University of North Carolina at Chapel Hill

Tuning RED for Web Traffic

Mikkel Christiansen
Kevin Jeffay
David Ott
Don Smith

SIGCOMM, August 2000

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

Tuning RED for Web Traffic
Research context

- RFC 2309 strongly advocates deployment of *random early detection* (RED) active queue management in routers

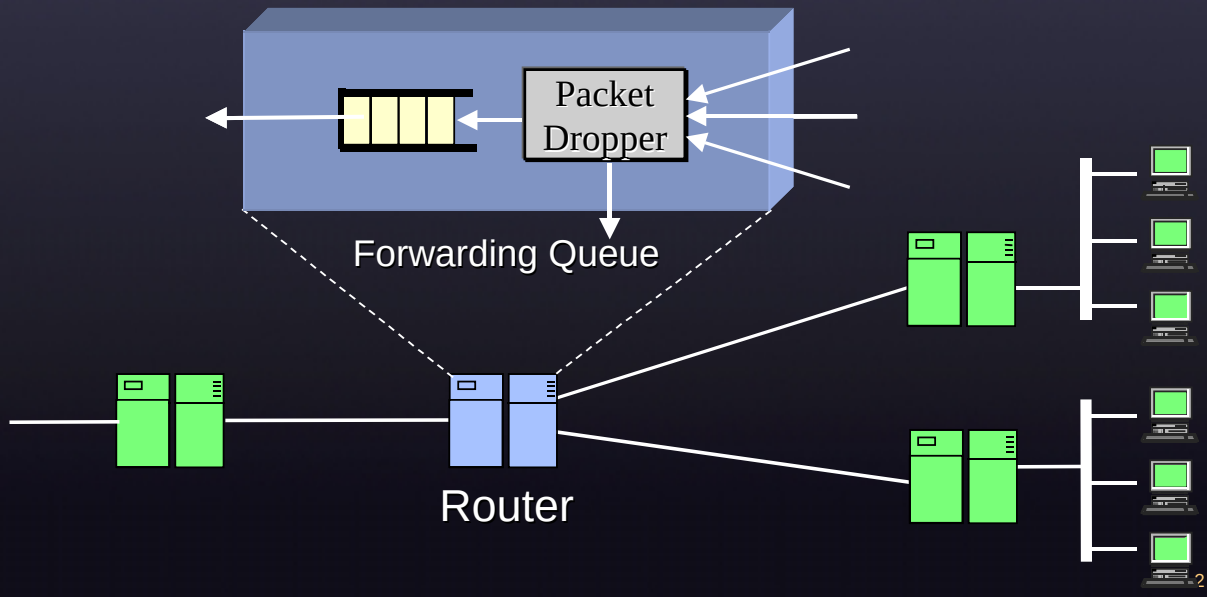

Tuning RED for Web Traffic
Research context

- RFC 2309 strongly advocates deployment of *random early detection* (RED) active queue management in routers

  "All available empirical evidence shows that the deployment of active queue management mechanisms in the Internet would have substantial performance benefits. There are seemingly no disadvantages to using the RED algorithm, and numerous advantages. Consequently, we believe that RED active queue management algorithm should be widely deployed."

- Measurement studies have shown that 60-80% of traffic in the Internet is HTTP
- How is HTTP response time performance affected by RED and can RED be tuned to optimize it?

Tuning RED for Web Traffic
Overview

- We've conducted an empirical evaluation of the effect of RED on response times of HTTP request/response transactions
- We conclude:
  - RED provides no advantage over drop-tail FIFO for offered loads up to 90% of link capacity
  - Above 90% RED can be tuned to provide better performance, however,
    - » We couldn't tune performance for short and long requests simultaneously
    - » The "best" RED parameter settings are a function of load
    - » We were unable to develop any systematic method for tuning RED
- Our results suggest that ISP's should focus on provisioning capacity rather than depending on RED

Tuning RED for Web Traffic
Outline

- RED active queue management
- Experimental methodology
  - HTTP traffic model
  - The network setup
- Experimental results
- Conclusions and future work

RED Active Queue Management
Algorithm description

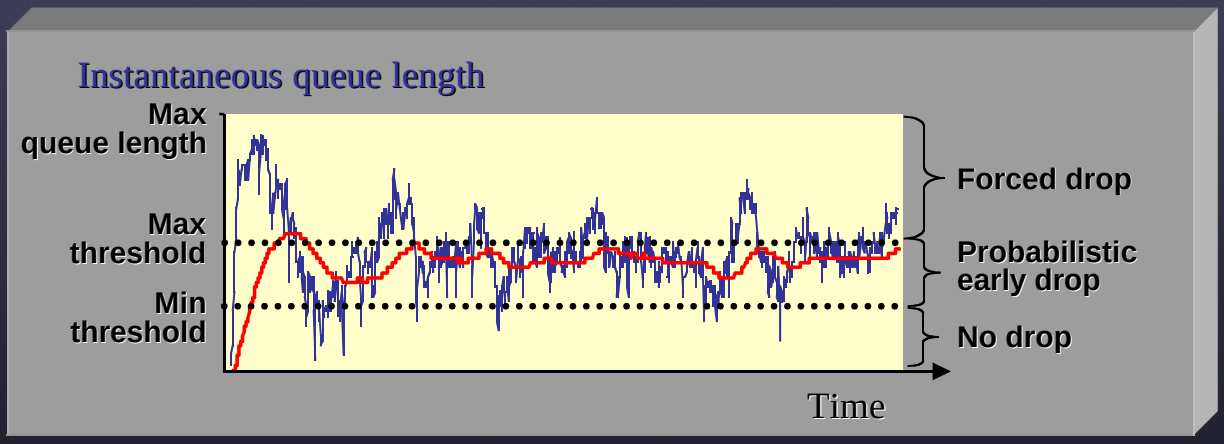

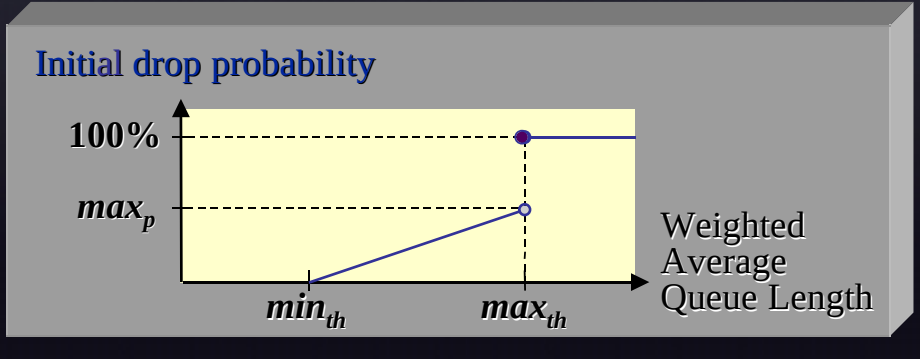

Experimental Methodology
HTTP traffic generation

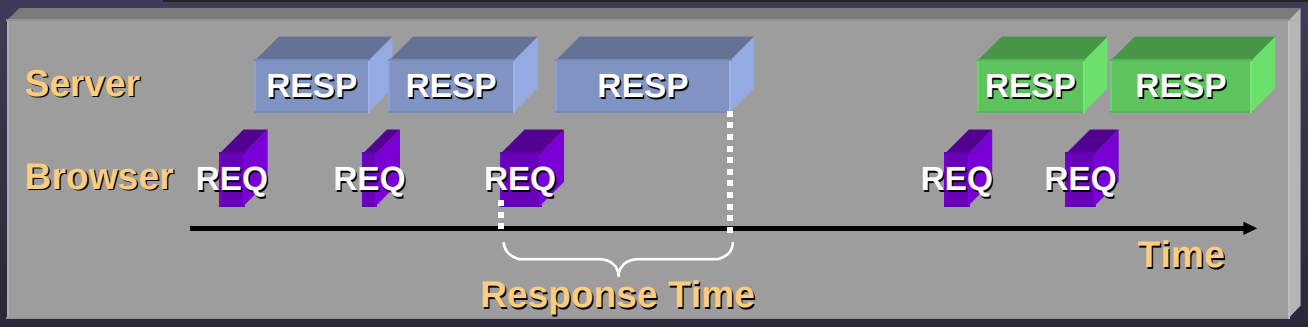

- We generate HTTP 1.0 traffic using the Mah document model and his empirical distributions of parameters
- Primary random variables:
  - Request sizes
  - Reply sizes
  - User inter-document-request think time
  - Number of objects per page
  - Consecutive documents per server
- Response time is the time it takes a "browser" to receive an object from a "server"

Experimental Methodology
Network Setup

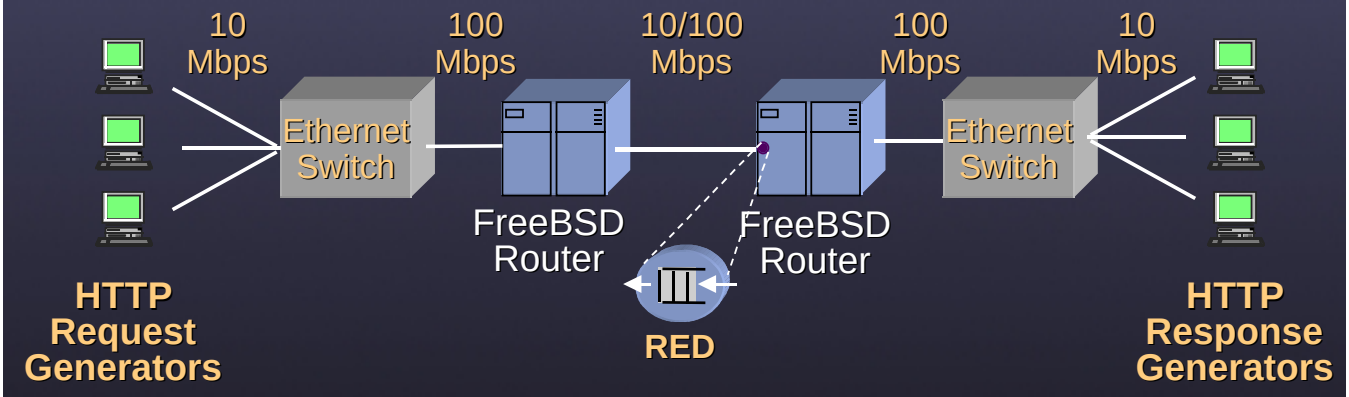

- Model a campus/enterprise network with a single wide area link to an upstream ISP
  - Browsers are local to the campus; servers are beyond the ISP
- RED is run on the path from servers to browsers

Experimental Methodology
Network Setup

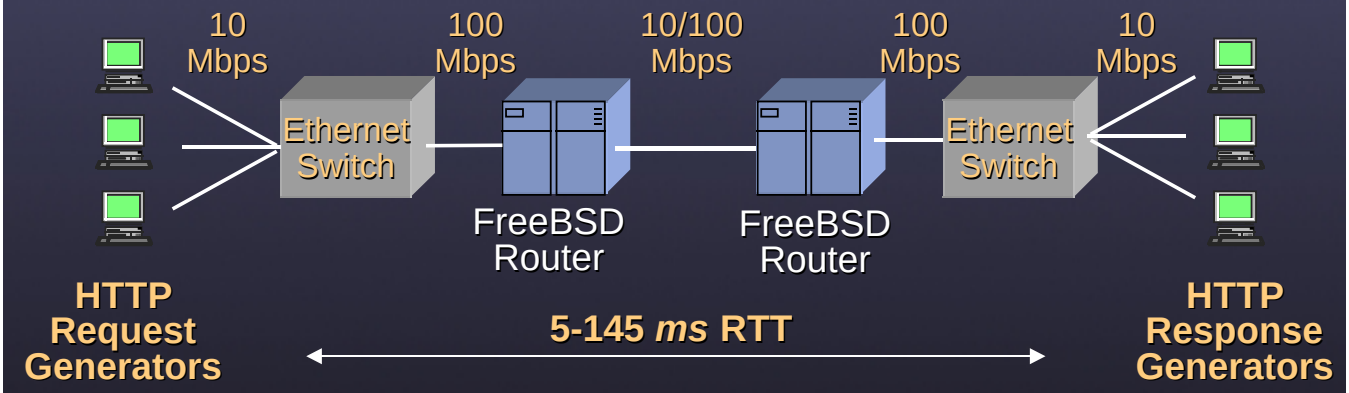

- Primary simulation parameters
  - Round-trip-time between each request/response generator pair
  - Offered load on an unconstrained 100 Mbps link as a percentage of the capacity of the 10 Mbps link

    E.g., 3,500 simulated users generate ≈ 9.8 Mbps on a 100 Mbps link, hence 3,500 users == 98% load on a 10 Mbps link

Experimental Results
Experimental plan

- For drop-tail FIFO we determine the queue length(s) that produce the "best" response time distribution
  - Determine queue length as a function of offered load
- For RED we determine the best parameter settings for a set of offered loads
- Use the performance (of drop-tail FIFO) on the unconstrained (100 Mbps) network as the baseline for evaluations

Experimental Results
Effect of offered load under drop-tail FIFO

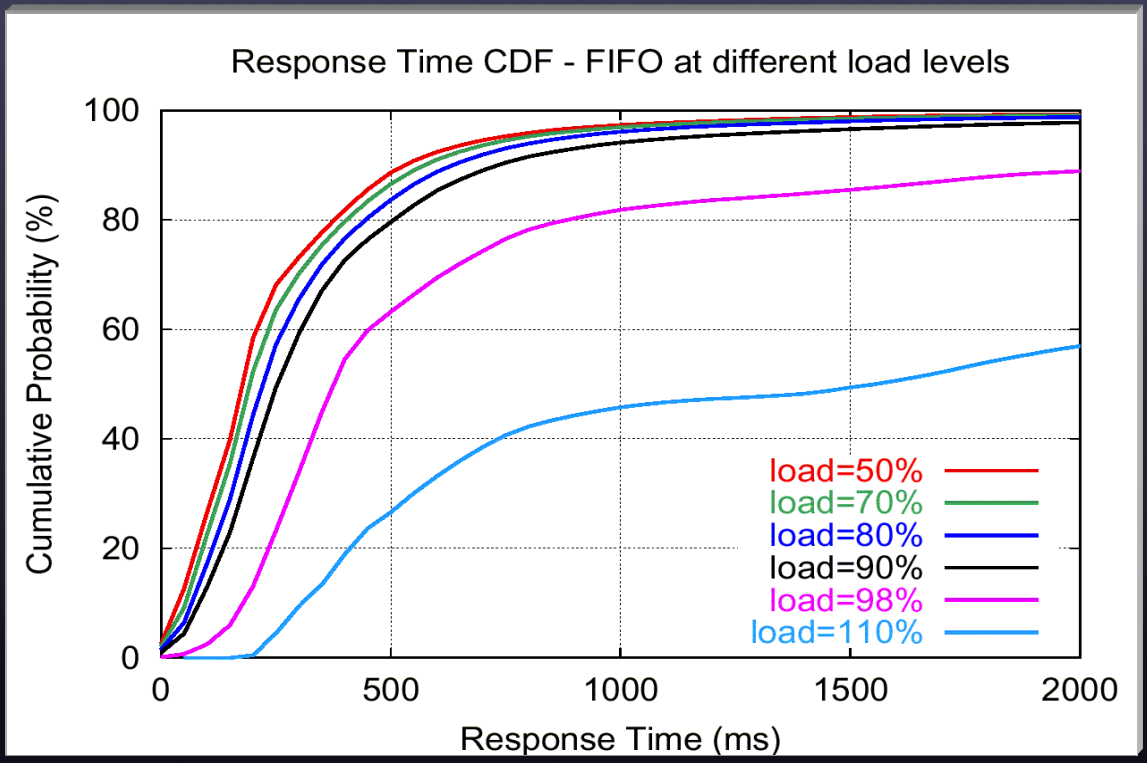

Experimental Results
Effect of offered load under drop-tail FIFO

- No significant difference in performance observed below 90% load
- Above 90% of link capacity, response time degrades quickly
- Therefore, we focus on optimizing performance at offered loads of 90 and 98% of link capacity

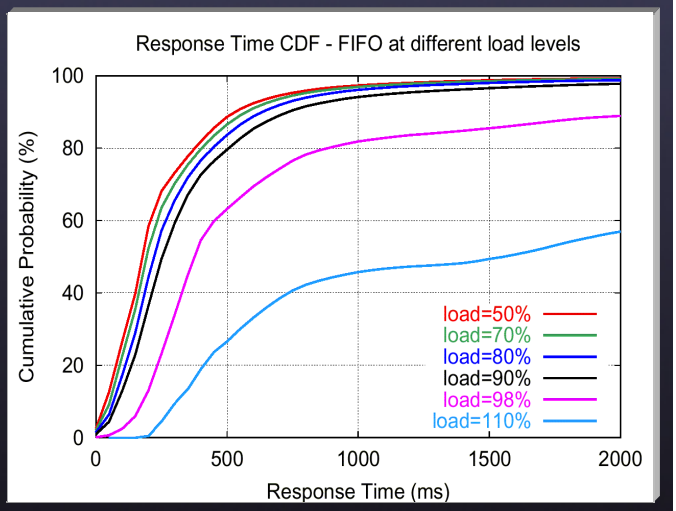

Experimental Results
Drop-Tail FIFO Parameter Determination

- Run experiments with queue lengths from 30-240 packets
- Results: Trade-off exists between optimizing for shorter versus longer responses
- We declare *qlen* = 120 to be the "winner"

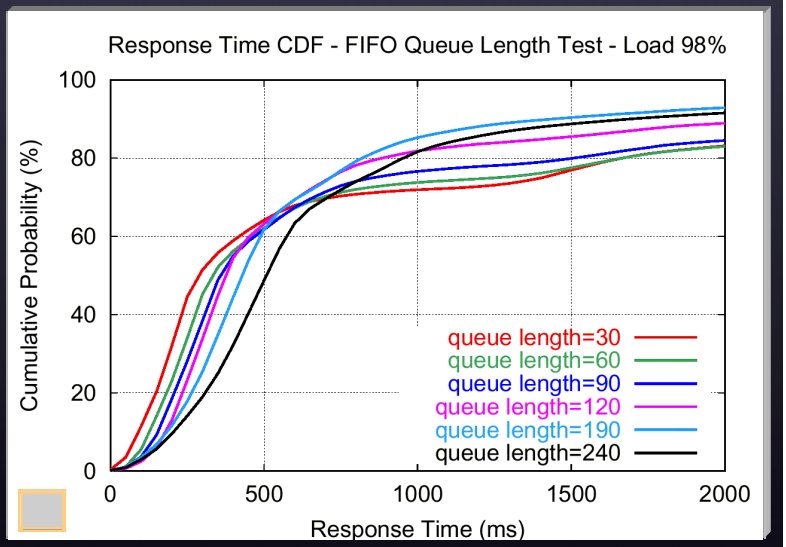

Experimental Results
RED parameter determination

- Ignore the effects of queue length
  - Set *qlen* to "infinity"
- Vary thresholds by the (recommended) function
  - $max_{th} = 3 \times min_{th}$
- Same trade-off exists between optimizing for shorter versus longer responses
- Best performance results from thresholds in the range (30, 90) to (60, 180)

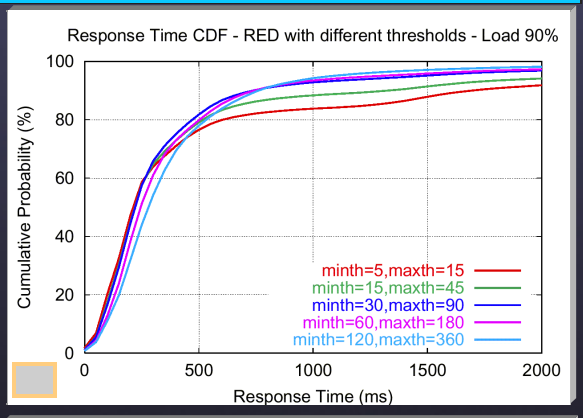

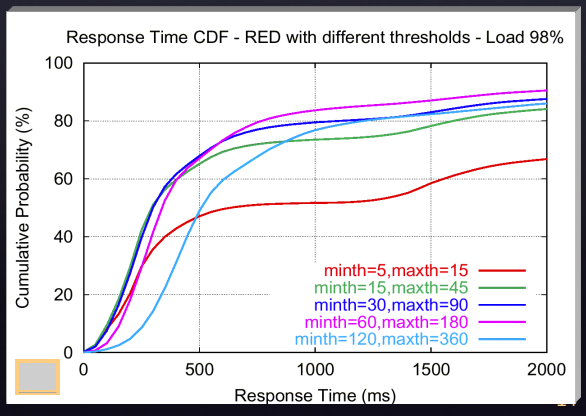

RED Parameter Determination

Effect of weight (w_q) and maximum drop probability (max_p)

- Combine testing of w_q and max_p
  - The two were determined to be closely related
- max_p = 1/4 is too aggressive
- Recommended settings work well
  - w_q = 1/512
  - max_p = 1/10

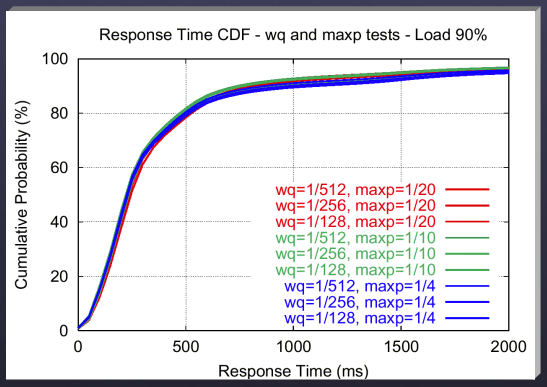

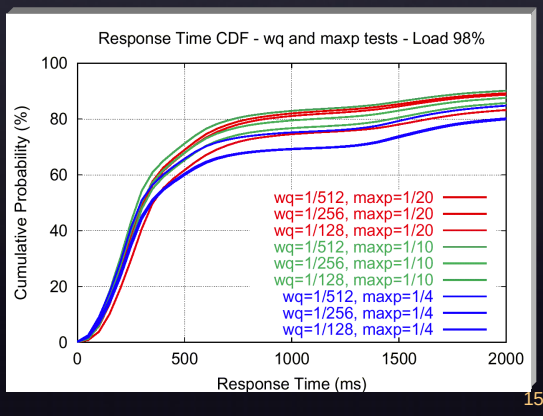

RED Parameter Determination

"Good" RED parameter settings

- Settings for the best response times at 90% and 98% load are significantly different:
  - 90%: $(min_{th}, max_{th}) = (30, 90)$, $w_q = 1/512$, $max_p = 1/10$
  - 98%: $(min_{th}, max_{th}) = (5, 90)$, $w_q = 1/128$, $max_p = 1/20$
- Tuning for better link utilization has a negative effect on response times
- Tuning for lowest drop rate also has a negative effect on response times

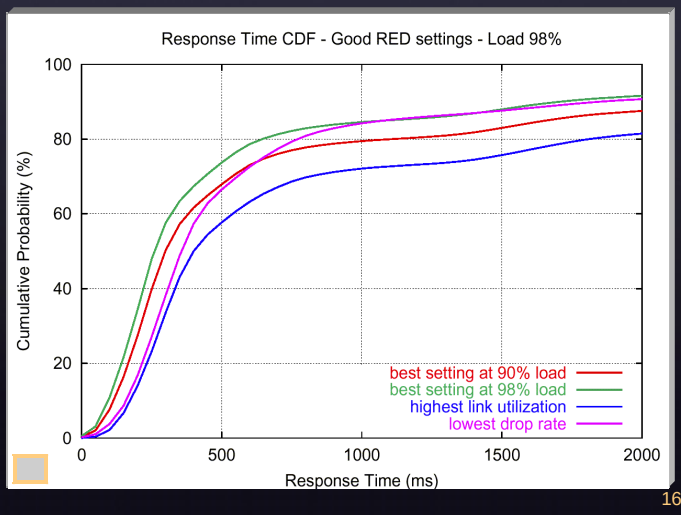

RED Parameter Determination

"Bad" RED parameter settings

- Reasonable RED settings can significantly decrease response time performance
- Example: Default settings for RED in FreeBSD:
  - $(min_{th}, max_{th}) = (5, 15)$
    $w_q = 1/512$
    $max_p = 1/10$
    $q_{len} = 60$

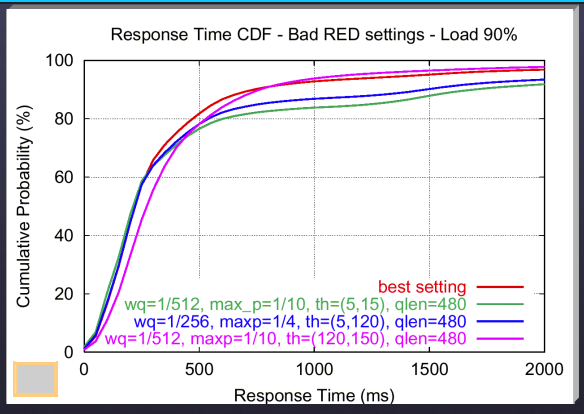

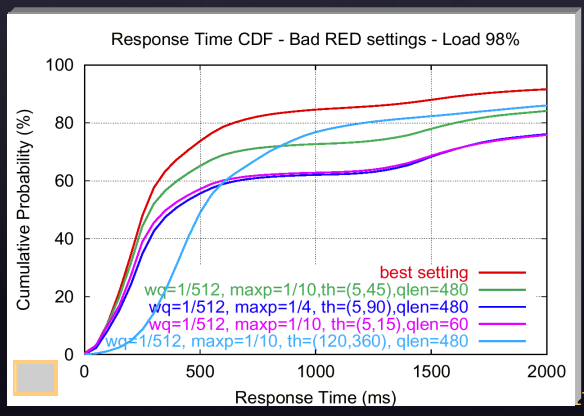

Drop-Tail FIFO *v.* RED

Comparison

- At 90% load and below, FIFO and RED have equal response time performance
- At 98% load RED can outperform FIFO
- (At 110% load RED and FIFO have equal performance)

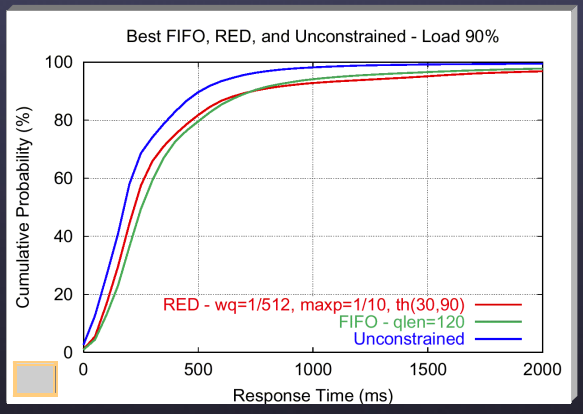

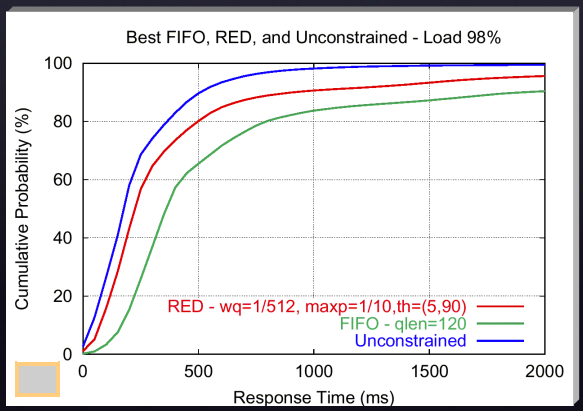

Tuning RED for Web Traffic
Summary and Conclusions

- RED provides no advantage over drop-tail FIFO for offered loads up to 90% of link capacity
  - » The “better performance” claim from RFC 2309 doesn’t hold for HTTP response times
- Above 90% RED can be tuned to provide better performance, however, ...
  - – Doing so is difficult & error prone
    - » The “no harm” claim from RFC 2309, doesn’t hold for HTTP response times
  - – “Better” is subjective
    - » We couldn’t tune performance for both short and long response times simultaneously
  - – Best parameter settings are a function of offered load

Tuning RED for Web Traffic
Summary and Conclusions

- RED provides no advantage over drop-tail FIFO for offered loads up to 90% of link capacity
  - » The “better performance” claim from RFC 2309 doesn’t hold for HTTP response times
- Above 90% RED can be tuned to provide better performance, however, ...
  - – Doing so is difficult & error prone
  - – “Better” is subjective
  - – Best parameter settings are a function of offered load
- In total, the results suggest that if Web performance is an issue, ISP’s should focus on provisioning capacity rather than depending on RED

Tuning RED for Web Traffic
Future Work

- Analysis of observed response time distributions
- Redo experiments with HTTP 1.1 and mixes of 1.0/1.1 traffic
  - But with updated HTTP model
- Redo experiments using a realistic mix of HTTP and other TCP (and UDP) traffic
- Study the impact of using RED with ECN markings
- Redo experiments on a multi-hop network

Experimental Results
FIFO queue length determination

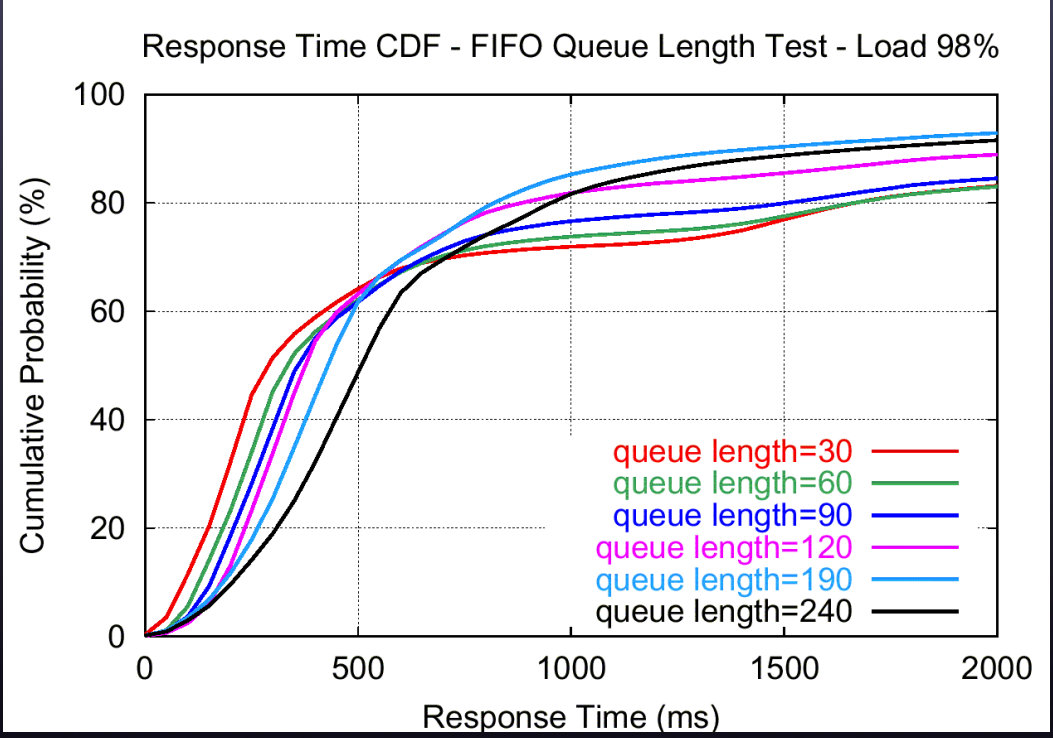

Experimental Results
RED parameter determination

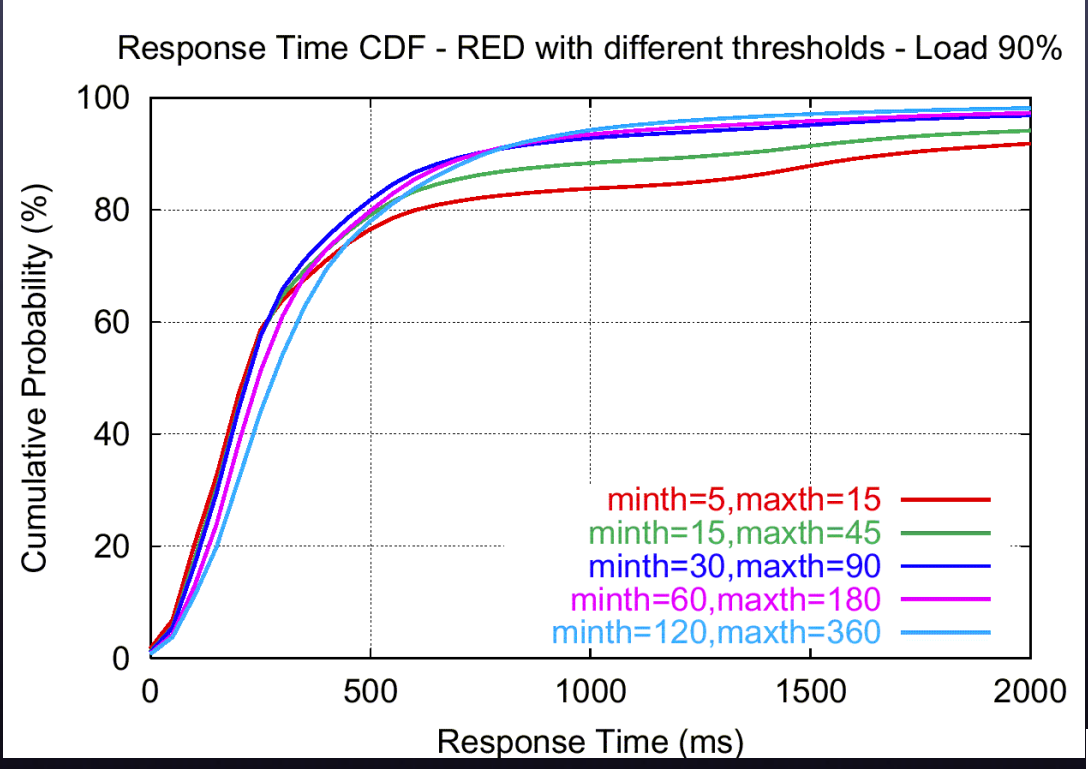

Experimental Results
RED parameter determination

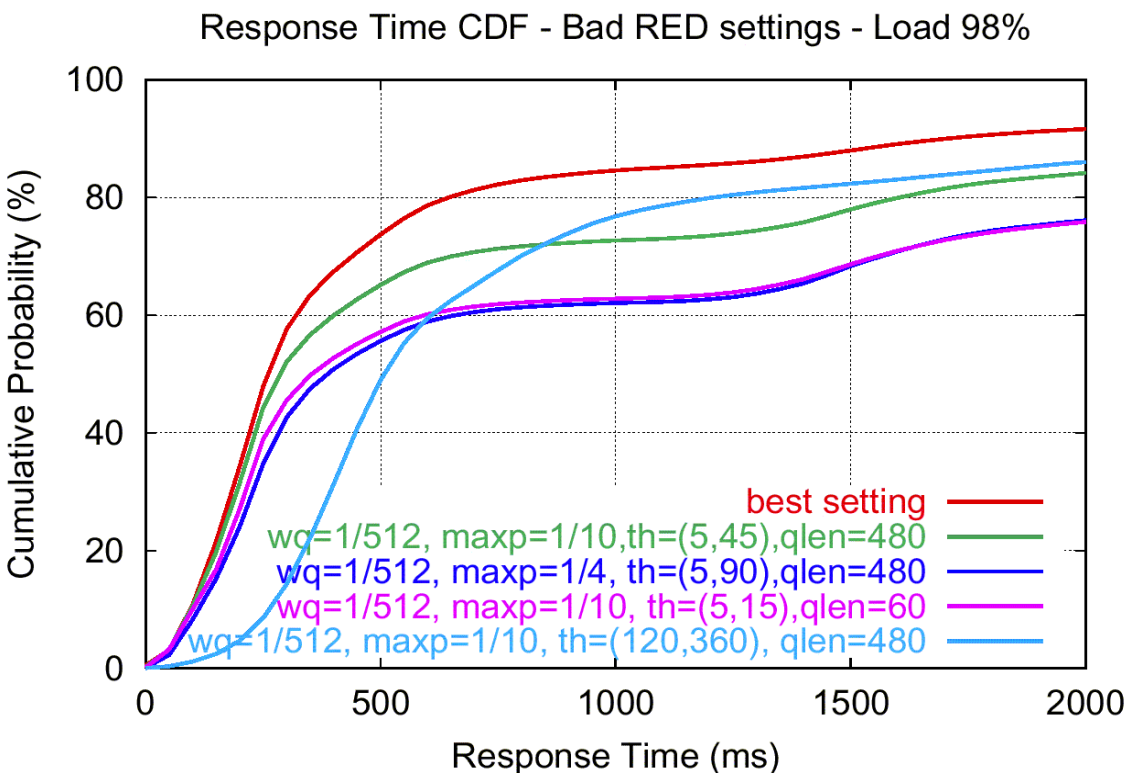

RED Parameter Determination
"Good" RED setting

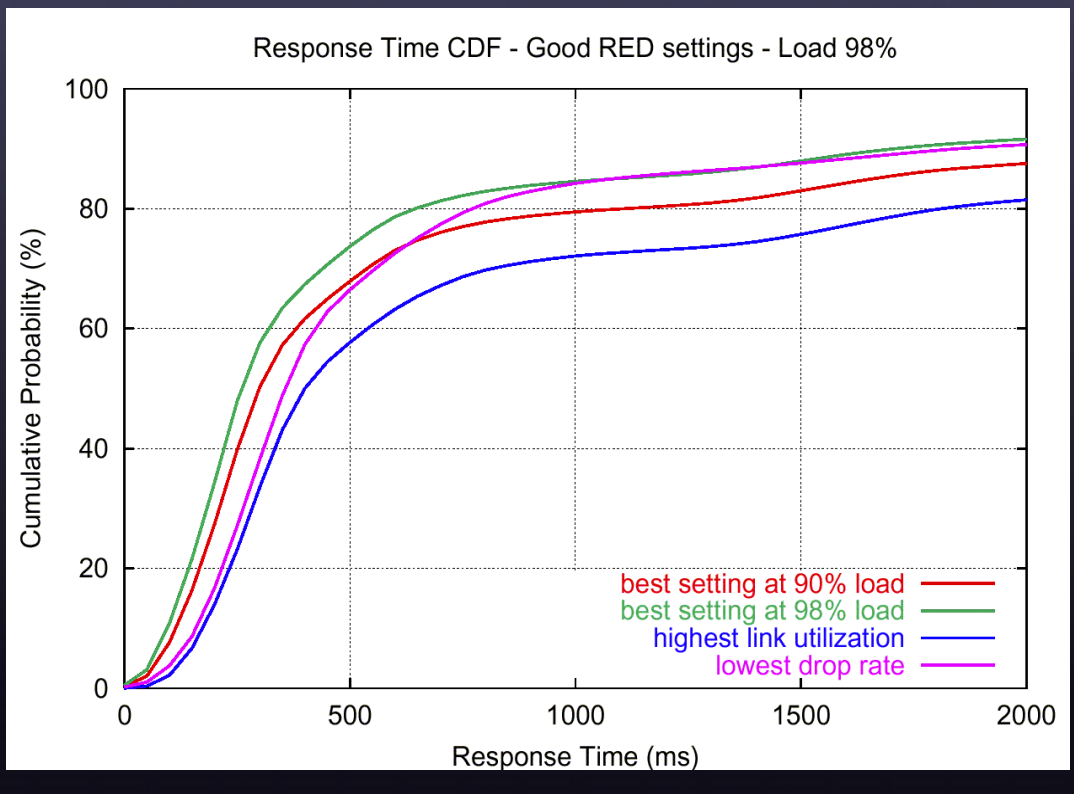

RED Parameter Determination
Bad RED setting

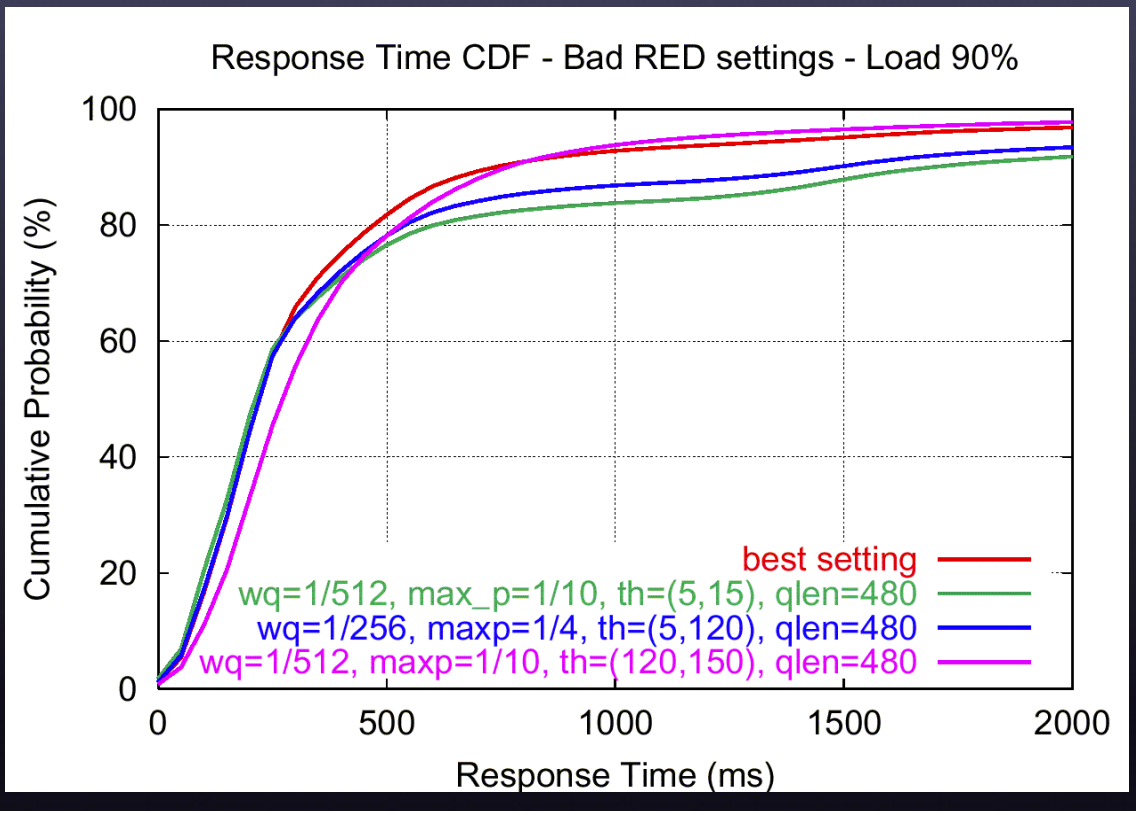

RED Parameter Determination
Bad RED setting

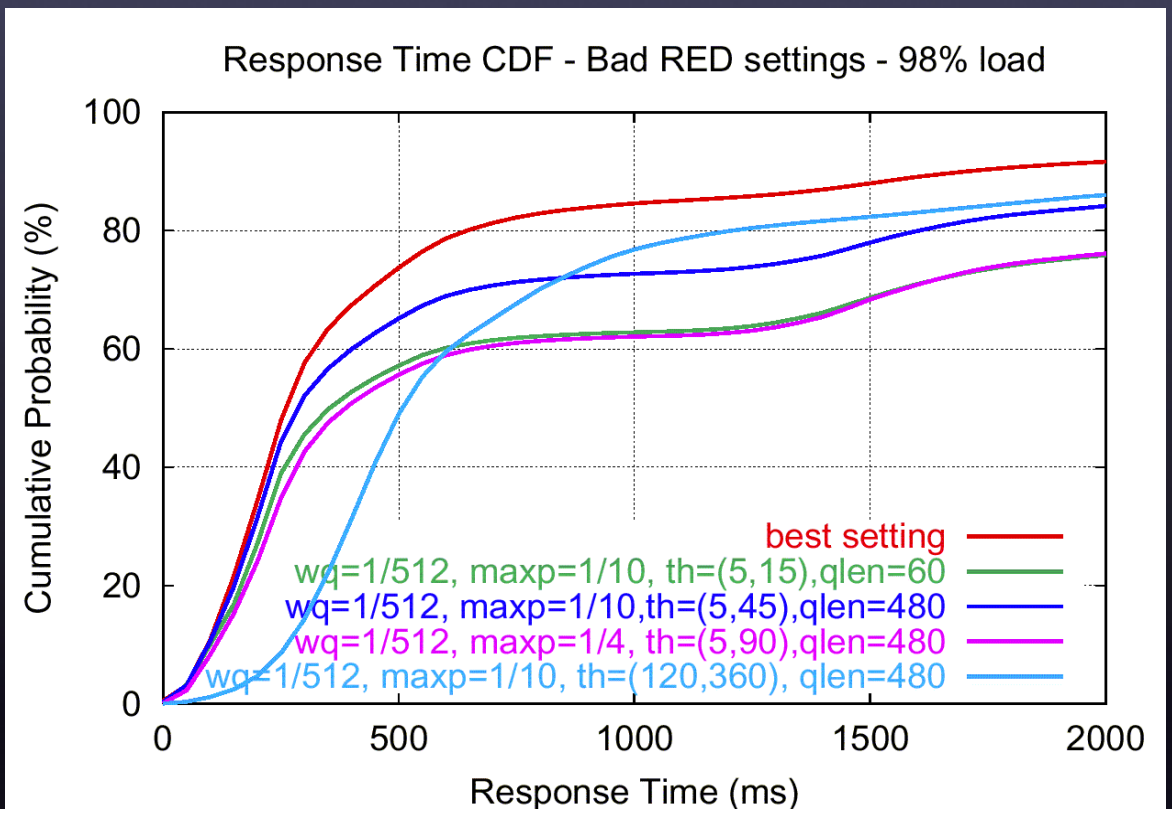

FIFO *v.* RED
Comparison

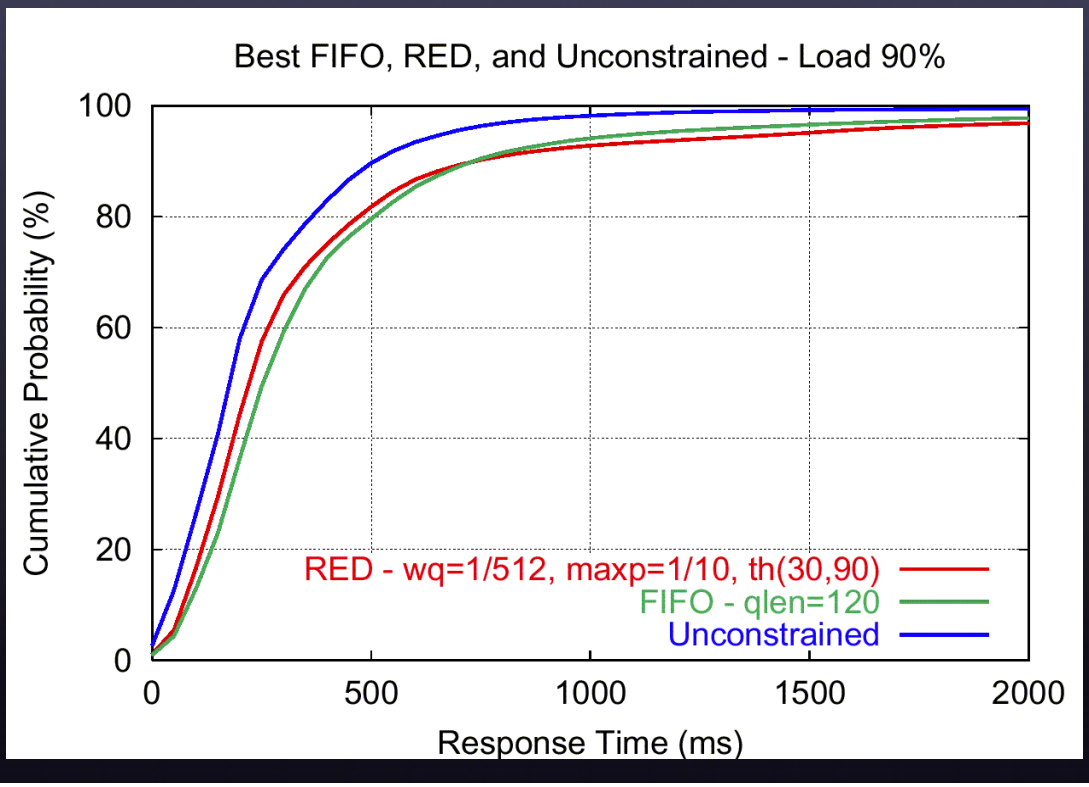

FIFO *v.* RED
Comparison

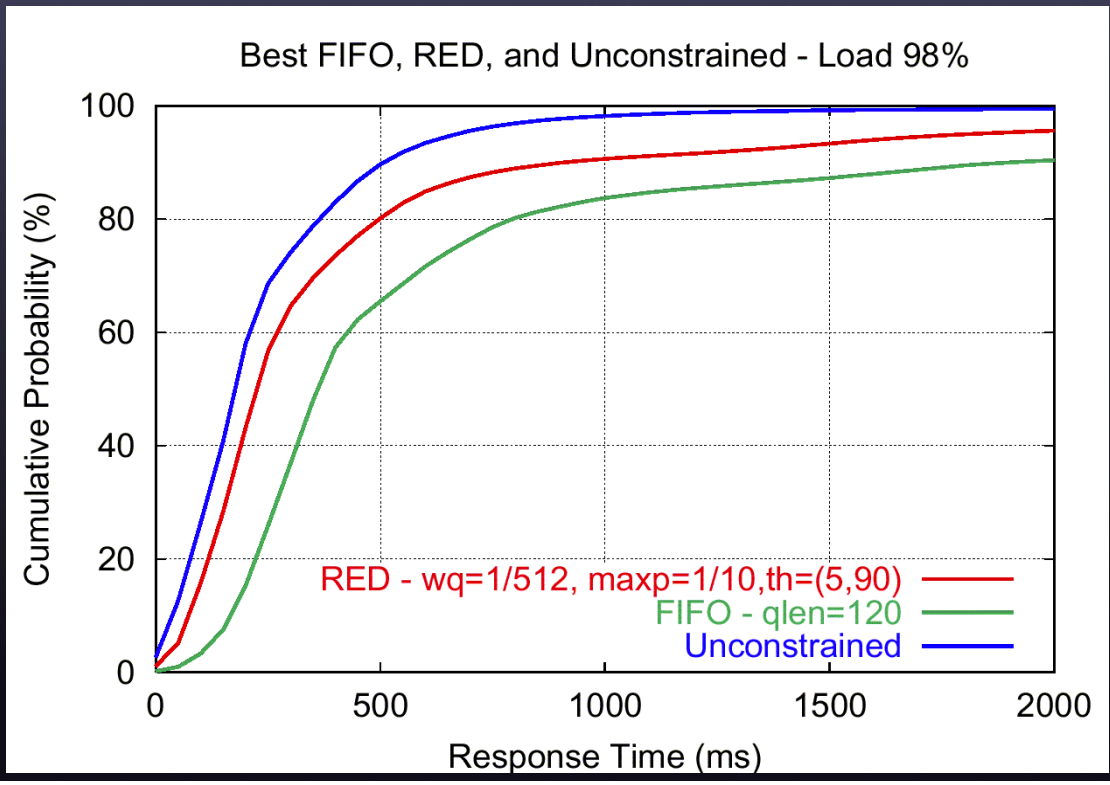

RED Response Time Analysis
Introduction

- Additional instrumentation allowing:
  - per flow statistics of:
    - » number of retransmissions
    - » type of retransmissions (SYN/FIN/Data)

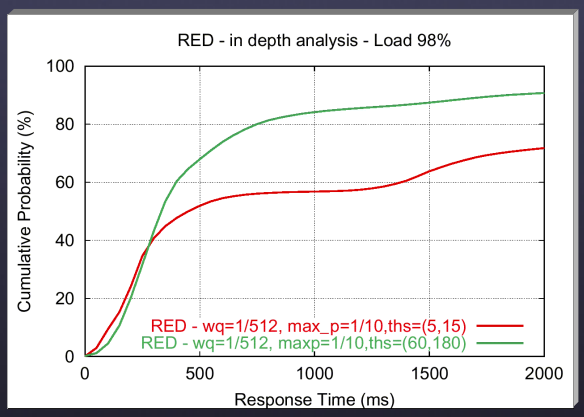

- Experiments with additional instrumentation:
  - *thresholds = (5,15) and (60,180)*

  w_q = 1/512, max_p = 1/10, and q_{len} = 480

- Offered load is set to 98% link utilization.

RED Response Time Analysis
Distribution by retransmission events

- Flows with FIN or Data retransmissions are shifted app. 1.5s
- Flows with SYN retransmissions are shifted more then 2-3s.
- Flows with multiple types of retransmissions are delayed even more.

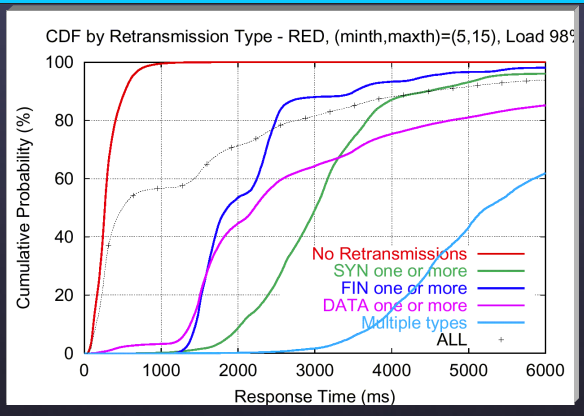

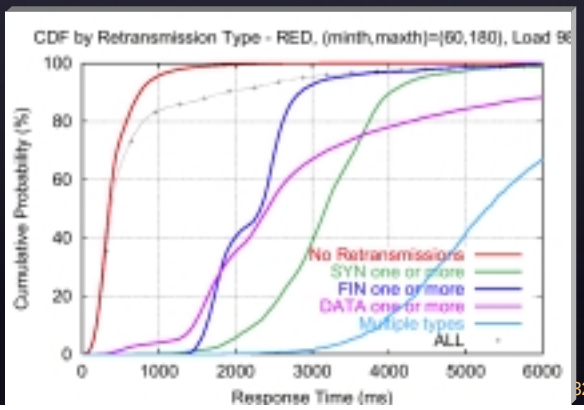

RED response time analysis
Cumulative effect

- Plots show the cumulative effect of each of the event classes combined with the *no retransmissions* class.
- Data retransmissions have the greatest cumulative effect.
- SYN and FIN retransmissions have a small but non negligible effect.

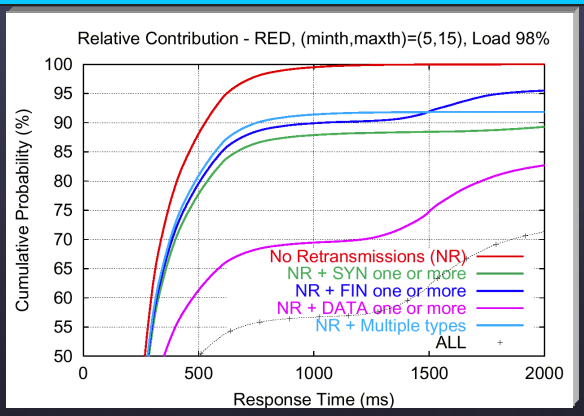

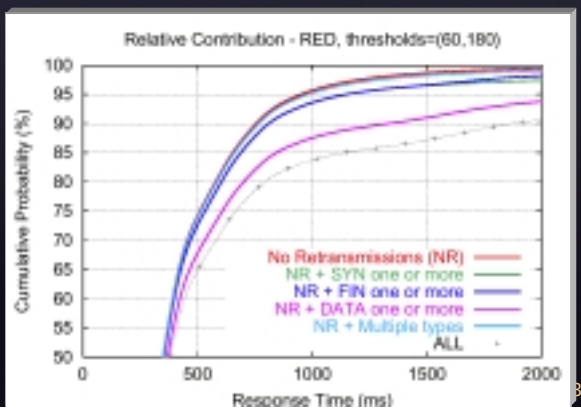

RED Parameter Determination
Scatter Plot

- Response times v. Reply Size
- One dot per connection
  - (~ 400,000 connections)
- Large influence of retransmissions on response times
- Regions of response times divided between flows with and without retransmissions

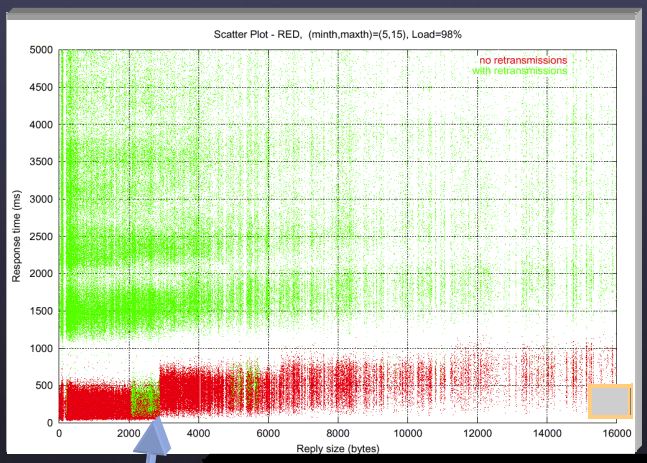

- No retransmissions
- Retransmissions

Initial window of 2.8KB