

Impact of bandwidth-delay product and non-responsive flows on the performance of queue management schemes

Zhili Zhao
Dept. of Elec. Engg., 214 Zachry
College Station, TX 77843-3128

A. L. Narasimha Reddy
Dept. of Elec. Engg., 214 Zachry
College Station, TX 77843-3128

Abstract

In this paper, we study the impact of bandwidth-delay products and non-responsive flows on queue management schemes. Our focus is on understanding the aggregate performance of various classes of traffic under different queue management schemes. Our work is motivated by the expected trends of increasing link capacities and increasing amounts of non-responsive traffic. In this paper, the impact on the performance of RED, RED with ECN enabled and DropTail routers is investigated. Our study considers the aggregate bandwidth of different classes of traffic and the delays observed at the router.

I. INTRODUCTION

A. Motivation of Our Work

Current network traffic consists of 10% - 20% UDP traffic and 80% - 90% TCP traffic [1]. Traffic can be further classified as long-term or short-term. Short-term TCP flows (STFs), usually referred as web mice, occupy about 30% of current network traffic. Long-term flows, UDP or TCP, still contribute the majority of the load in current networks.

Traditionally, long-term flows have been used to model network traffic in conducting experiments for the performance evaluation of queue management schemes. With the increased knowledge of the important role of STFs on the network performance, STFs have gained growing attention of researchers. Previous research has shown that: 1) the aggregated behavior of STF at the router appears to be non-responsive [2], [3]; 2) the performance of queue management schemes is different under a mixed workload, including STFs, when compared to workloads consisting of only long-term flows [4], [5].

With the growth of multimedia applications over the network, such as streaming audio/video and Voice over IP (VoIP), UDP traffic is expected to increase in the future. Some of these applications may be designed not to respond to congestion for delivering packets in time or for assumed need for maintaining constant bandwidth. Those non-responsive flows have been shown to impact goodput of the network and the performance realized by other responsive applications [6]. This trend in workloads

points to a need for studying the impact of increasing non-responsive UDP loads on the network performance. For example, does the early dropping of packets in RED allow non-responsive flows to force the responsive flows to realize smaller throughputs?

The link capacities are increasing at a steady rate. Larger link capacities allow the deployment of new applications with higher bandwidth demands. The terabit grid is an example of such a network. Typically, the routers are configured with buffers, the size of which is determined based on the link bandwidth-delay product (BWDP). With the increased link capacities, BWDP or the required buffer size also grows. Each interface card of a router may not be able to hold that much buffer memory at higher link capacities. So it is possible that future routers may be configured with smaller amounts of buffer than the traditional rule of 1 BWDP. This points to the need to study the impact of buffer sizes on the performance of various classes of traffic. If the routers are configured with smaller buffer sizes in the future, does this impact the decisions on buffer management algorithms?

The study in this paper tries to address the resulting issues of these two trends of increasing link capacities and increasing non-responsive loads. We are motivated by the following questions: (a) what impact do higher non-responsive loads have on different queue management schemes?, (b) is one queue management scheme better at protecting responsive traffic over the others?, (c) are there differences in the performance of queue management schemes at different buffer sizes?, and (d) can we observe any discernable trends in the performance of queue management schemes with the expected trends in workloads and link capacities?

In this paper, we study the performance of three different queue management schemes, DropTail (DT), RED and RED with ECN enabled (RED-ECN) under different workloads, link capacities and buffer sizes. The workloads have different fractions of Long Term Non Responsive Flows (LTNRFs), Long Term Responsive Flows (LTRFs) and Short Term Flows (STFs). We consider different performance metrics, such as realized throughput for responsive flows, delay and link utilization.

B. Related Work

In [5], authors analyze TCP goodput, loss rate, and average queueing delay and deviation by changing number of LTRFs and STFs under a fixed RED/DT configuration. In [7], authors analyze TCP average throughput, number of bytes sent, and UDP loss% by changing RED parameters, however, those changes hardly reveal the relation between RED configurations or buffer sizes and TCP throughput. And in both papers, a fixed 10% UDP load is used.

In [4], authors illustrate the impact of STFs on RED queue dynamics and stress the importance to consider STFs while conducting any network modeling or simulation.

In [8], authors investigate the RED performance with different configurations under only web traffic. The response time is the performance metric. Authors compare the CDF of response times between RED and DT and concluded that the performance difference is fairly small between them and that tuning RED configurations gains little in performance.

The rest of the paper is organized as follows: In section II, the impact of non-responsive flows and bandwidth-delay product over droptail and AQM schemes is investigated; In section III, simulation results are summarized and discussed.

II. PERFORMANCE EVALUATION

In this section, the impact of bandwidth-delay product and non-responsive flows on the performance of DT, RED and RED-ECN is investigated under various configurations of workloads, buffer sizes and link capacities.

Different traffic workloads are considered in our evaluation. The workloads differ in the amount of the long-term non-responsive load and the amount of load from short-term flows. We consider workloads with both long-term only flows and a mixture of long-term and short-term flows, so we can study the impact of short-term flows on both long-term flows and buffer management schemes. We consider high loads of LTNRFs, from 30% to 60% of the total traffic load, to study their impact on the buffer management schemes.

We study the impact of link capacities by employing three different link capacities of 5Mb, 35Mb and 100Mb. This is expected to allow us to study the performance trends over a range ($\times 20$ times difference) of link capacities. The number of flows (STFs, LTRFs, and LTNRFs) is scaled correspondingly (based on the capacities) to result in similar workload mix from different classes of traffic. For each link capacity, we considered three different buffer sizes of 1/3, 1 and 3 times the BWDP. These buffer sizes are chosen to study the impact of under provisioning and over provisioning as well as the normal rule of provisioning the buffer sizes.

The realized throughput of responsive flows (long-term TCP throughput), along with average queueing delay, link utilization and standard deviation of queueing delay, are considered as performance metrics.

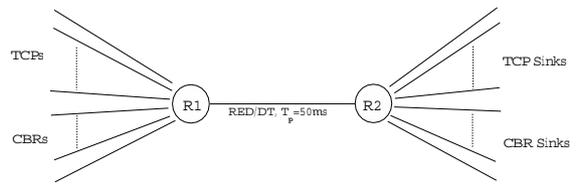


Fig. 1. Simulation Topology

TABLE I
BUFFER SIZES OF DIFF. LINK CAPACITIES AND BWDPs

Multiple of BWDP	Link Capacity (Mb)		
	5	35	100
1/3	25	200	500
1	75	500	1500
3	225	1500	4500

A. Simulation Setup

All simulations are conducted in ns-2 simulator [9]. Simulation topology is shown in Figure 1. The link between R1 and R2 is the bottleneck link, deployed with DT, RED or RED-ECN. Link capacities employed at the bottleneck link are 5Mb, 35Mb and 100Mb. Rest of the links have capacities high enough to avoid packet drops. T_p is the one-way propagation delay of the bottleneck link. The one-way propagation delay of each ingress link of R1 and each outgoing link of R2 is $5ms$. So the total round trip propagation delay is $120ms$.

Based on the link capacity and the round trip propagation delay, 1 BWDP is calculated and rounded up in units of packets (1 packet=1000 bytes), as shown in the second row of Table I. Buffer sizes under 1/3 and 3 BWDPs cases are also listed in Table I.

RED or RED-ECN is deployed at the bottleneck link as a typical AQM scheme. It is configured by following the recommendations of [10]: $Max_{th} = 3 * Min_{th}$; $Max_p = 0.1$; $w_p = 0.002$. Its performance then is compared to that of a droptail router of the same buffer size based on different performance metrics.

TCP flows include both long-term responsive and short-term flows. Long-term TCP connections(FTP) represent LTRFs. Short TCP connections, sending 10 packets every 10s on average, represent typical STFs(web mice). 0%, 5%, and 30% STF loads are generated to change the proportion of traffic mix. CBR flows represent typical LTNRFs. Each CBR flow sends at 1Mbps under 35Mb and 100Mb links and 0.5Mbps under 5Mb link for easily adjusting long-term non-responsive load across different link capacities.

In current Internet traffic, total non-responsive load of LTNRFs and STFs is about 40-50%. We intend to investigate scenarios when the non-responsive load is high as explained in Section I, so total non-responsive load, including loads of both LTNRFs and STFs, is set to 60%. 60%, 55% and 30% LTNRF loads are generated corresponding to respective STF loads. A fixed number of LTRFs are employed to complete the traffic mix. The packet size is chosen to be 1000 bytes for all flows. Table II shows the actual number of LTRFs, LTNRFs and STFs deployed

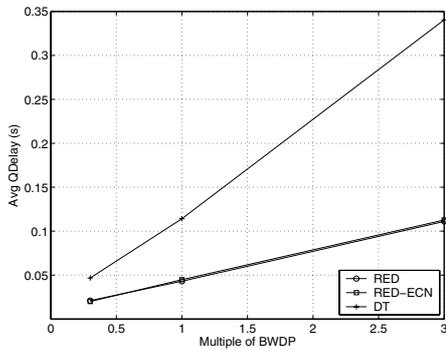


Fig. 2. Avg QDelays with Diff. Multiples of BWDP (30% STF Load and 35Mb Link)

in each simulation with different STF loads under 35Mb Link. For 5Mb and 100Mb links, the number of flows are scaled down or up according to the link capacity.

TABLE II

CHARACTERISTICS OF DIFF. WORKLOADS FOR 35Mb LINK

STF Load	35Mb Link		
	# of LTRFs	# of STFs	# of LTNRFs ¹
0%	55	0	22
5%	55	250	22
30%	55	1300	14

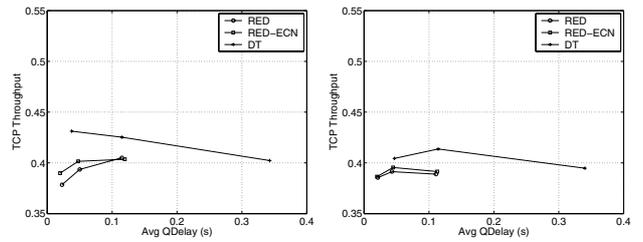
¹ Each LTNRF sends at 1 Mbps

B. Changing buffer sizes

By changing buffer sizes, realized long-term TCP throughput, link utilization and standard deviation of queueing delays are collected for analysis. 0%, 5% and 30% STF loads were generated. The performance impact of 30% STF load with different buffer sizes is presented and analyzed in the following simulations.

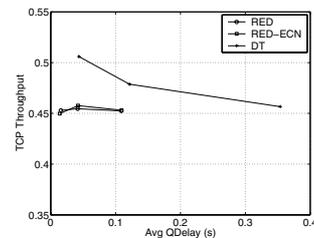
Figure 2 shows the correlation between average queueing delays and different multiples of BWDP (buffer sizes) under 30% STF load and 35Mb link. It is noticed that the average queueing delay linearly increases with the increase of the buffer size. Because of this linear relationship, in order to clearly illustrate the performance difference among different buffer management schemes with more information in each plot, average queueing delays (instead of buffer sizes) are used.

In Figure 3, realized long-term TCP throughput and average queueing delays are shown in different configurations under 30% STF load. It is noticed that TCP throughput of a droptail router is always higher than that of either RED or RED-ECN. But with the increase of BWDP, the throughput difference gets smaller. Average queueing delay of the droptail router under higher BWDPs is more than 3 times higher than that of RED under the same BWDP. It is undesirable to have high average queueing delays, especially for delay sensitive multimedia applications. So, under large BWDPs, RED or RED-ECN is a better choice considering the realized TCP throughput and lower average



(a) Link Cap.=5Mb

(b) Link Cap.=35Mb



(c) Link Cap.=100Mb

Fig. 3. TCP Throughput with Diff. Avg QDelay under 30% STF Load

TABLE III

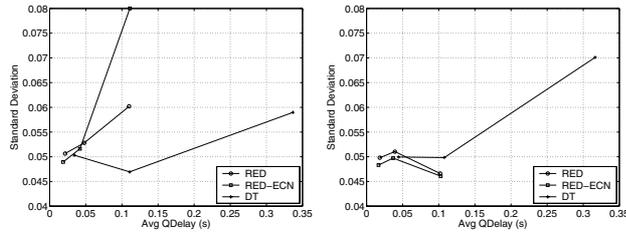
LINK UTILIZATION IN DIFF. CONFIGURATIONS UNDER 30% STF LOAD

Multiple of BWDP	5Mb Link			35Mb Link			100Mb Link		
	RED	RED-ECN	DT	RED	RED-ECN	DT	RED	RED-ECN	DT
1/3	.943	.947	.974	.961	.955	.968	.967	.959	.971
1	.963	.965	.975	.967	.967	.971	.971	.971	.972
3	.973	.973	.976	.969	.970	.972	.972	.972	.973

queueing delay. It is also noticed that using RED-ECN has a marginal gain of TCP throughput over RED, while RED-ECN requires ecn-compatible TCP sources.

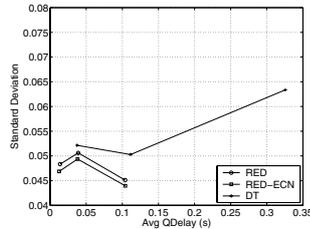
Link utilization under 30% STF load is listed in different configurations in Table III. Under low BWDP cases, link utilization of the droptail router is higher than that of either RED or RED-ECN. It is intuitive, since RED tends to drop packets earlier than DT. And the impact gets magnified when the buffer is smaller. So the utilization difference between DT and RED under 5Mb link is larger than those under 35Mb and 100Mb links. With the increase of BWDP, the utilization difference between DT and RED becomes smaller. RED-ECN gains marginally on link utilization compared to RED.

Besides the average queueing delay, standard deviation of queueing delay is also interesting, when jitter becomes an important factor for the performance of applications. In Figure 4, standard deviation under 30% STF load and 30% ON/OFF LTNRF load is shown with different configurations. Each ON/OFF long-term non-responsive flow was "ON" for 20s and "OFF" for the next 20s. By configured like this, the queue length fluctuated much higher than it did under constant non-responsive load. This can help us to clearly illustrate the differences among different buffer



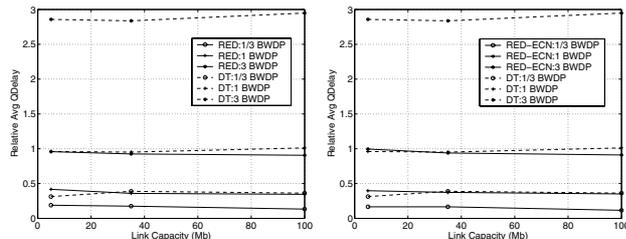
(a) Link Cap.=5Mb

(b) Link Cap.=35Mb



(c) Link Cap.=100Mb

Fig. 4. Std. Dev. of QDelay with Diff. Avg QDelays under 30% STF Load and ON/OFF LTNRF Load



(a) ECN Disabled

(b) ECN Enabled

Fig. 5. Relative Avg QDelay with Diff. Link Capacities under 30% STF Load

management schemes.

It is observed from Figure 4: 1) Under 5Mb link, DT has comparable standard deviations to RED and RED-ECN; 2) With the increase of the link capacity and buffer sizes, RED and RED-ECN have much smaller standard deviations (i.e. less queue length fluctuation) than DT. RED and RED-ECN seem more suitable for high link bandwidths and larger buffer sizes when delay jitter is an important consideration under dynamically changing workloads.

C. Changing link capacities

When compared across three plots of Figure 3, it is noticed that: 1) the increase of link capacities has minor impact on the differences of TCP throughputs among queue management schemes; 2) TCP throughputs are higher under 100Mb link than those under other link capacities.

In Figure 5, *relative* average queueing delays with different link capacities are compared under different configurations. *Relative* average queueing delay is defined as

$Avg\ QDelay / Round\ Trip\ Propagation\ Delay$.

It is noticed that, for a droptail router under different BWDPs, the relative average queueing delay tends to be close to the buffer size. For example, DT under 3 BWDP has the average queueing delay around 3 times of 1 round trip propagation delay. Similar tendencies can be observed in other DT cases. The reason is that, in a droptail router, with high link utilizations, the queue tends to stay fully occupied most of the time. So queueing delay of the droptail router is close to $Buffer\ Size / Serving\ Rate$. RED, however, has noticeable lower average queueing delays. For example, under 3 BWDP, the relative average queueing delay of RED is around 1, while that of DT is around 3. The average queueing delay is about 3 times smaller with RED than with DT under the same multiple of BWDP. It is also noticed that RED-ECN has similar average queueing delays to RED. Changing link capacities has almost no impact on the trends of average queueing delays.

TABLE IV

DROP/MARKING RATES UNDER 30% STF LOAD AND 1 BWDP

QM	Type of Flow	Link Capacity (Mb)		
		5	35	100
RED	LTRF ¹	.03627	.03112	.02503
	LTNRF	.03681	.03891	.02814
RED-ECN	LTRF ²	.00352/.04256	0/.04123	0/.03036
	LTNRF	.04688	.05352	.03406
DT	LTRF ¹	.01787	.01992	.01662
	LTNRF	.10229	.09954	.12189

¹ Format: Drop Rate

² Format: Drop Rate/Marking Rate

Table IV compares the drop rate or marking rate of long-term responsive and non-responsive flows from RED, RED-ECN and DT under 30% STF load and 1 BWDP with different link capacities. It is observed that drop rates of LTNRFs are much higher than those of the other flows in a droptail router, while the drop rates of RED router and the marking/drop rates of RED-ECN router show no significant difference across different types of flows. RED-ECN has slightly higher *marking rates* of LTRFs than *drop rates* of LTRFs in RED, but with much smaller actual drop rates of LTRFs. The drop rates of LTNRFs in RED-ECN are higher than those in RED, so RED-ECN gains marginally long-term TCP throughput improvement over RED (see Fig. 3).

D. Changing STF loads

In Figure 6, *normalized* long-term TCP throughputs with different STF loads are compared under 1 BWDP of each link capacity. The normalized TCP throughput is defined as $Total\ TCP\ Throughput / (Total\ TCP\ Throughput + Total\ UDP\ Throughput)$ so that TCP throughput is always compared to the amount of throughput from both TCP and UDP flows regardless of added STF loads. It is observed that with the increase of STF load, the throughput difference decreases or stays the same between DT and AQM schemes (RED or RED-ECN).

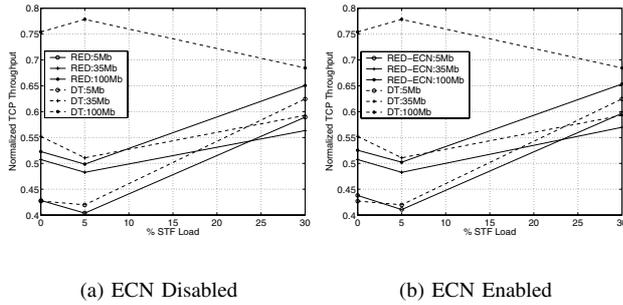


Fig. 6. Normalized TCP Throughput with Diff. STF Loads under 1 BWDP

TABLE V
COMPARISON OF THROUGHPUTS WITH DIFF. STF LOADS UNDER 1 BWDP AND 100MB LINK

STF Load	RED			RED-ECN			DT		
	LTRF	STF	LTRNF	LTRF	STF	LTRNF	LTRF	STF	LTRNF
0%	.505	0	.461	.507	0	.458	.730	0	.238
5%	.457	.051	.460	.460	.051	.456	.729	.051	.190
30%	.454	.272	.244	.457	.271	.242	.478	.272	.220

Because of the similar TCP throughputs between RED and RED-ECN in Figure 6, the throughputs of different classes of traffic in different configurations is listed in Table V again. Data in the table were collected with different STF loads and the buffer size of 1 BWDP at 100Mb link. It is noticed that STF load is almost constant across 3 buffer management schemes and that long-term TCP (LTRF) throughput under DT is noticeably higher than REDs when STF load is low (0% and 5%). Short-term flows claimed the proportion of the link capacity (5% or 30%) with almost no impact from long-term flows, since the aggregated behavior of short-term flows at the router are more aggressive than long-term flows. RED-ECN has marginal improvement in TCP throughput compared to RED. Under higher STF load (30% in our experiment), TCP throughputs from RED and RED-ECN become very close to that of DT.

III. DISCUSSION

In this section, we summarize and discuss simulation results within the scope of our investigation:

- 1) With the existence of STFs, the investigated performance metrics of both RED and RED-ECN with the recommended configuration is comparable to or prevails over that of DT under higher BWDP cases. Change of link capacity or STF load has minor impact on trends of long-term TCP throughputs and relative average queueing delays. With the increase of STF load, TCP throughputs of AQM schemes become very close to that of a droptail router.
- 2) RED-ECN has marginal long-term TCP throughput improvement compared to RED, but provides lower drop rate. It was shown earlier that, with ECN mechanism enabled, TCP throughput will be benefit significantly in a highly congested network [11]. The congestion is moderate

(drop rates are between 1% and 10%) in our simulations. Secondly, although our simulations employ recommended TCP-Sack(with or without ECN), the performance differences between RED and RED-ECN are still very small under the workloads in our simulations. However, it is possible that RED-ECN would perform better with different workload configurations. So advantages of ECN mechanism are less significant within the scope of our investigation.

3) Between DT and AQM schemes(RED, RED-ECN), for lower BWDPs or smaller buffers ($\ll 1$ BWDP), droptail routers provide better realized long-term TCP throughput, but at slightly higher average queueing delays; for high BWDPs or larger buffers (≥ 1 BWDP), such as those long-distance high-capacity links, AQM schemes have better performance in terms of acceptable realized long-term TCP throughput and significantly lower average queueing delay and jitter.

IV. CONCLUSION

In this paper, the impact of bandwidth-delay product and non-responsive flows on the performance of both droptail and AQM schemes was investigated. Simulations under combinations of different link capacities, buffer sizes and loads of short-term flows were conducted. Observations and discussions of simulation results were provided. From the experiment results, we observe that, with the existence of short-term flows, AQM schemes(RED or RED-ECN) have comparable long-term TCP throughput to a droptail router and have the advantage of lower queueing delays and jitters, over links with high bandwidth-delay products. Droptail routers provide better long-term TCP throughput with higher delays over links with smaller bandwidth-delay products.

Acknowledgements: This work was supported by an NSF Grant ANI-0087372, the State of Texas, TITF, and Intel Corp.

REFERENCES

- [1] National Laboratory for Applied Network Research, "Network traffic packet header traces," 2000, <http://moat.nlar.net/Traces/>.
- [2] End to-end discussion group, "Web mice versus elephants," *Mailing list discussion*, 2000-2001.
- [3] N. Caldwell, S. Savage, and T. Anderson, "Modeling TCP latency," *Proc. of IEEE Infocom*, March. 2000.
- [4] S. Floyd and E. Kohler, "Internet research needs better models," *Proc. of HotNets-I*, Oct. 2002.
- [5] G. Iannaccone, M. May, and C. Diot, "Aggregate Traffic Performance with Active Queue Management and Drop from Tail," *ACM SIGCOMM Computer Communication Review*, 2001.
- [6] S. Floyd and K. Fall, "Promoting the use of end-to-end congestion control in the internet," *ACM Trans. on Networking*, Aug. 1999.
- [7] M. May, J. Bolot, C. Diot, and B. Lyles, "Reasons not to deploy RED," *Proc. of IWQoS'99, London*, pp. 260-262, Jun. 1999.
- [8] M. Christiansen, K. Jeffay, D. Ott, and F. D. Smith, "Tuning RED for web traffic," *Proc. of SIGCOMM'00*, Aug. 2000.
- [9] University of California at Berkeley, "Network simulator (ns)," 1997, Available via <http://www.nrg.ee.lbl.gov/ns/>.
- [10] S. Floyd, "RED: Discussion of Setting Parameters," Available via <http://www.icir.org/floyd/REDparameters.txt>.
- [11] S. Floyd, "TCP and explicit congestion notification," *ACM SIGCOMM Computer Communication Review*, Oct. 1994.