Lightweight Active Router-Queue Management for Multimedia Networking

Mark Parris
Kevin Jeffay
F. Donelson Smith

Department of Computer Science
University of North Carolina at Chapel Hill

http://www.cs.unc.edu/Research/dirt

Outline

- Problem - Supporting multimedia over the Internet
- Context - Active Queue Management
  - » Drop-tail on Full (FIFO)
  - » Random Early Drop (RED)
  - » Flow Random Early Drop (FRED)
- Towards better support for multimedia flows
  - » Class Based Thresholds (CBT)
- Empirical evaluation

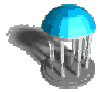

Interactive Multimedia on the Internet

- Interactive multimedia over the Internet
  - » video teleconferencing, distributed VR
- Media requirements and characteristics
  - » low latency
  - » periodic transmission times
  - » tolerant of some loss
  - » elastic but lower bounded bandwidth requirement
- Application specific congestion management
  - » Ranges from none to responsive

Congestion Avoidance

- Drops are the only widely used indicator of congestion
- Drops and retransmissions
  - » Congestion collapse
  - » TCP's congestion avoidance (Jacobson)

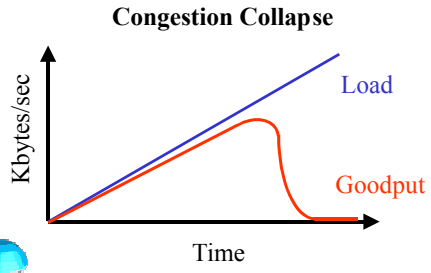

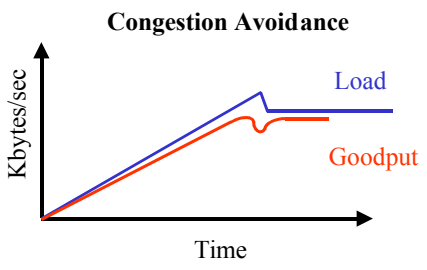

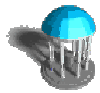

Traditional Queueing in the Router (FIFO)

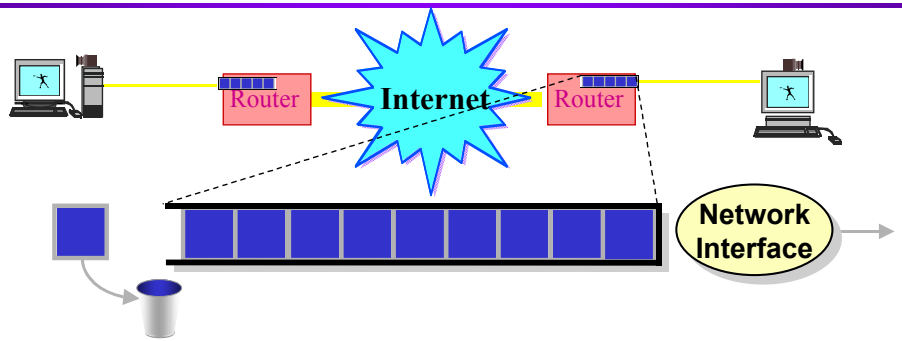

- Full queues
  - no room for bursts of packets
  - Synchronized back-off (TCP/IP)
- High latency
- Lock-out phenomena
- High drop-rate

Active Queue Management Random Early Detection (RED)

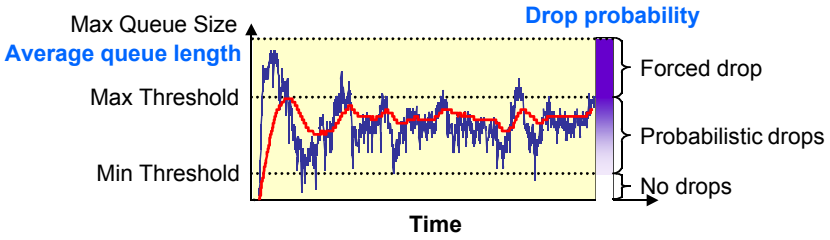

- Weighted average accommodates bursty traffic
- Multiple modes based on threshold values
- Probabilistic and forced drops
  - » avoid consecutive drops
  - » drops proportional to bandwidth utilization
    - – (drop *rate* equal for all flows)

FIFO and RED Vulnerable to Misbehaving Flows

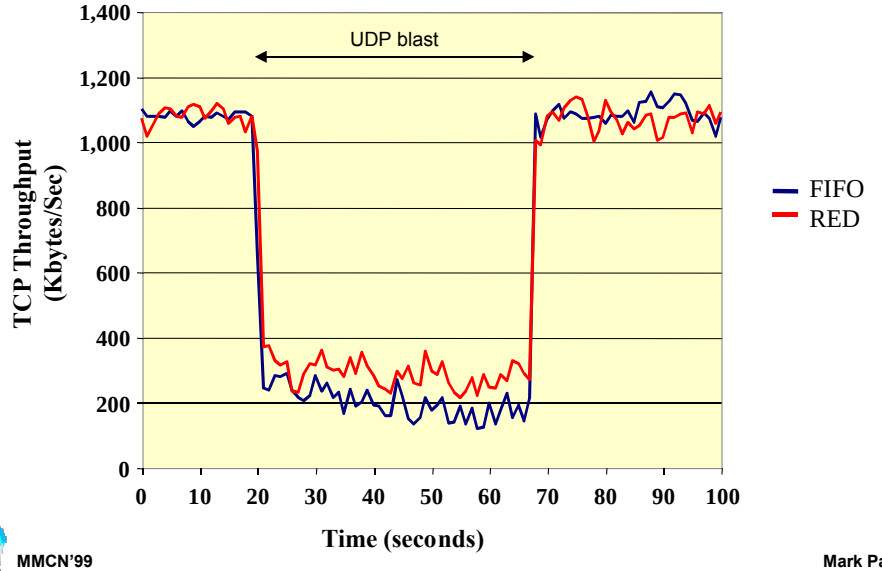

Queue Management and Congestion Avoidance

- Braden, et al (1998) recommend:
  - » Deploy Active Queue Management (e.g. RED)
    - – Avoid full queues, reduce latency, reduce packet dropping, avoid lock-out phenomena
  - » Continue research into mechanisms to deal with unresponsive or aggressive flows. (Floyd & Fall)
    - – Unresponsive traffic's impact on RED, FIFO
- Multimedia is "unresponsive" for a reason.
  - » can tolerate some loss
  - » price for unnecessary reliability is too-high

Flow Random Early Detect (FRED) Goals

- **Fairness**
  - » Protect all TCP flows from effects of aggressive flows
  - » Protect fragile TCP flows
  - » Establish drop *rate* proportional to load

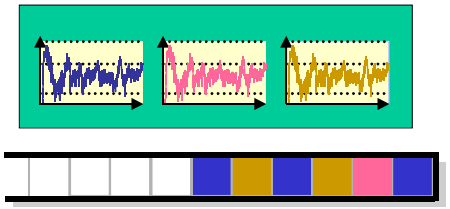

Flow Random Early Detect (FRED)

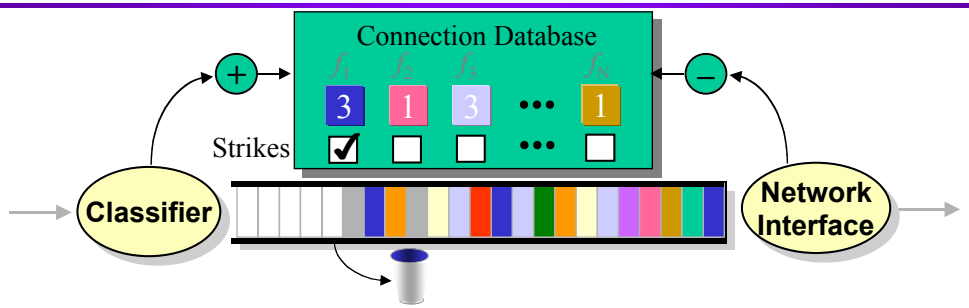

- **Maintain a single FIFO queue but track the number of packets in the queue from each connection**
- **Drop packets from a connection when the connection exceeds its share of the queue's capacity**
  - » Drops are proportional to bandwidth used
  - » Unresponsive flows are identified and penalized

TCP Performance with FRED

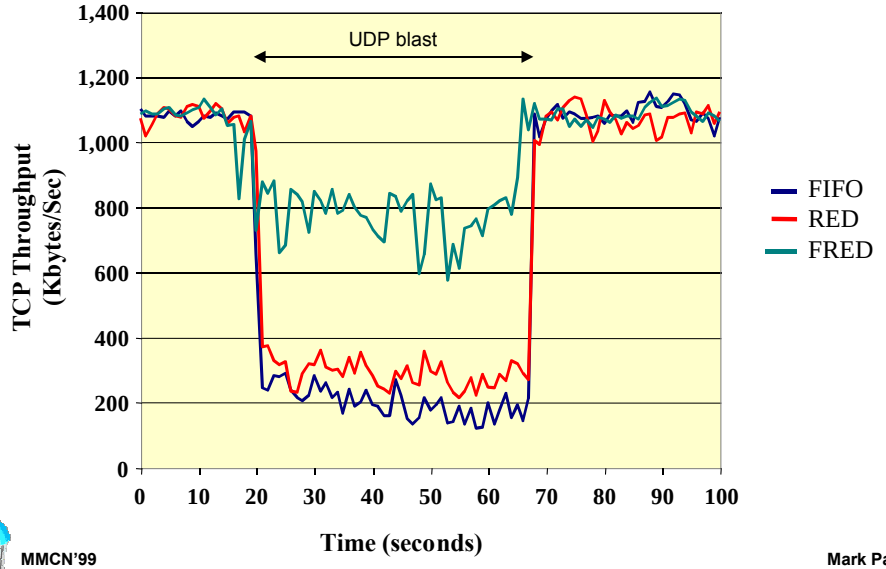

Class Based Thresholds (CBT) Goals

- **Isolation**
  - » Responsive from unresponsive
  - » Unresponsive: Multimedia from aggressive
- **Flexible Fairness**
  - » Something more than equal shares
- **Light weight**
  - » minimal state
- **Maintain benefits of RED for responsive traffic**
  - » feedback
  - » distribution of drops

Class Based Thresholds (CBT) Design

- **Isolation:**
  - » Packets are classified into one of several classes
  - » Separate statistics kept for each class
  - » Marked (multimedia) and unmarked non-TCP traffic are monitored with aggregate statistics
- **Flexible Fairness:**
  - » Configurable thresholds determine ratios between classes during periods of congestion
- **Light weight:**
  - » State per class instead of per flow, one queue
- **Maintain benefits of RED for responsive traffic:**
  - » Continue to apply RED policies to TCP

Class Based Thresholds (CBT)

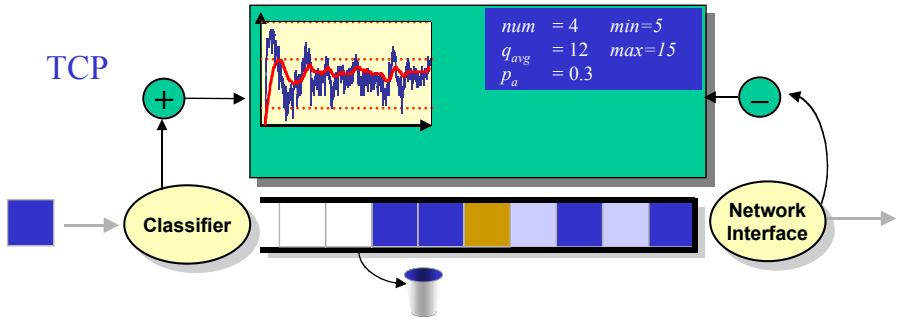

- **One queue**
- **Packet classification**
- **Specific drop policies for each class**
  - » RED for TCP

Class Based Thresholds (CBT)

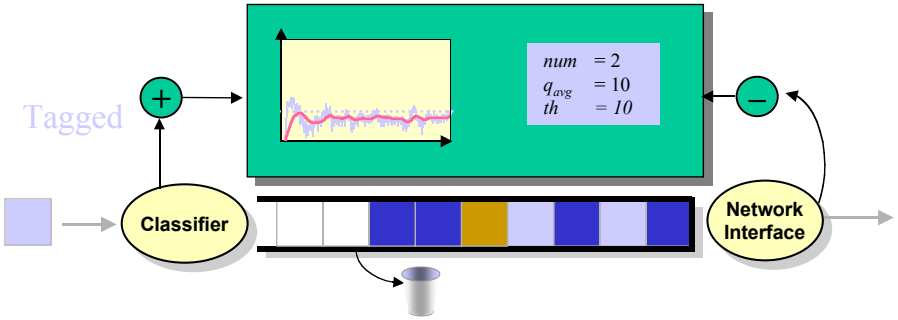

- **Tagged traffic is allocated a share of the queue using a threshold on average queue occupancy**

Class Based Thresholds (CBT)

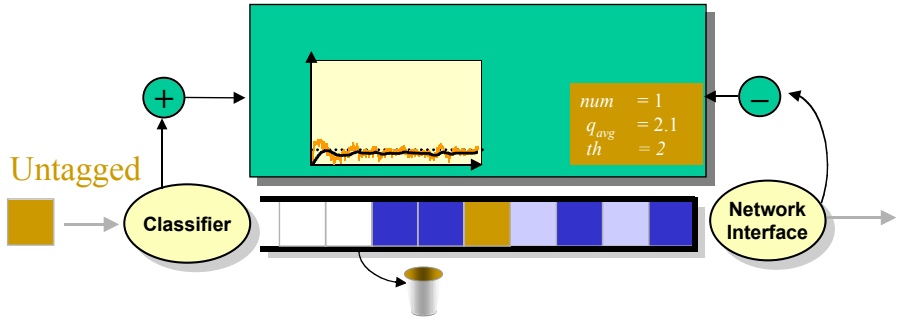

- **Other traffic is limited to a share of the queue using a (small) threshold on average queue occupancy**

Class Based Thresholds (CBT)

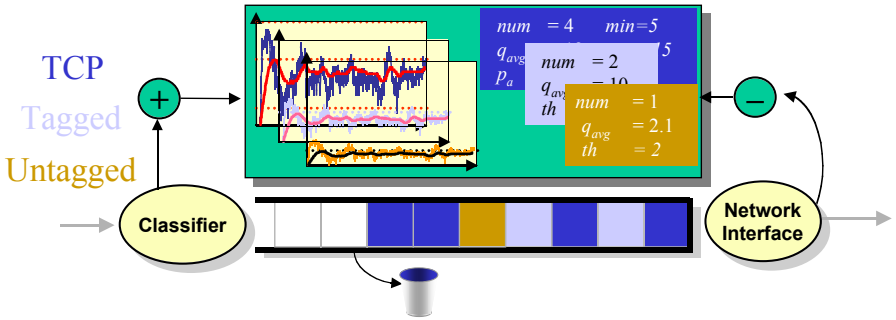

- Maintain a single FIFO queue but track the average number of packets in the queue from each class
- Drop packets from a class based on policy for that class:
  - » During times of congestion, each class is limited to a fixed share of network bandwidth in proportion to the threshold.

Empirical Evaluation
Experimental Setup

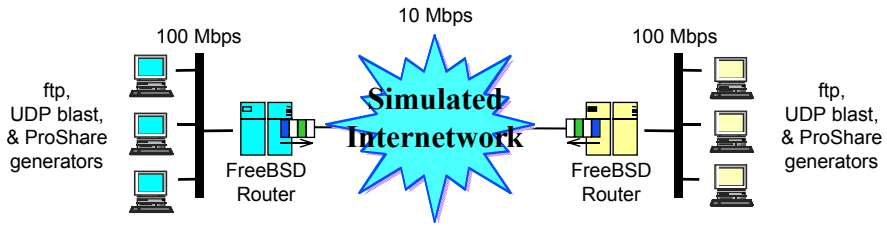

- Senders transmit a mix of TCP, real-time, and non-real-time UDP traffic

Experimental Results
TCP Throughput

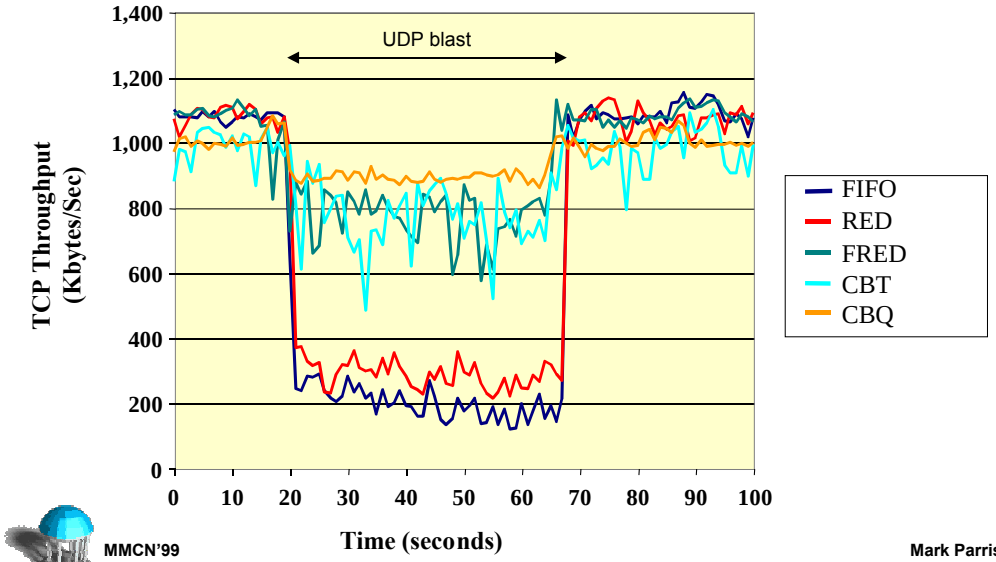

Experimental Results - TCP
CBT vs. Packet Scheduling

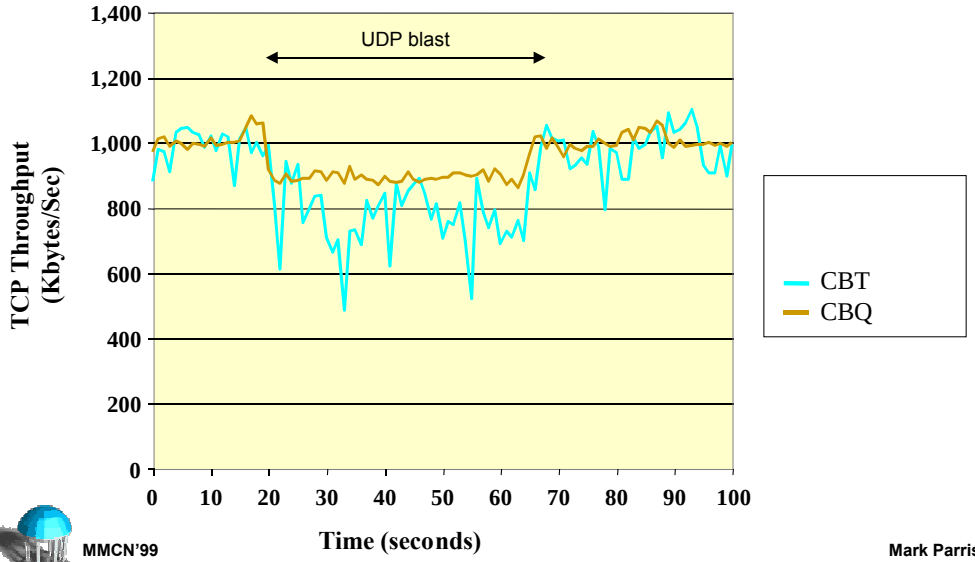

Multimedia Latency

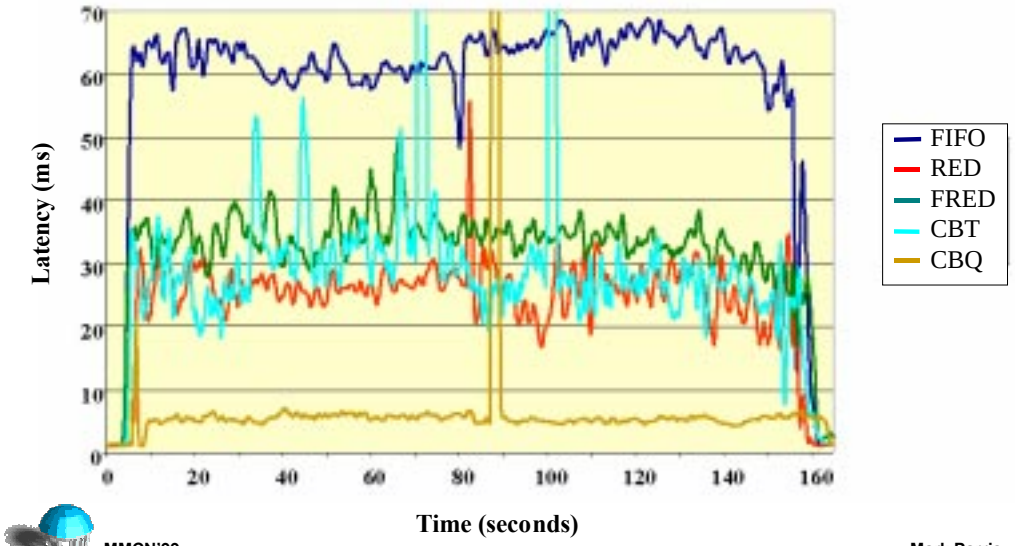

Drop Rate for Continuous Media

Queue Management Scheme	Drop Rate for Continuous Media
FIFO	32.4%
RED	30.0%
FRED	35.7%
CBT	1.3%
Packet Scheduling	0.0%

- Proper resource allocation
- Only Packet Scheduling has a lower drop rate for continuous media

Results

- Active Queue Management comparison
  - » Better TCP performance
  - » Less state than FRED (no per flow state)
  - » Lower latency to multimedia applications
  - » Lower drop rate for multimedia
- Packet scheduling comparison
  - » comparable performance

Summary

- Current queue management schemes are vulnerable to aggressive flows.
- Calls for mechanisms to identify and penalize unresponsive flows
- Class-Based Thresholds provides isolation for TCP and multimedia while providing a useful level of service for multimedia.
- Allocation of the buffer offers comparable performance and is simpler than packet scheduling.