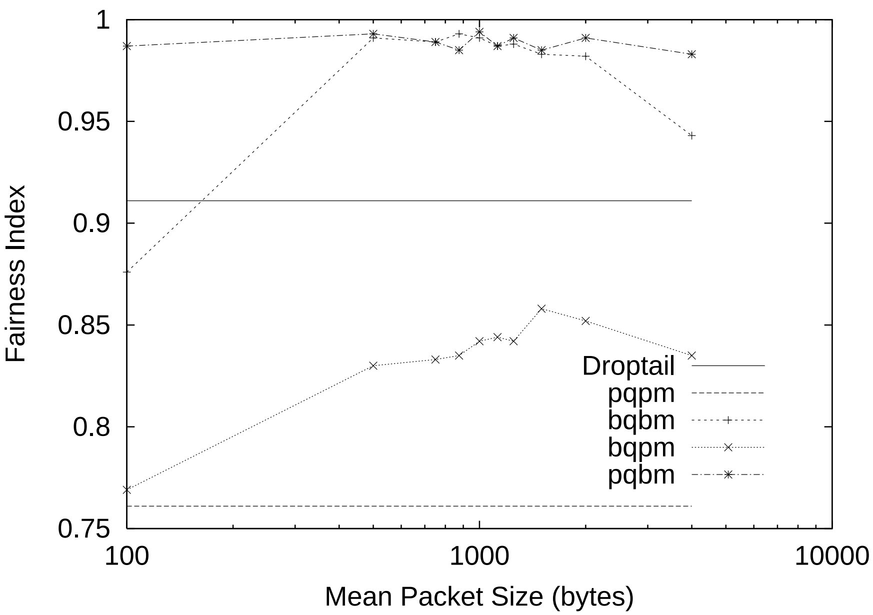

Real World Mean Packet Size

- Distribution of difference between mean packet size in a given sample and one previous to it:

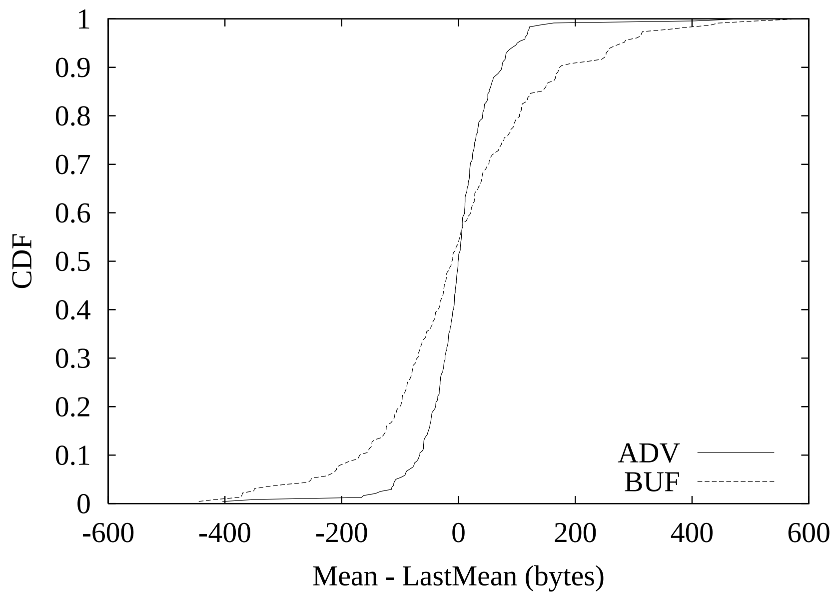

- Mean packet size can be dynamic over time and network
  - No universal setting will be good for everyone
  - Must remail site-configurable

Comparing Performance

- We use a formula to rate modes based upon their performance in terms of link utilization and fairness, favoring some degree of balance between the two.

$$R \leftarrow U^{\alpha} \cdot F^{(1-\alpha)}$$

- alpha is used to select preference for higher utilization or fairness

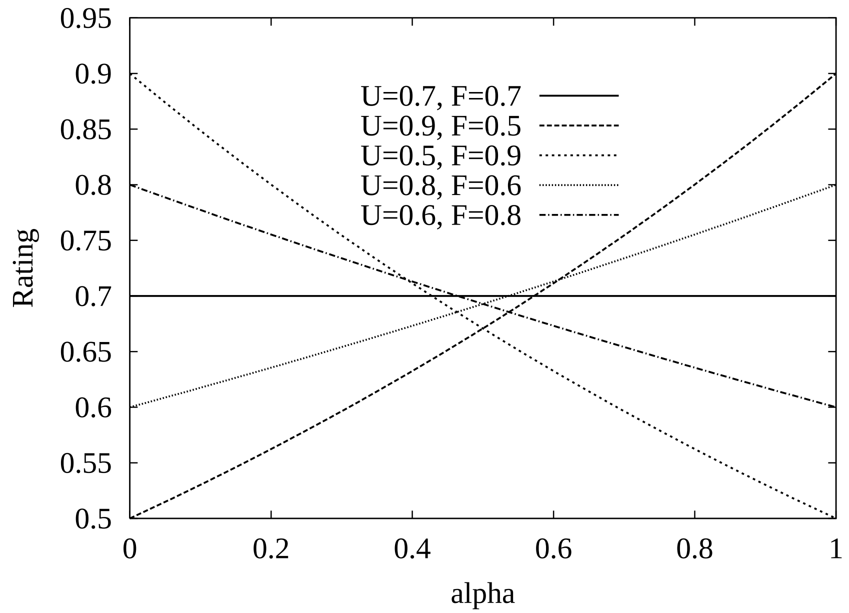

Performance Comparison

Experiment	bq_bm	bq_pm	pq_bm	pq_pm
Extreme Mix	1	3	2	4
M_s	1	4	2	3
M_e	4	3	2	1
Light Load	1	4	3	2
Moderate Load	2	1	4	3
Heavy Load	1	3	2	4
MPS-Bulk	1	4	2	3
MPS-WWW	1	2	3	4

Performance Comparison

- If we value all performance metrics equally, byte marking variants consistently rank highest across several different simulation types
- The byte based queue measurement option combined with the byte based marking mode is the top performer in many cases
- This is a fairly strong argument for using at least the byte marking mode in the real world
- Any mode you pick is still generally better than droptail

Summary

- There are certainly large differences in performance between modes
- Byte-based marking leads to greater fairness among competing flows
- Link utilization is harder to maximize through mode selection
  - Packet modes are generally best
  - Byte-based queue measurement seems good in many tests with WWW-centric traffic
  - Byte modes with ludicrous MPS settings work well if you don’t care about delay