Tuning RED for Web Traffic

Mikkel Christiansen, Kevin Jeffay, David Ott, F.Donelson Smith

University of North Carolina at Chapel Hill
Department of Computer Science
Chapel Hill, NC 27599-3175

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

SIGCOMM 2000

Presented by Luba Sakharuk

Kevin Jeffay -- Fearless leader, purveyor of truth & virtue. Unfortunately, doesn't consider beer runs as progress towards degree.

David Ott -- Night owl in training. Probably sniffing your packets while you are reading this.

F. Donelson Smith -- The grown-up that we can't get out of our sandbox. Left high-paying job with IBM to join us. (So how smart can he really be?)

DiRT stands for DIstributed and Real-Time systems. Our purpose in life is to understand how to better build, you guessed it, distributed real-time systems.

“Forgive our shabby appearance but we're too busy doing cutting edge research to create a sexier home page. “

Current research projects include:

- Active queue management for multimedia networking
- TCP congestion control with GPS synchronized clocks
- Traffic measurement and characterization of the NC GigaPOP
- Adaptive, rate-based scheduling
- Proportional share resource allocation in real-time operating systems
- Tuning RED for web traffic

Tuning RED Outline

- Introduction
- Background and Related Work
- Experimental Methods
  - Experimental Network
  - Web-like Traffic Generation
  - Experiment Calibrations
  - Experimental Procedures
- FIFO Results
- RED Results
- Analysis of RED Response Times
- Comparing FIFO and RED
- Conclusions
- Future Directions

Introduction

Main idea:: **to show that RED has no effect on the performance of Web browsing**

They evaluate RED across a range of parameter settings and offered load, and their results show that:

•Compared to FIFO queue, RED has a minimal effect on HTTP response times for offered loads up to 90% of link capacity

•Response times at loads in this range are not substantially effected by RED parameters

•Between 90% and 100% load, RED can be carefully tuned to yield performance somewhat superior to FIFO, however, response times are sensitive to the actual RED values selected

•In heavy congested networks, RED parameters that provide the best link utilization produce poorer response times.

Background and Related Work

•Reviewed RED algorithm and parameters (*avg*, *qlen*, min_{th}, max_{th}, w_q , max_p) and pointed to Sally Floyd guidelines.

•RED is effective in preventing congestion collapse when TCP windows configured to exceed network storage capacity.

•Buffer size should be 1-2 times the bandwidth-delay product at a bottleneck link

•RED issues (shortcomings) studied through alternatives: BLUE, Adaptive RED, BRED, FRED, SRED, and Cisco's WRED

•ECN not considered in this paper, so results are not comparable

•INRIA :TCP goodput does not improve significantly with RED and this effect is independent of the number of flows.

•INRIA : RED has lower mean queuing delay but higher variance.

•Research elements missing : Web-like traffic and worst-case studies where there are dynamically changing number of TCP flows with highly variable lifetimes.

Experimental Network

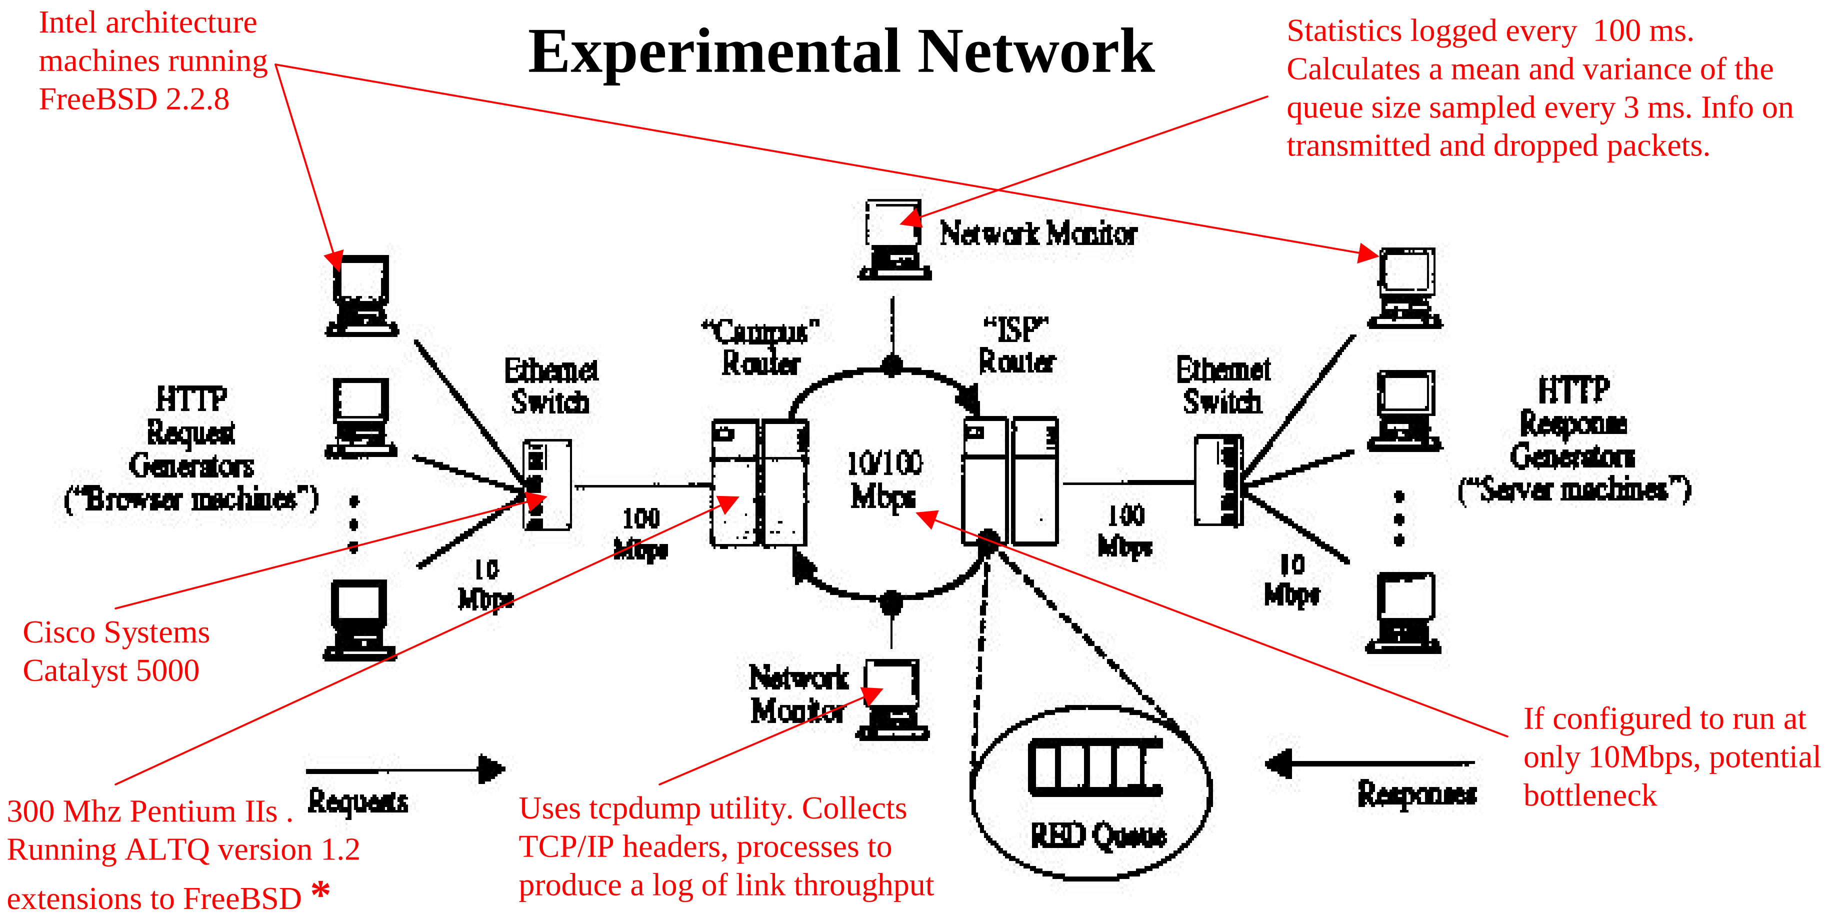

Figure 2: Experimental laboratory network diagram.

*** ALTQ extends the network-interface output queuing discipline to include FIFO, RED, CBQ, WFQ queue management**

Web-like Traffic Generation

•Used Mah model to write Web-traffic generating programs using socket system calls provided in FreeBSD

•Mah’s model is an application-level description of the critical elements that characterize how HTTP 1.0 protocols are used

•For each request, a message of random size(sampled form the request size distribution) is sent to the server program

•Response time defined as the elapsed time in milliseconds between the time of the socket connect() and the time the response is completed and the connection is closed

The elements of the HTTP model are:

- HTTP requests length in bytes
- HTTP reply length in bytes
- Number of embedded (file) references per page
- Time between retrieval of two successive pages (user “think” time), and
- Number of consecutive pages requested form a server

Experiment Calibrations

Two critical element of experimental procedures that had to be calibrated before performing experiments:

1) Ensuring no bottleneck other than when the links connecting two routers are limited to 10Mbps

2) The offered load on the network can be predictably controlled using the number of emulated browsing users

---

Note: 64 socket descriptors limitation was never encountered

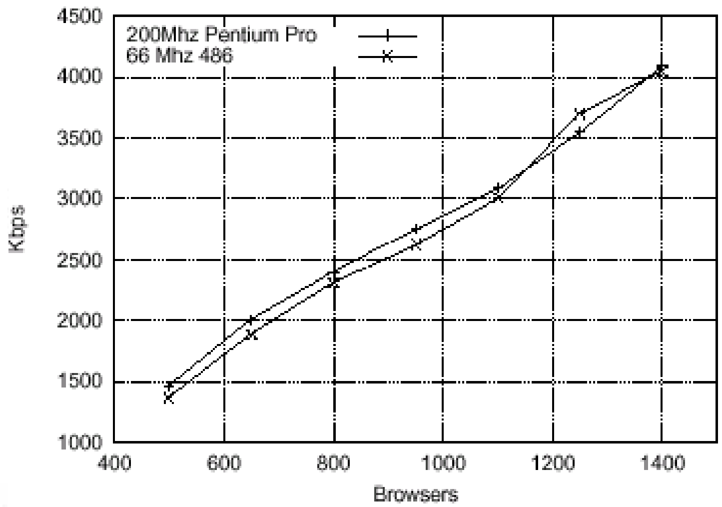

Figure 3: Offered load as a function of the number of simulated users on one machine.

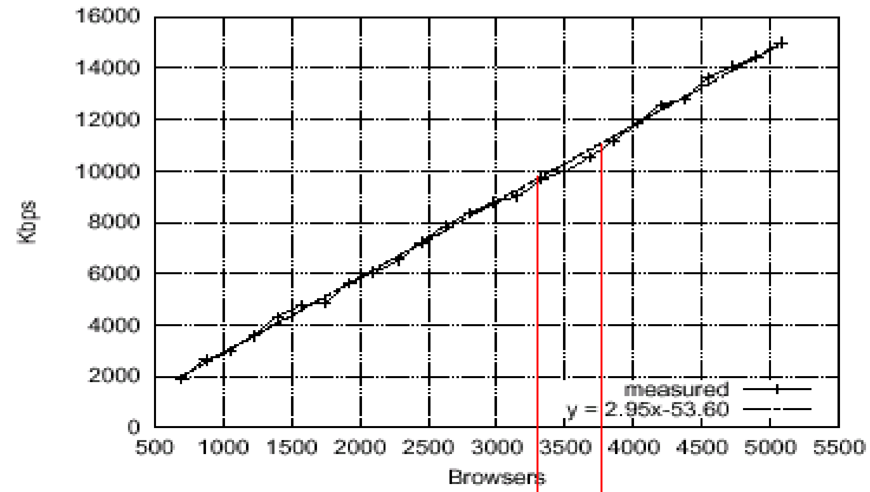

Figure 4: Offered load as a function of the number of simulated users on 7 machines.

Emulate 3750 for 110% load (11Mbps)

Emulate 3400 for 100% load (10Mbps)

Figure 3

- **CPU and interface speeds of the end system are not resource constraints**
- **If traffic generators are limited to simulating no more than 1400 users each, we can be confident that the number of users simulated in an experiment is accurate and reproducible**

Figure 4

- **No fundamental resource limitation in the system and generated load can exceed the capacity of 10 Mbps link.**
- **These data is used to determine the number of emulated browsers that would generate a specific offered load**

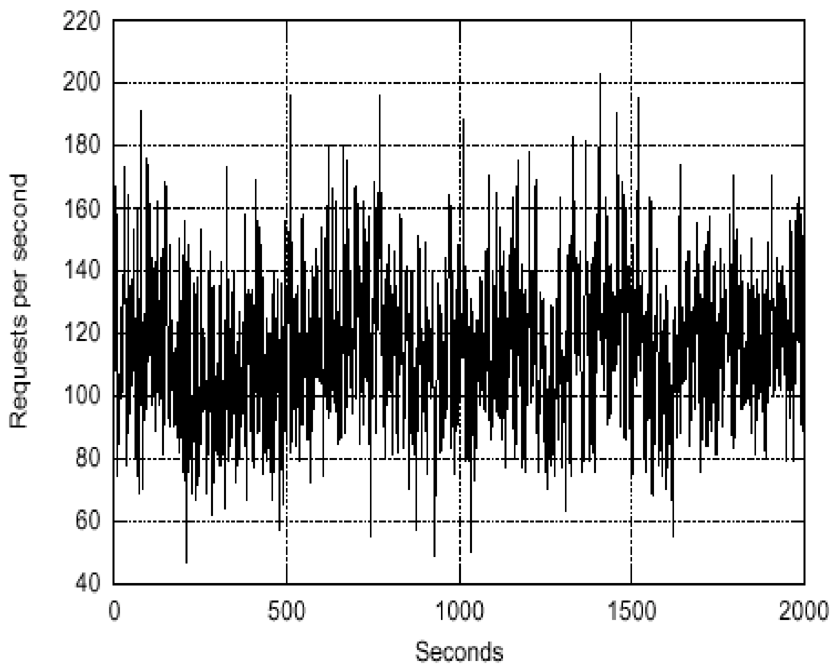

Figure 3: Requests per second from 3,500 users.

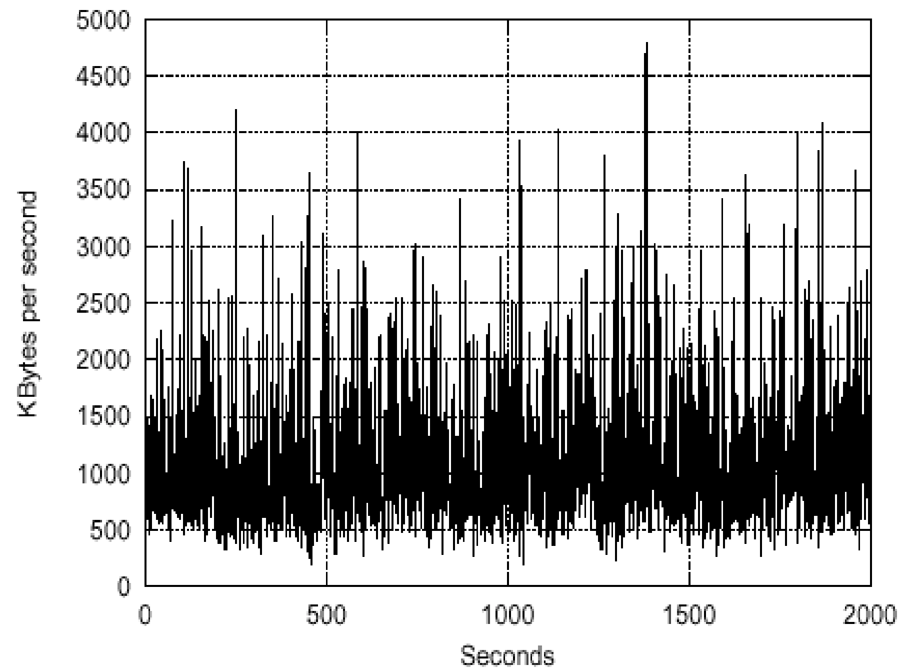

Figure 4: Bytes requested per second from 3,500 users.

•These plots show the number of requests initiated during each one second interval and the number of bytes requested in each one second interval

•These show the highly bursty nature of traffic actually generated

Experimental Procedures

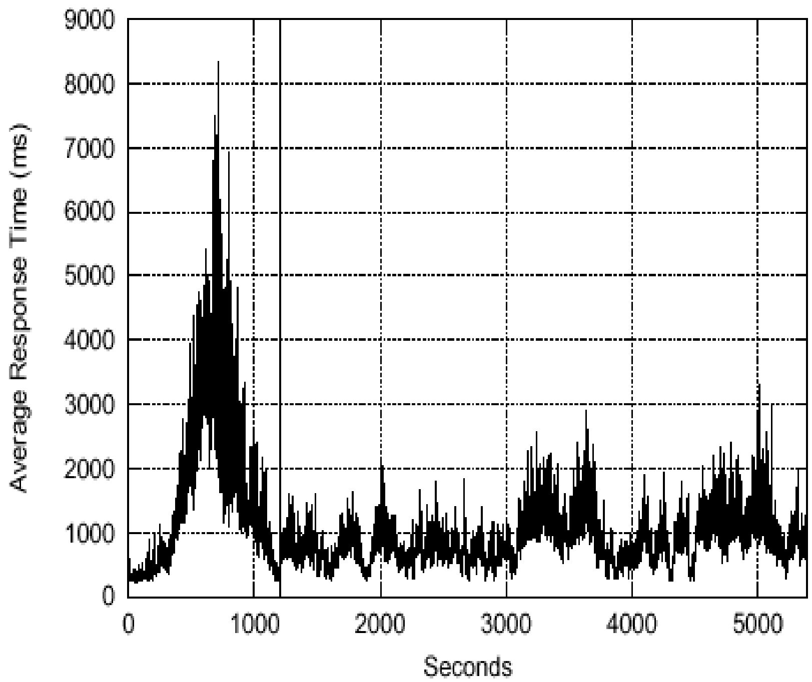

Figure 5: Average response time per second during an experiment. The plot includes the initial 20 minutes, where the traffic generators are started and stabilize.

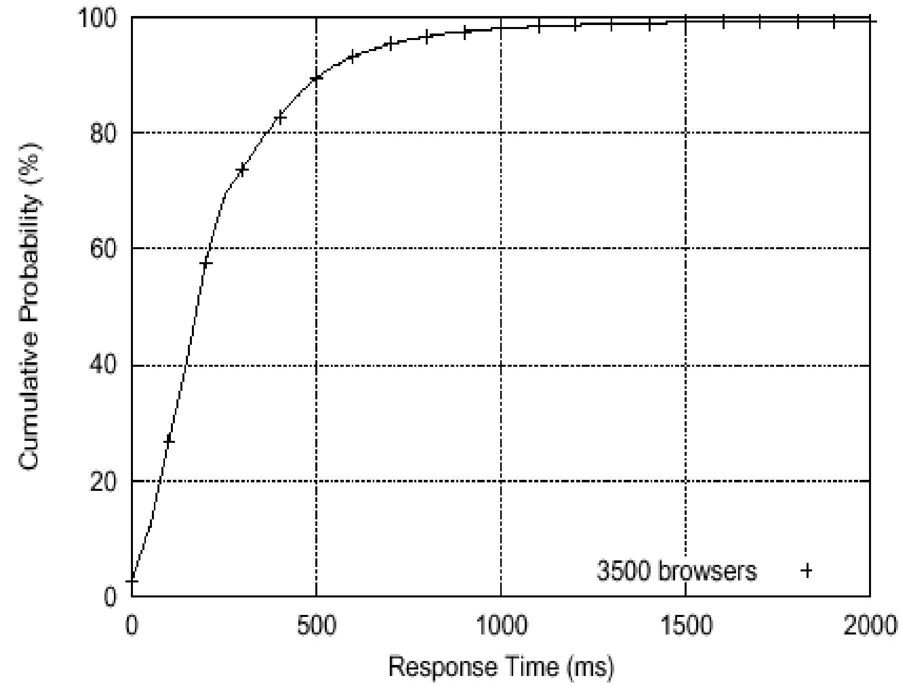

Figure 6: Cumulative response time distribution for 3,500 users on the unconstrained (100 Mbps) network.

Figure 6

•Represents the best-case performance for HTTP request/response pairs, will be used as a basis for comparison with experiments on the congested network link

•About 90% of requests complete in 500 milliseconds or less

•Decided not to consider loads beyond 110%, due to unbearable response times

FIFO Results

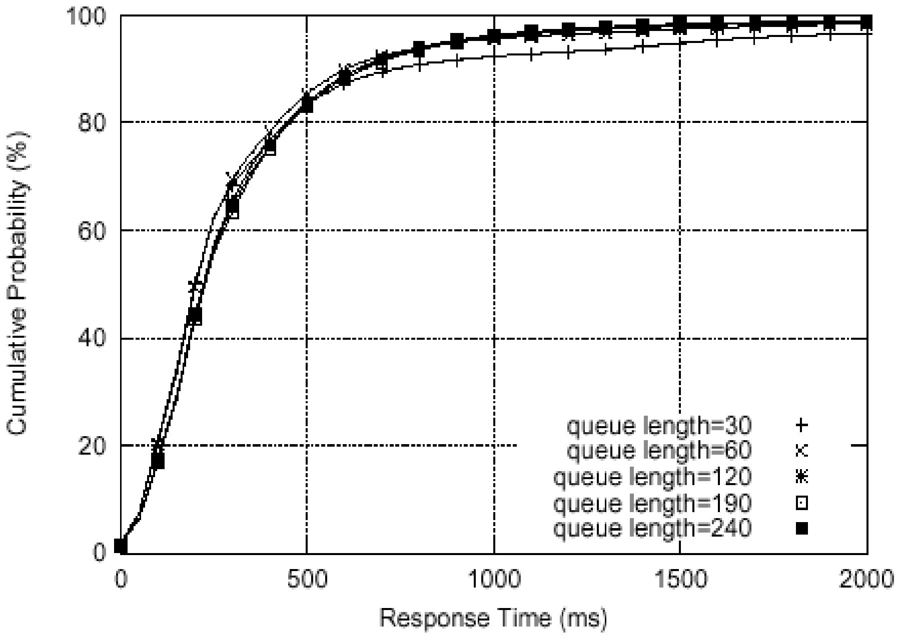

Figure 7a: FIFO performance at 80% load.

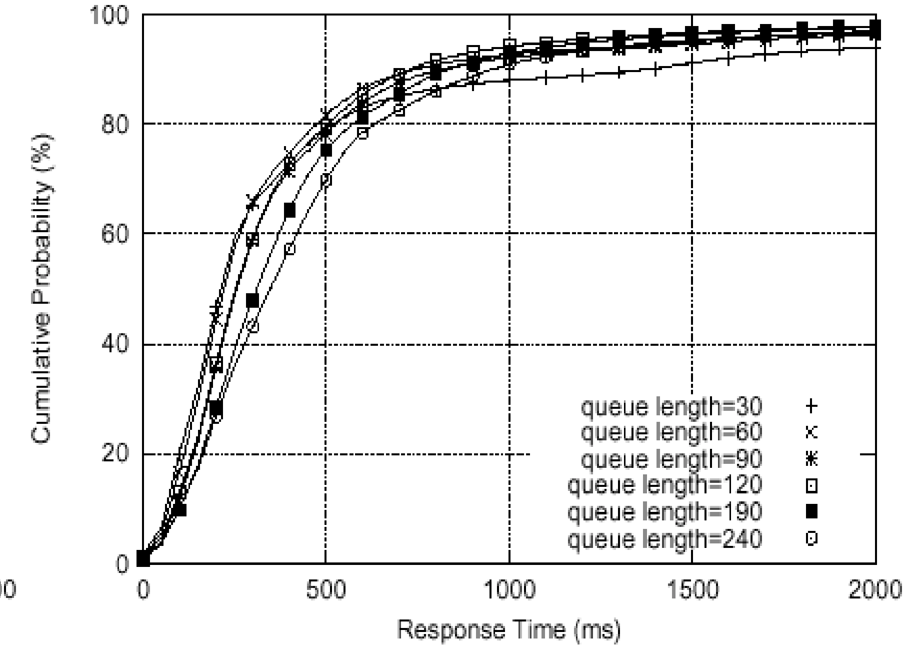

Figure 7b: FIFO performance at 90% load.

Figure 7a

- Little effect from increasing the queue size from 30 to 240 elements

Figure 7b

- Queue size has more significant effects on response times
- Queue size of 120 elements is a reasonable choice for this loading

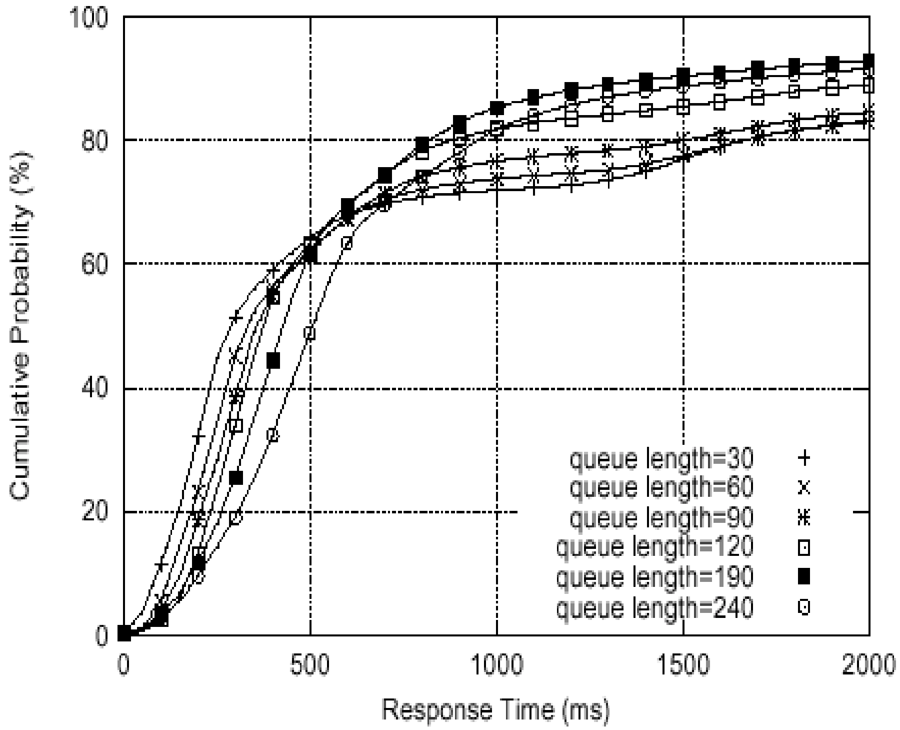

Figure 7c: FIFO performance at 98% load.

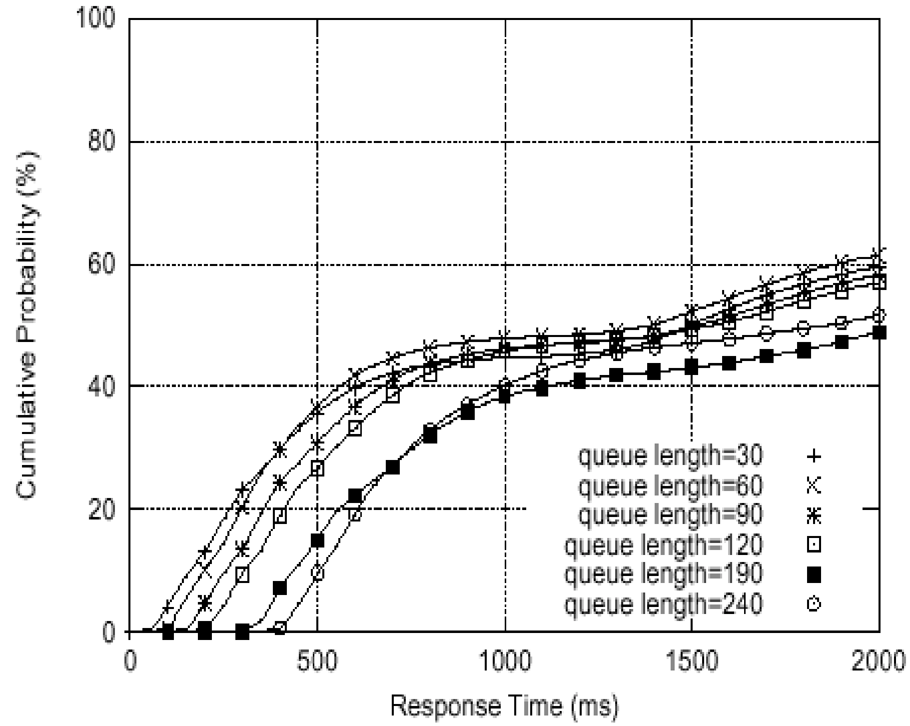

Figure 7d: FIFO performance at 110% load.

Figure 7c

Increasing the queue size from 30 to 120 has negative effect on short responses, however, it reduces response times significantly for long requests

Figure 7d

Queue sizes beyond 120 exacerbate an already bad situation

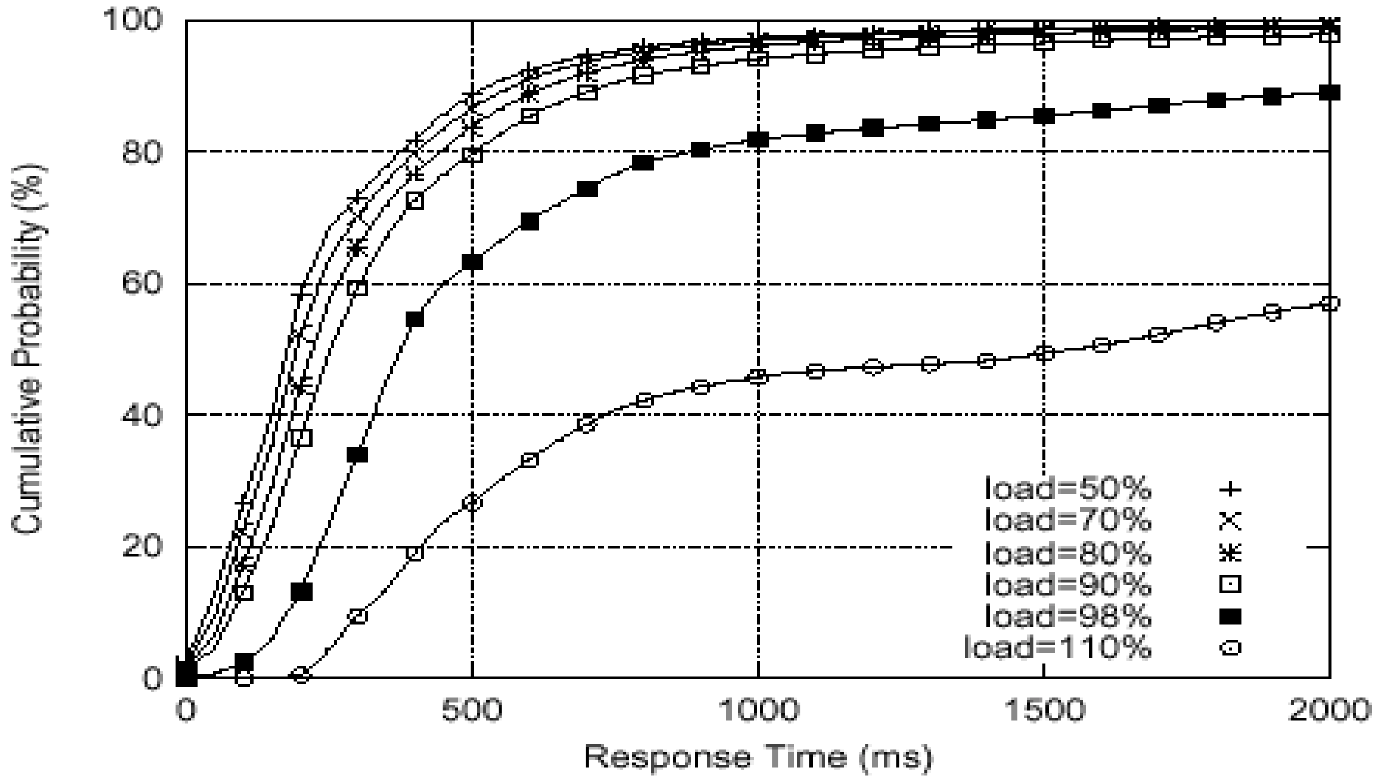

Figure 8: FIFO performance for different loads with a queue length of 120 elements.

•Response times degrade sharply when load approaches or exceeds link capacity

•Below 80%, no significant change in response times as a function of load

RED Results

•The goal for the experiments with RED was to determine parameter settings that provide good performance for Web-traffic

•Examined the effects of varying each parameter

Table 1: RED control parameters.

qlen	The maximum number of packets that can be enqueued.
min_{th}	Queue length threshold for triggering probabilistic drops.
max_{th}	Queue length threshold for triggering forced drops.
w_q	Weighting factor for the average queue length computation.
max_p	The maximum probability of performing an early drop.

Initial experiments:

•setting queue length to 480, fixing w_q at 0.002, max_p at 0.10

•varying min_{th} to 5~120

•max_{th} at 3 times min_{th}

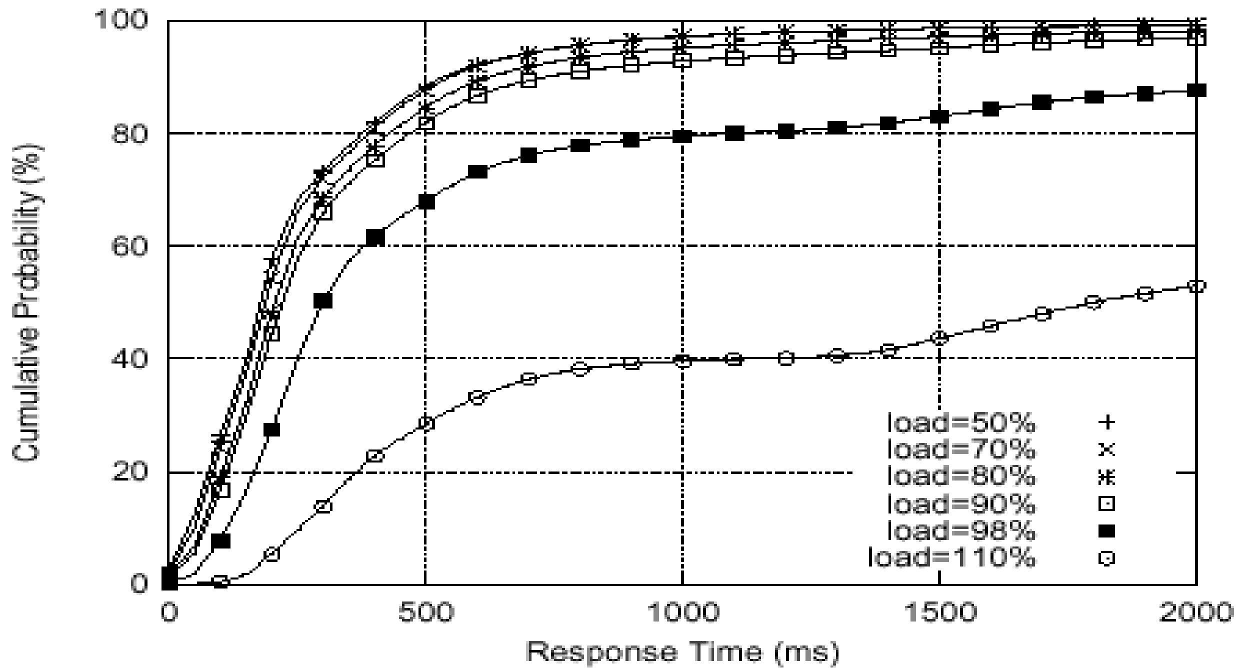

Figure 9: The performance of RED at different loads. w_q=1/512, max_p=1/10, min_{th}=30, max_{th}=90, $qlen$=480.

- Performance degradation occurs at loads greater than 70%
- Drop rates at 50% load never exceeds .01% of the packets received at the router
- Parameter tuning will have limited effect until loads reach 70-80% of link capacity
- Significant performance decrease occurs at load levels of 90-110% (most interesting)

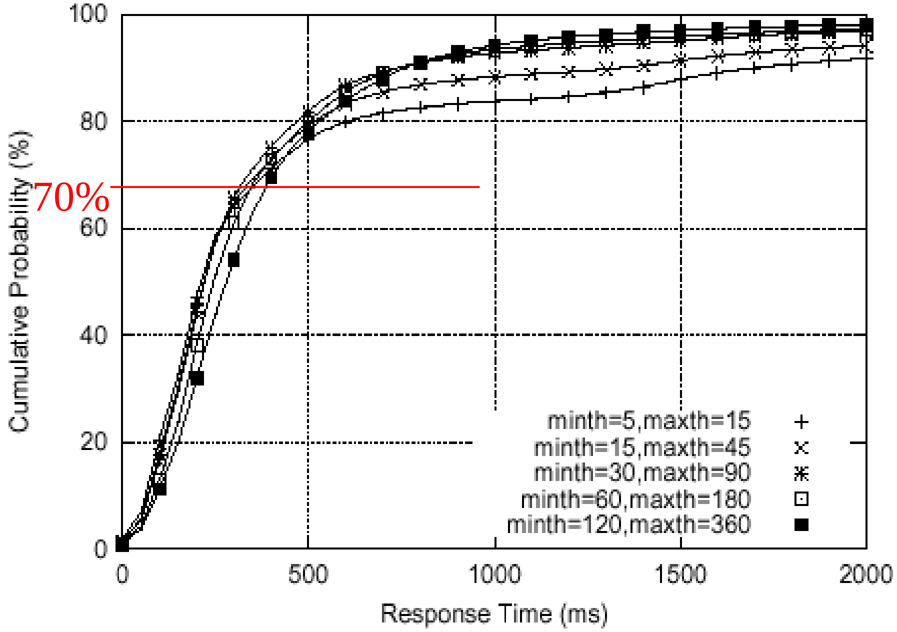

Figure 10a: Response time CDF for offered load at 90% of link capacity (w_q=1/512, max_p=1/10, *qlen*=480).

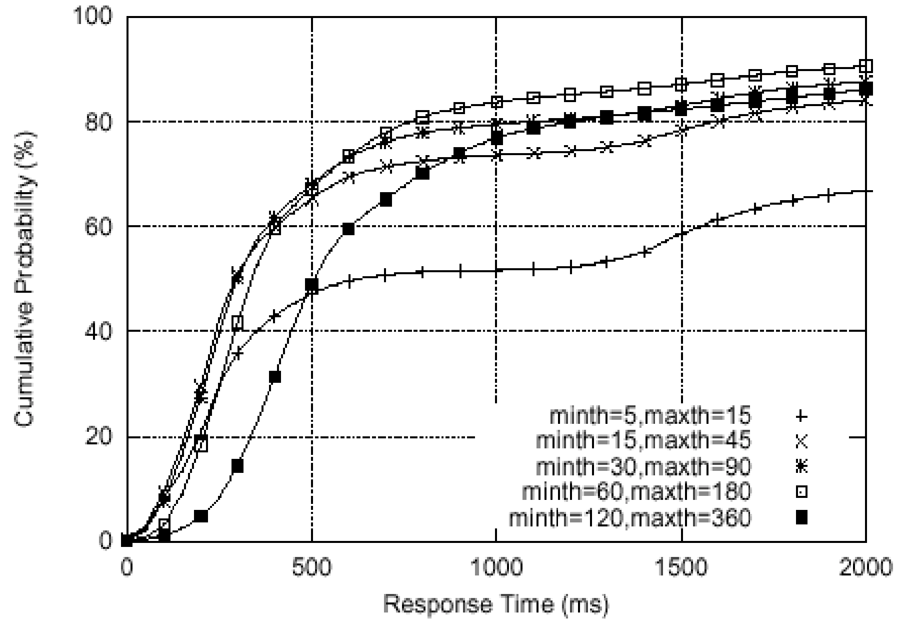

Figure 10b: Response time CDF for offered load at 98% of link capacity (w_q=1/512, max_p=1/10, *qlen*=480).

•Guidelines in …/floyd/REDparameters.txt with a min_{th} of 5 result poor performance

•About 70% of requests experience better response times with (30,90) than (60,180)

•Tradeoff between better response times for short responses at (30,90) and improving response times for longer ones at (60,180)

•Like the FIFO results, response times at loads of 110% are bad and are not improved by changing RED settings for (min_{th}, max_{th})

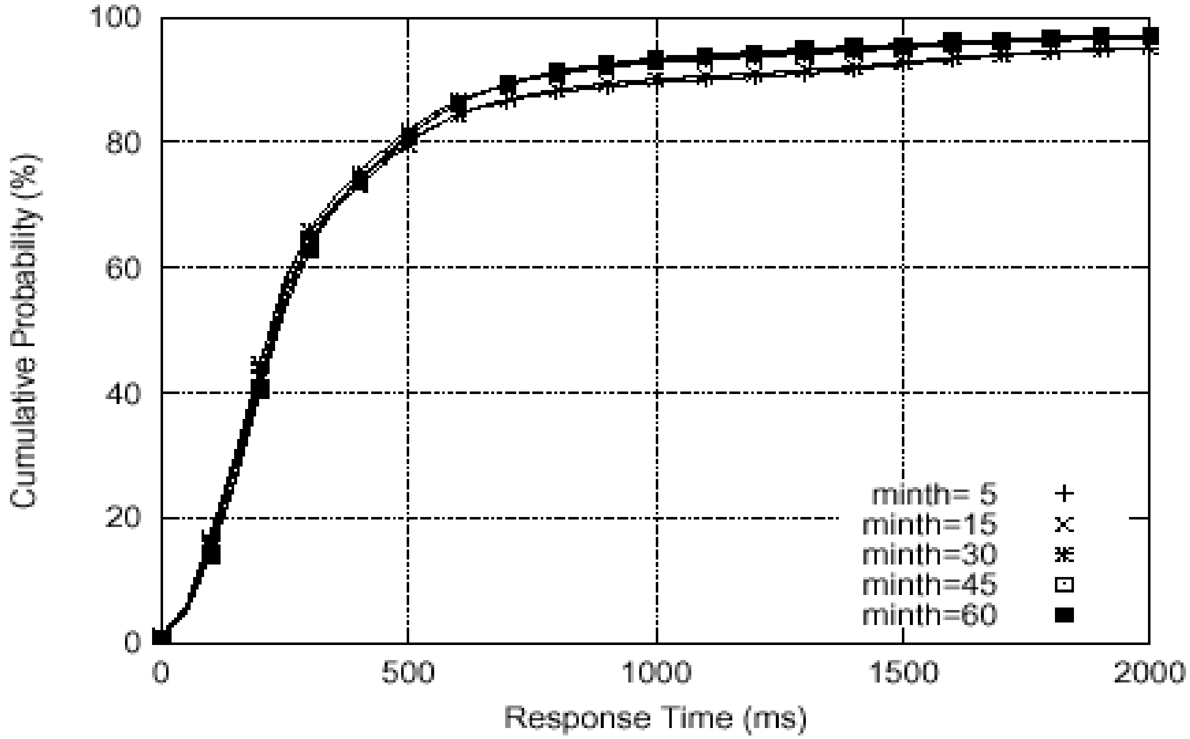

Figure 11: The effect of changing min_{th}. Load = 90% and max_{th} = 90, w_q=1/512, max_p=1/10, $qlen$=480.

•Best balance of response times for all sizes of responses (with loads considered here) are achieved with (min_{th}, max_{th})=(30,90)

•Results obtained by varying max_{th} are similar

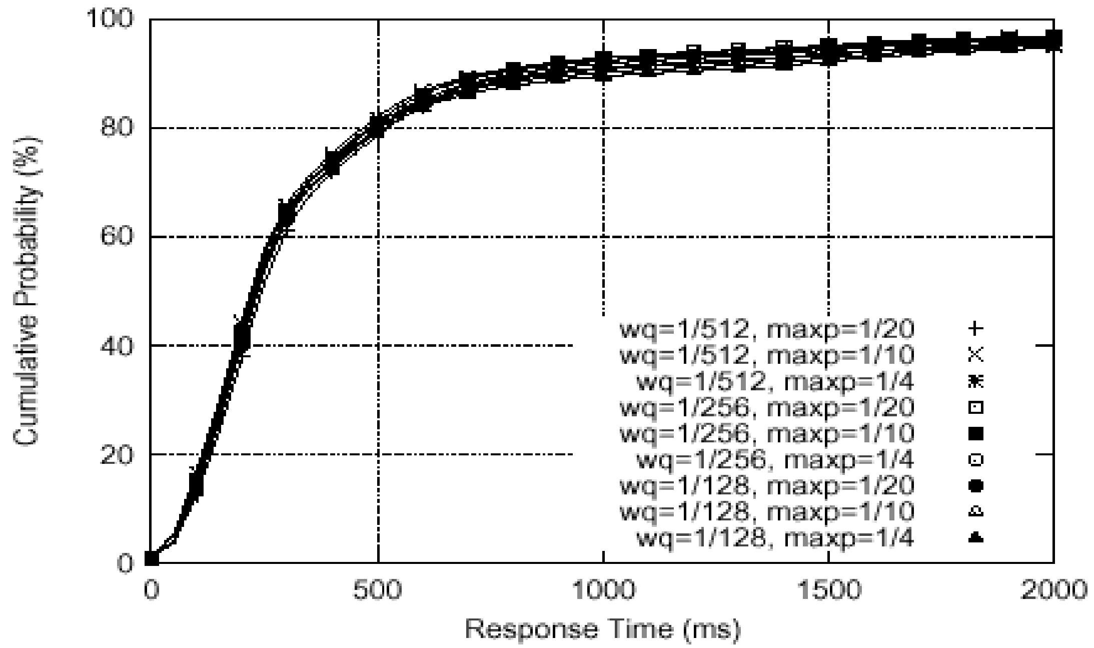

Figure 12: Results for different values of w_q and max_p. Load = 90%, and $qlen$ = 480, min_{th} = 30, max_{th} = 90.

•Decreasing w_q to 1/1024 was found to be an unrealistic setting that causes reaction to congestion to be slow

•Setting of max_p to 0.25 has a negative impact on performance(too many packets dropped)

•Changes of w_q and max_p mainly impact the longer flows

•There is no strong evidence against suggested w_q=1/512 and max_p=0.10

Table 3: RED performance with recommended parameters and queue lengths.

Load %	Queue Length	KB/s	% drop	Mean queue	Median resp.(ms)	% ≤ 1 sec	1<%≤ 2 sec	2<%≤ 3 sec	% > 3 sec
90	480	1079	0.8	20.2	266	92.5	4.3	2.0	1.3
90	160	1093	1.1	22.2	278	91.2	4.7	2.4	1.7
90	120	1066	0.7	18.8	266	93.0	4.1	1.7	1.2
98	480	1164	4.1	39.4	345	79.2	8.2	6.3	6.3
98	160	1175	5.9	46.3	397	72.4	9.7	8.2	9.7
98	120	1171	5.5	44.3	377	74.2	9.2	7.7	8.9
110	480	1187	19.7	76.0	1846	39.4	12.9	12.1	35.5
110	160	1188	19.5	76.6	1864	39.1	13.0	12.2	35.7
110	120	1188	18.9	77.0	1840	39.3	13.2	12.5	34.9

- Results similar to the FIFO - the 120 element queue (1.25 times bandwidth delay) is a reasonable choice at 90% and 110% loads
- Longer queues of 2-3 times bandwidth-delay might provide some advantage at load just below link saturation
- min_{th} should be set to larger value to accommodate the bursty character of Web traffic
- Attempting to tune RED parameters outside the guidelines is unlikely to yield significant benefits.

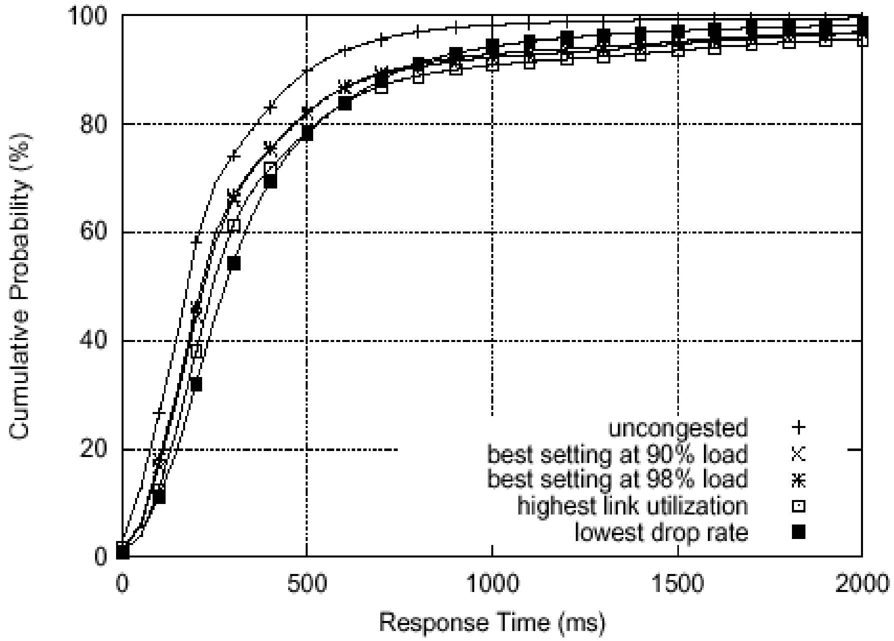

Figure 13a: "Good" RED parameter settings at 90% load.

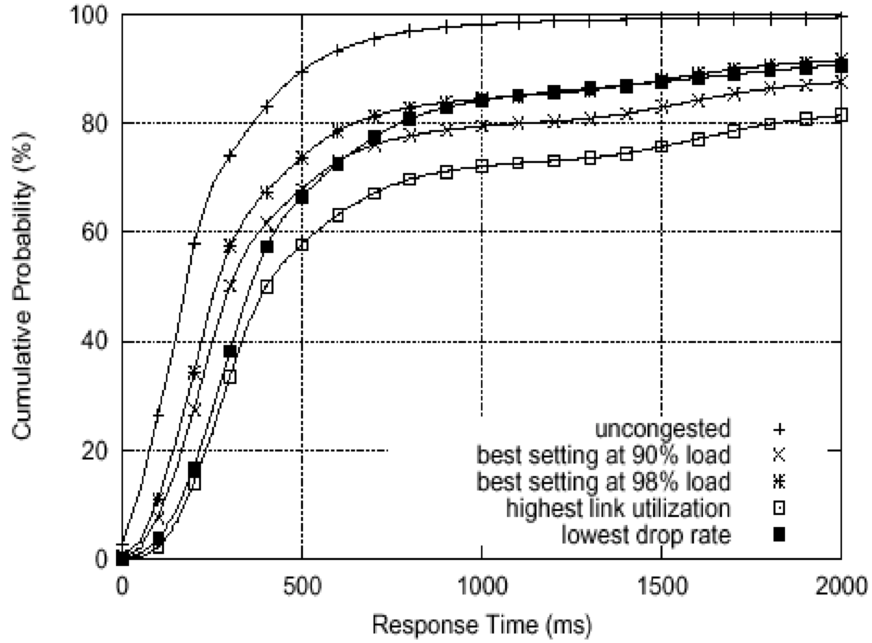

Figure 13b: "Good" RED parameters settings at 98% load.

Figure 13a

Relatively small differences between tuning for highest link utilization or lowest drop rate and tuning for response times

Figure 7b

Tuning for highest link utilization has potentially serious effects of increasing response times

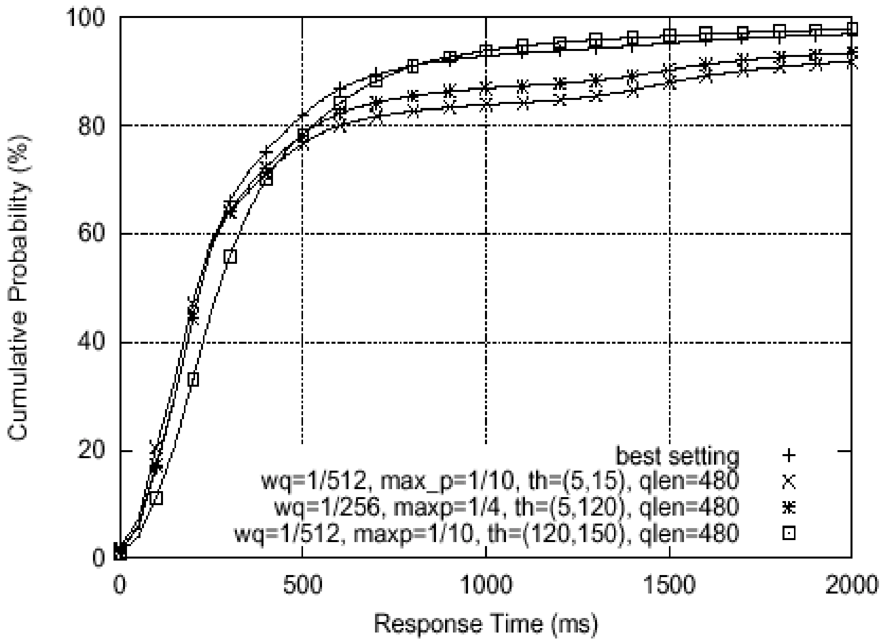

Figure 14a: "Bad" RED parameters settings at 90% load.

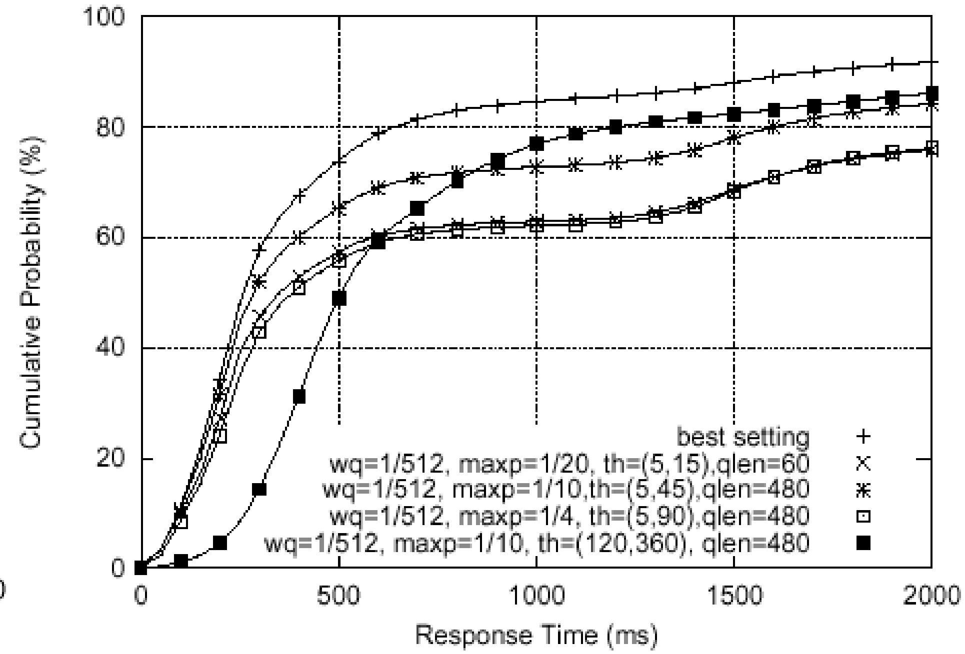

Figure 14b: "Bad" RED parameters settings at 98% load.

Table 4: Empirically determined "best" RED parameter values.

Load	min_{th}, max_{th}	w_q	max_p	Notes
90	30,90	1/512	1/10	best overall response
90	30,90	1/512	1/20	highest link utilization
90	120,360	1/512	1/10	lowest drop rate
98	5,90	1/128	1/20	best overall response
98	30,180	1/512	1/10	highest link utilization
98	90,150	1/512	1/10	lowest drop rate

Analysis of RED Response Times

- Detailed analysis of retransmission patterns for various TCP segments (e.g., SYN, FIN)
- This section reinforces the complexity of understanding the effects of RED for HTTP traffic

Summary retransmission statistics for experiments with more detailed instrumentation

Class of retransmission event	% of all TCP connections (min_{th}, max_{th}) = (5,15)	(60,180)
No retransmission	56.1	87.1
1 or more retransmissions	43.9	12.9
1 or more SYN segments	7.4	2.0
1 or more FIN segments	6.0	2.0
1 or more data segments	25.5	8.5
Combined SYN/FIN/data	5.0	0.4
Total TCP connections	439,979	460,022
Total segments lost	12.4%	2.4%

Comparing FIFO and RED

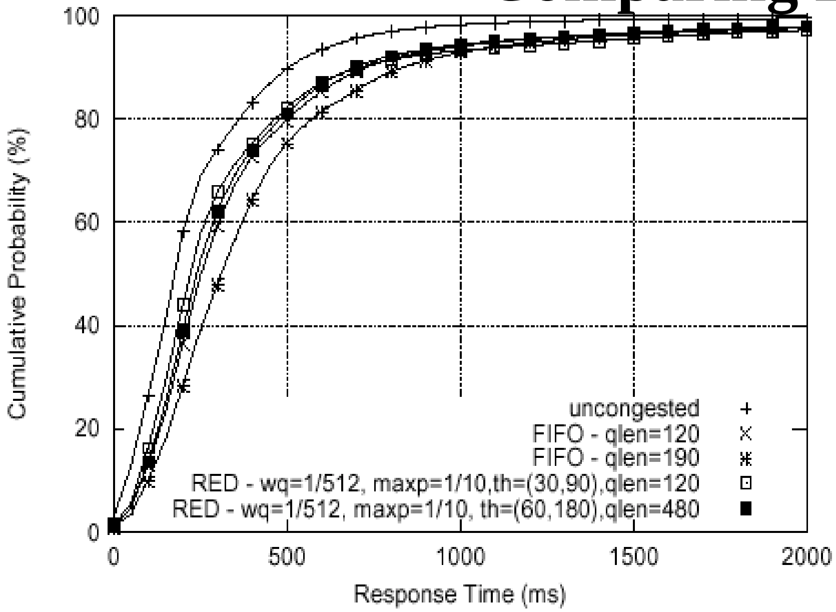

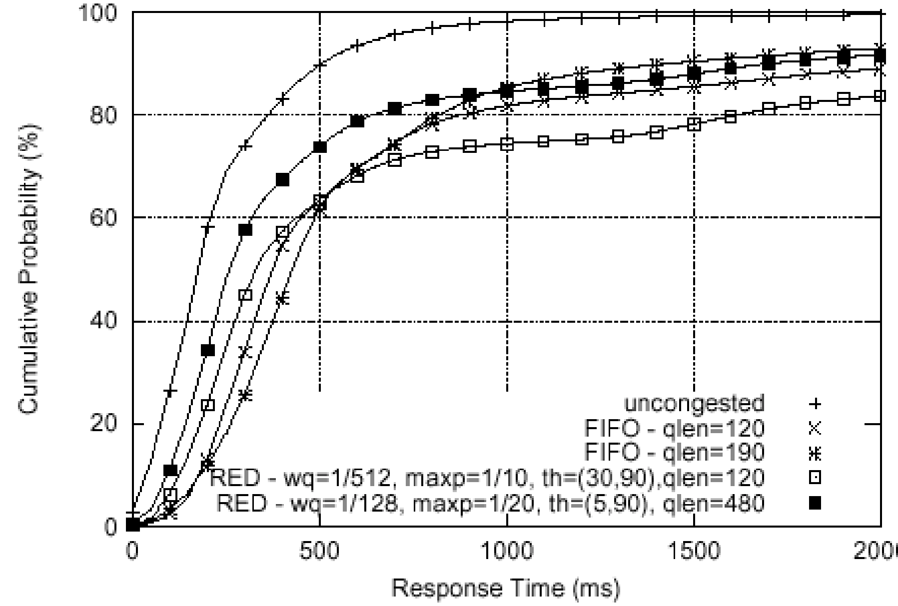

Figure 15b: FIFO and RED at 98% load.

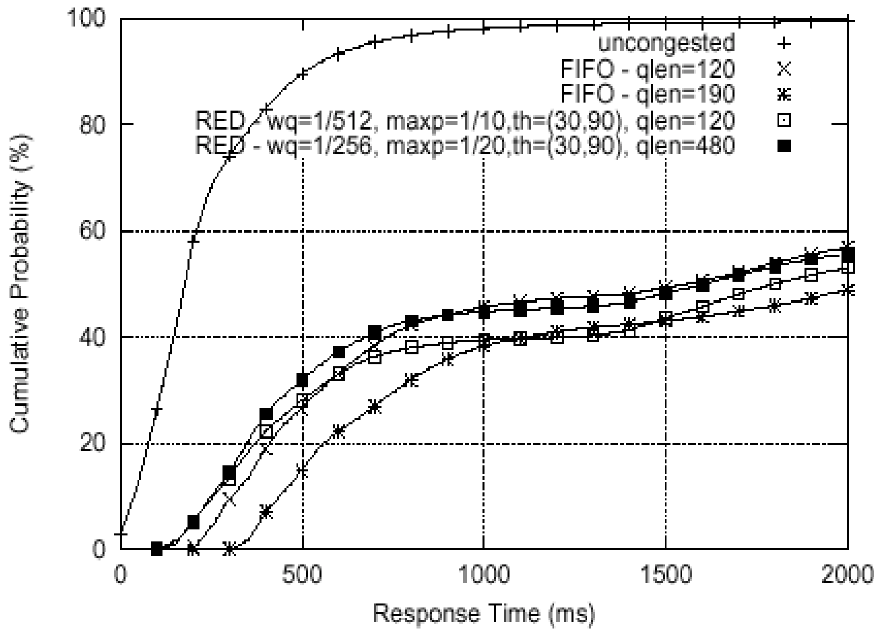

Figure 15c: FIFO and RED at 110% load.

The only distinct advantage from using RED is at the 98% where response times for shorter responses (80% of requests) are improved with carefully tuned RED parameters

Conclusions

•Tuning of RED parameters produce little gain (or loss) in response time performance

•It is possible to produce poorer performance

•Providing adequate link capacity (utilization less than 90%) is far more important for Web response times than tuning queue management parameters

•There seems to be no advantage to RED deployment on links carrying only Web traffic

More Conclusions

•Compared to FIFO queue, RED has a minimal effect on HTTP response times for offered loads up to 90% of link capacity

•Response times at loads in this range are not substantially effected by RED parameters

•Between 90% and 100% load, RED can be carefully tuned to yield performance somewhat superior to FIFO, however, response times are sensitive to the actual RED values selected

•In heavy congested networks, RED parameters that provide the best link utilization produce poorer response times.

Future Directions

•They used packet-drops as the only "marking" behavior of RED. Explicit marking by RED for ECN-capable TCP implementations is likely to produce better results

•They examined only HTTP 1.0 protocols. The interaction of RED with a mix of HTTP 1.0 and HTTP 1.1 traffic should also be analyzed

•They studied a link carrying only Web-like traffic. More realistic mixes of HTTP and other TCP traffic as well as traffic from UDP-based applications need to be examined

•Congestion on both paths on a full-duplex link and over multiple router hops should also be considered

"The DiRT Method of Conflict Resolution"

"Spot the New Asst. Prof"

What's the difference between a new prof (Ketan Mayer-Patel) and a 10-year graduate student (Mark Parris)? Clearly neither of them do any work!

"The DiRT Method of Conflict Resolution"