Effectiveness of ECN and RED Based Congestion Detection

Yury Puzis

Vasily Tarasov

Stony Brook University

12/15/2008

Abstract

This project evaluates performance of of TCP RENO under Random Early Detection (RED) and Explicit Congestion Notification (ECN) -based schemes of congestion avoidance. The evaluation is done in a number of Wide Area Network (WAN) scenarios, which are analyzed for throughput, drop-rate and fairness.

1 Introduction

Standard behavior of TCP protocol dictates that in order for source node to find out that the network is congested is to detect the congestion by detecting packet drops, either from triple-ack sent by the recipient, or from a packet timeout. This approach has several drawbacks.

First, if packets are being dropped it means that the network is already congested, and other sources might also be experiencing packet drops (extent of impact on other sources depends on network topology). As a consequence, a significant number of packets might have to be retransmitted on the network, and multiple sources might throttle the congestion window synchronously, leading to immediate significant under-utilization of the network. If network congestion could be detected preemptively, and reacted to with better coordination between the traffic sources, those problems would be minimized.

Second, source node is only listening to it's own traffic, and therefore cannot make well-informed decisions about conditions of the network as a whole. The difference in impression among traffic sources about conditions of the network before it finally becomes congested for the majority of nodes, leads for further lack of coordination inability to preemptively prevent network congestion.

RED gateways [1] are designed to help source nodes handle network congestion. Since gateways are able to monitor traffic from multiple sources to multiple destinations, they are much better positioned then individual nodes to make conclusions about state of the network as a whole. RED gateways detect possible network congestion by monitoring their output queue, and provide feedback to transport layer protocol of the traffic source. While participation of transport layer in this scheme is not required, well-designed response in the transport layer can be very beneficial to performance of the network. In this report we investigate impact of ECN scheme described in [2], which is designed specifically for TCP RENO style protocol.

Section 2 of the project describes the explicit congestion notification scheme used in our experiments. Section 3 describes the experimental setup, presents results and analysis of the simulations. Section 4 presents the conclusions.

2 Description of Explicit Congestion Notification Scheme

Since ECN scheme uses RED gateways for explicit congestion notification, we describe the RED scheme first. RED gateway maintains two queue thresholds, *min*, and *max,* and continuously updates average queue size. When number of packets is less than *min*, no action is taken. When number of packets is greater than max, every arriving packet is marked or dropped. When number of packets is between min and max, RED gateway calculates probability that the packet should be dropped or marked, proportional to the connection's share of the gateway's bandwidth. In addition, if gateway measures the queue size in bytes rather than packets, the probability packet is to be marked is also proportional to the packet size in bytes. Finally, if gateway's output queue is empty for some period of time, the algorithm estimates how many packets could have been transmitted during that time, and updates the average queue size accordingly.

The goal of ECN scheme is to define response of the transport layer of the sender node to reciet of congestion notification from RED router. When ECN scheme is active, RED gateway is configured to mark rather than drop packets. On one side, the ECN scheme used in this report uses a single message as an indication of network congestion. On the other side the scheme tries to make sure TCP does not respond too frequently by reacting to congestion notification at most once per round trip time (this includes triple-acks). Following reciet of packet with ECN bit set, the sender will halve the congestion window and the slow-start threshold. The protocol does not halves those parameters again in response to triple-ack or another packet with ECN bit, until all packets outstanding at time of response to ECN have been acked.

3 Experiments and Analysis

In all experiments we study behavior of WAN with two gateways, and several nodes attached to each gateway. The link between gateways is configured to either simply drop packets when the queue is full (DropTail in graph), or behave as a RED gateway (described above, RED in graph). The nodes are configured to work as an unmodified TCP Reno (DropTail and RED), or TCP Reno with ECN activated (ECN in graph). Nodes on one side of the network send FTP traffic to nodes on the other side of the network. Bandwidth between the nodes and the gateways is set to 100Mbps, and bandwidth between the gateways is set to 1Mbps, designed to create bottleneck effect for traffic.

In the first experiment (Figures A, B) we measure additive and average throughput of the network as a function of time. Settings for the RED gateway are 100 packets queue size, 5 packets min threshold, and 15 packets max threshold. Here one node is connected to each of the two gateways in the network.

As can be observed, the network becomes congested after about 3rd sec. Observe the difference between DropTail and RED gateways. DropTail exhibits higher throughput, but uses it's queue completely. This affects latency of the network. RED on the other hand provides lower throughput but keeps the queue under-utilized (Figures C, D) that would allow other possible connections to experience lower latency. ECN gives same benefits as RED in terms of latency, but throughput is higher.

Pay attention that the average throughput in RED network drops twice (Figure A), as a consequence of the node throttling it's congestion window two times.

Also notice that since ECN scheme provides an explicit congestion notification to the sender, it's congestion window is reduced faster, less packets are dropped, and the network restores to better bandwidth utilization faster. This is why we can observe difference between RED and ECN graphs on Figure C.

On Figure D we can clearly see that after initial significant reduction is number of packets in queue (as a consequence of earlier congestion), queue starts to oscillate at around 3000 - 4000 bytes, or 5 – 7 packets. The reason is that minimum threshold of the queue is only 5 packets, after which packet drops / ECN flags cause repeated congestion window drop and restoration, which results in buffer size oscillation.

Figure A. DropTail, RED, ECN: Average throughput behaviour in a congested WAN

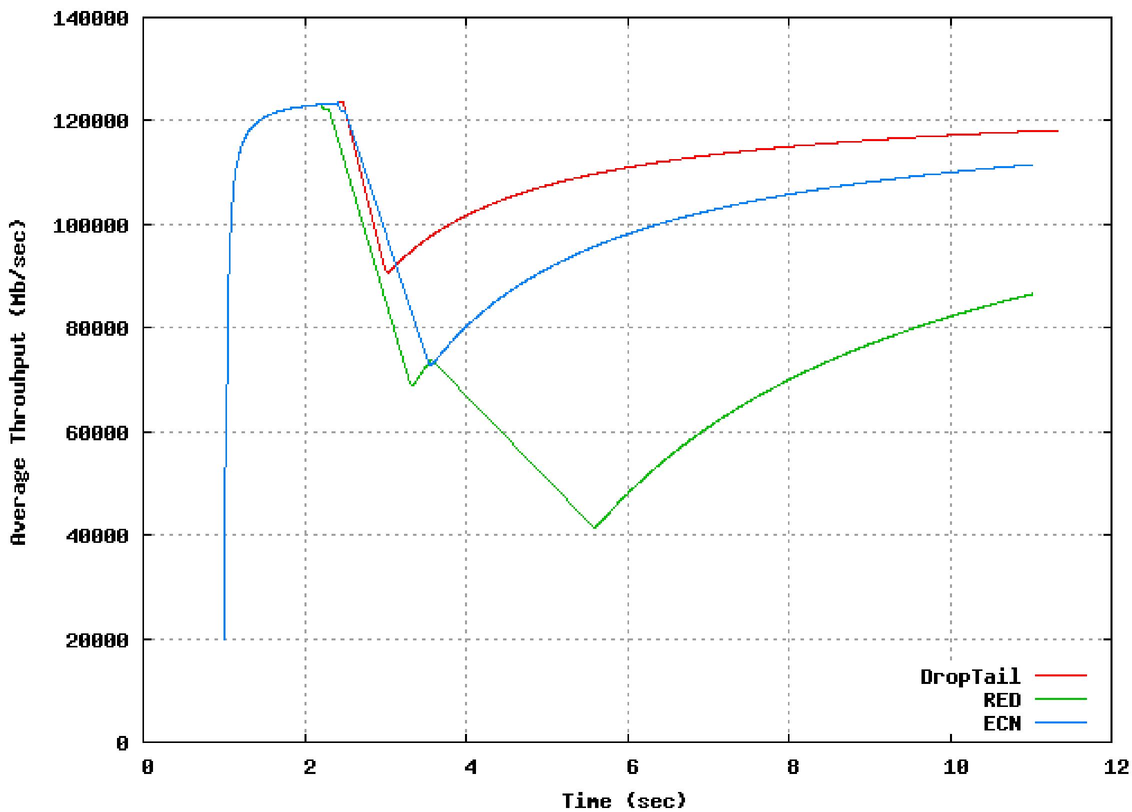

Figure B. DropTail, RED, ECN: Additive throughput behaviour in a congested WAN

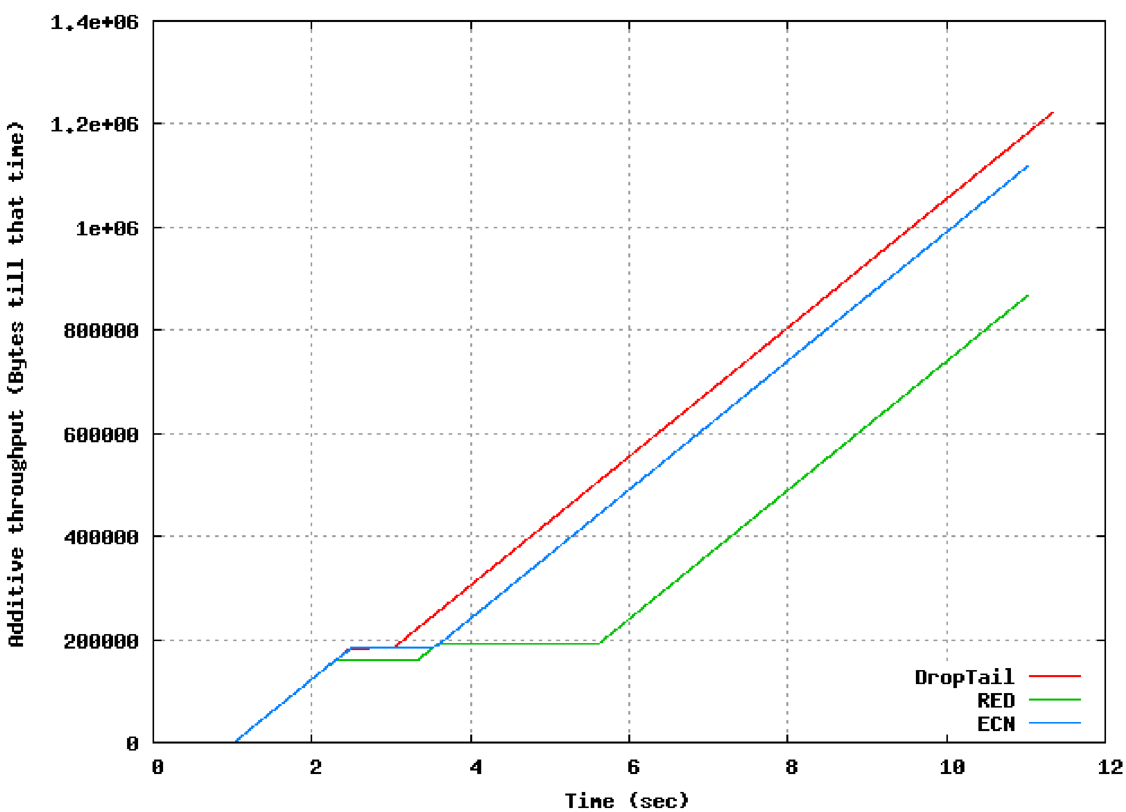

Figure C. DropTail, RED, ECN: Average Buffer Usage in a congested WAN

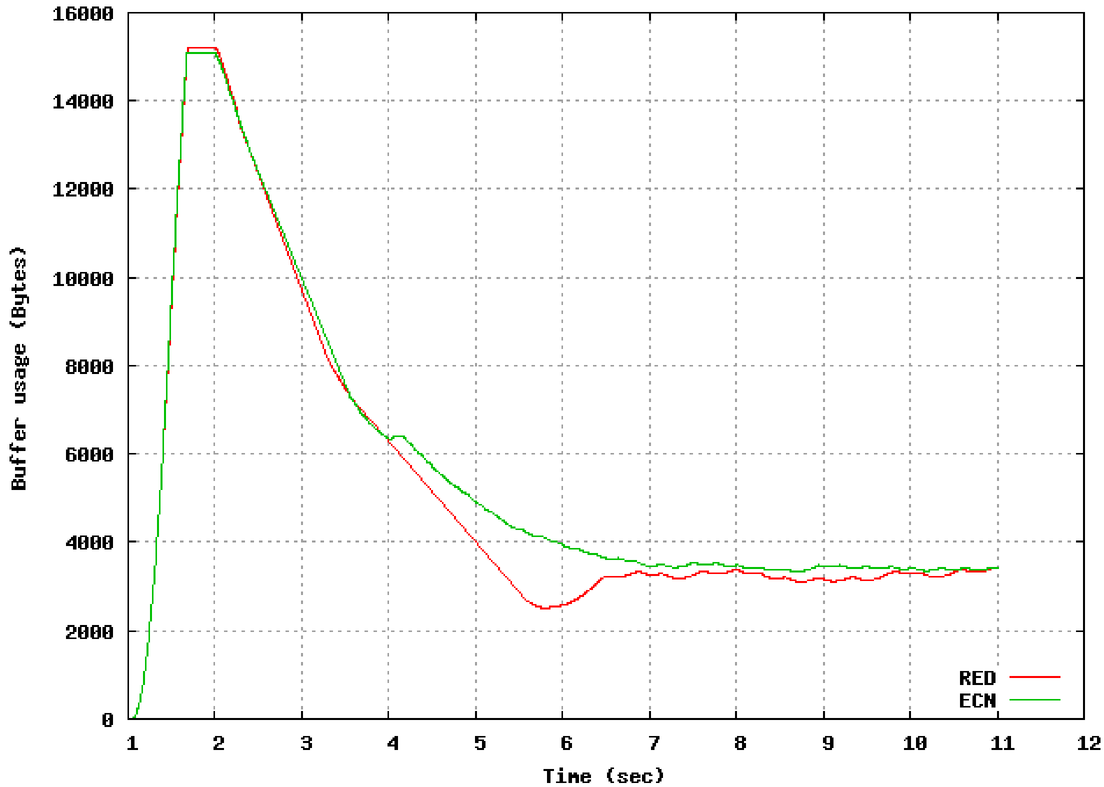

Figure D. DropTail, RED, ECN: Current Buffer Usage in a congested WAN

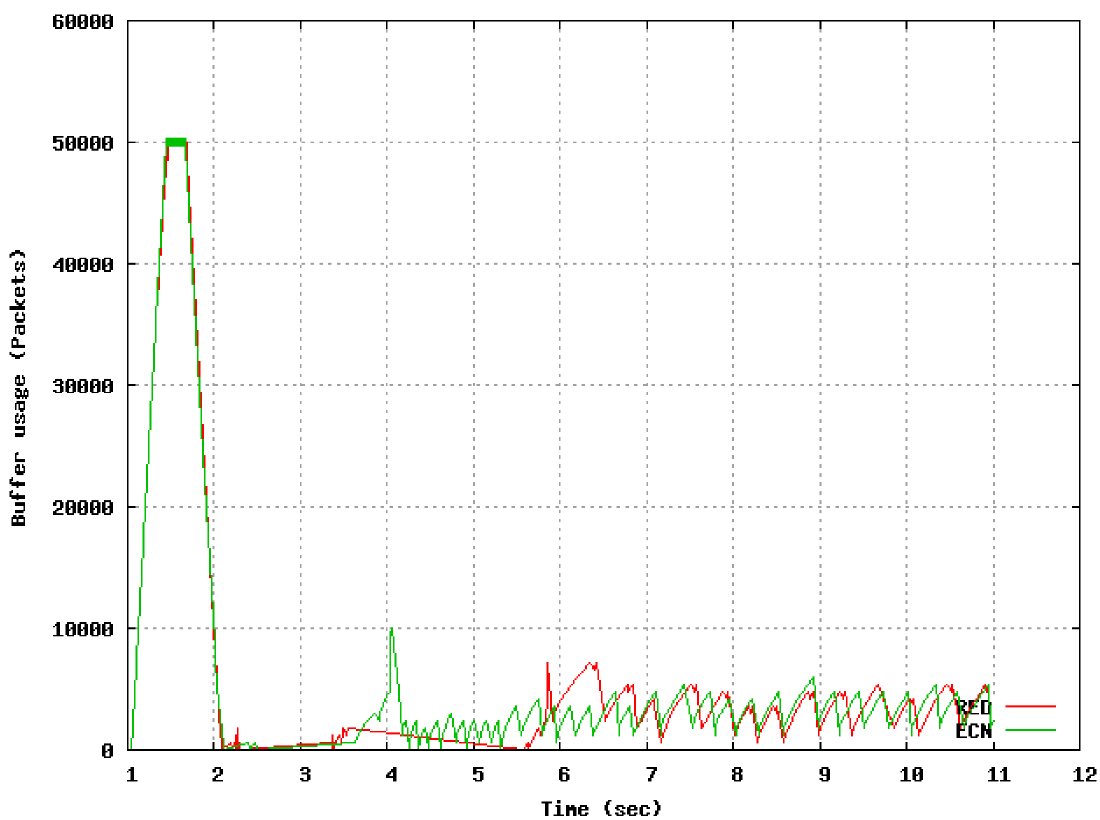

In the next experiment (Figures I,J,K,L) we modify the RED gateway thresholds, setting minimum to 40 and maximum to 70 packets. The first thing to notice here is that throughput of RED and ECN schemes are nearly identical to throughput of DropTail scheme. This is a direct consequence if higher thresholds of the RED gateway. The result of this effect is evident in Figures I and J.

Second observation is that buffer usage of both RED and ECN schemes is much higher than in previous experiment. After initial peak at nearly full buffer capacity the buffer size drops and oscillates between 30000 bytes (about 55 packets) and 0 bytes in case of RED, or 5000 bytes (9 packets) in case of ECN. The average is near the minimum threshold of the gateway, or 40 packets. The scenario is similar to the one described above for figure D.

Notice the advantage of ECN revealed in this experiment. After reaching the region between minimum and maximum thresholds (when the protocol starts to drop packets) the buffer usage in case of ECN protocol drops down less dramatically than in case of RED protocol. This is due to fact that explicit notification provided by the ECN protocol allows better reaction on congestion than just double reduction of congestion window.

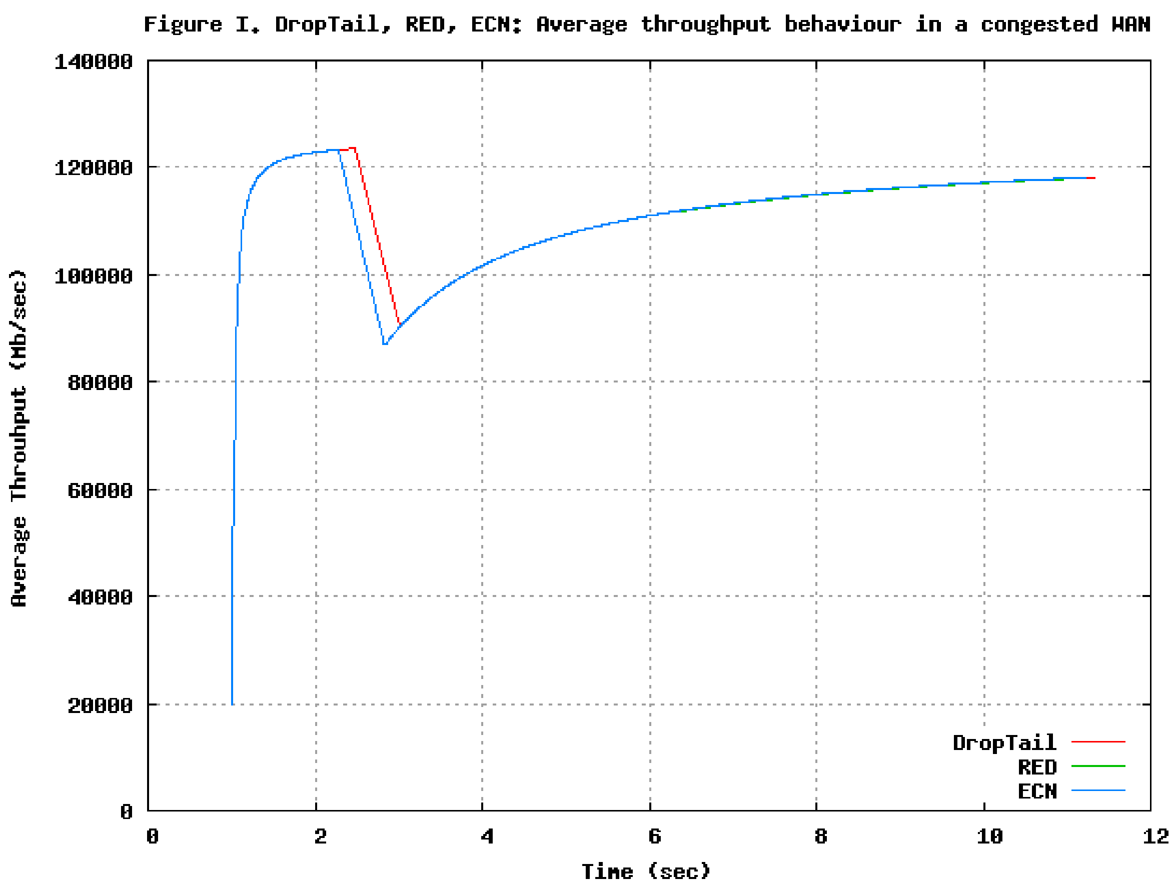

Figure J. DropTail, RED, ECN: Additive throughput behaviour in a congested WAN

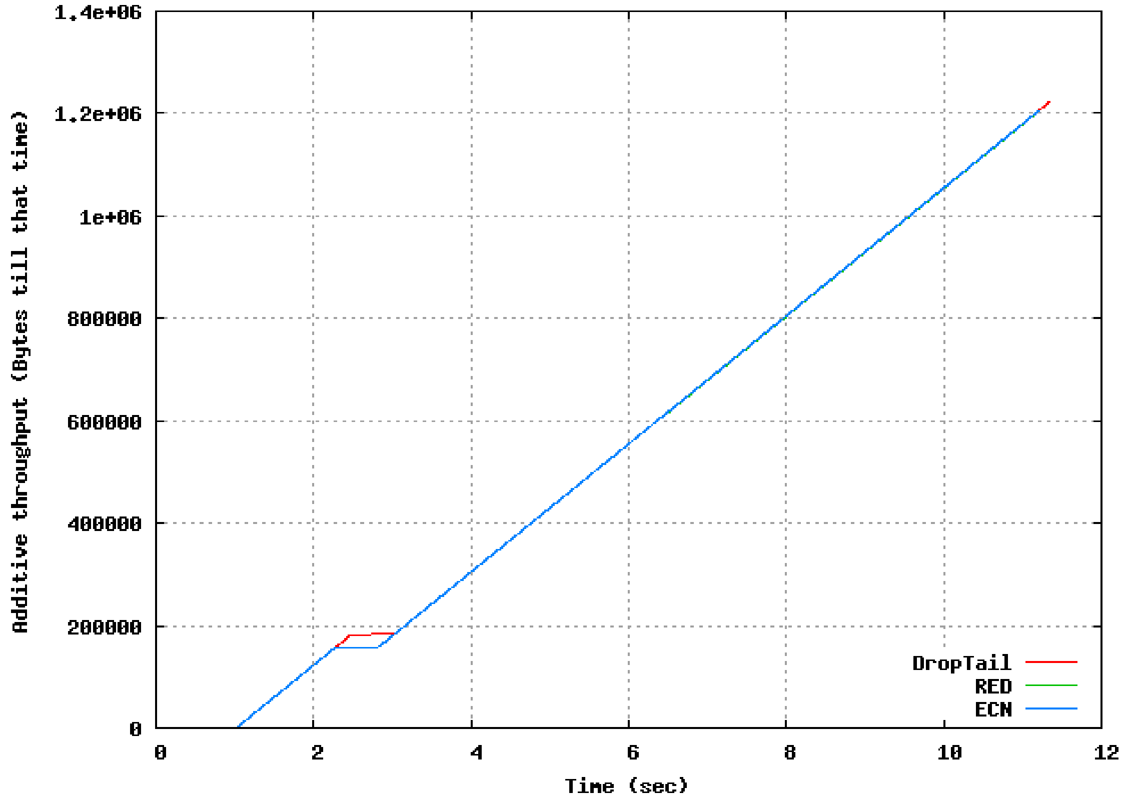

Figure K. DropTail, RED, ECN: Average Buffer Usage in a congested WAN

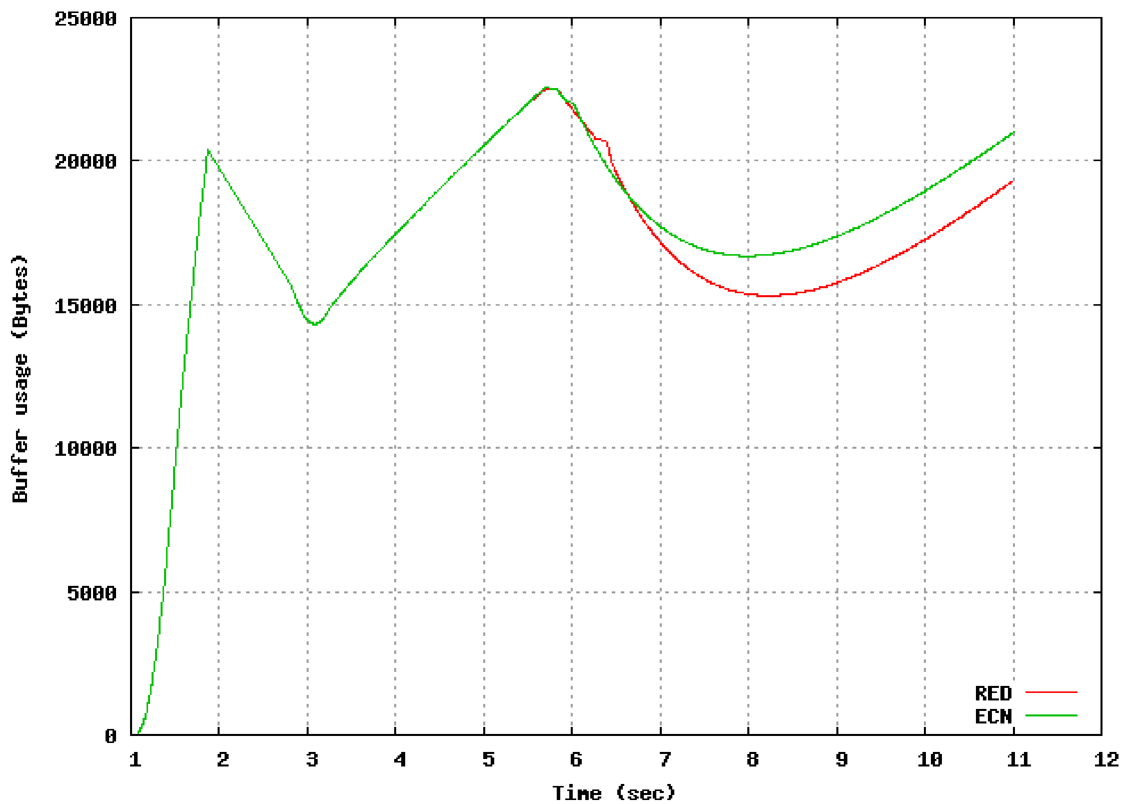

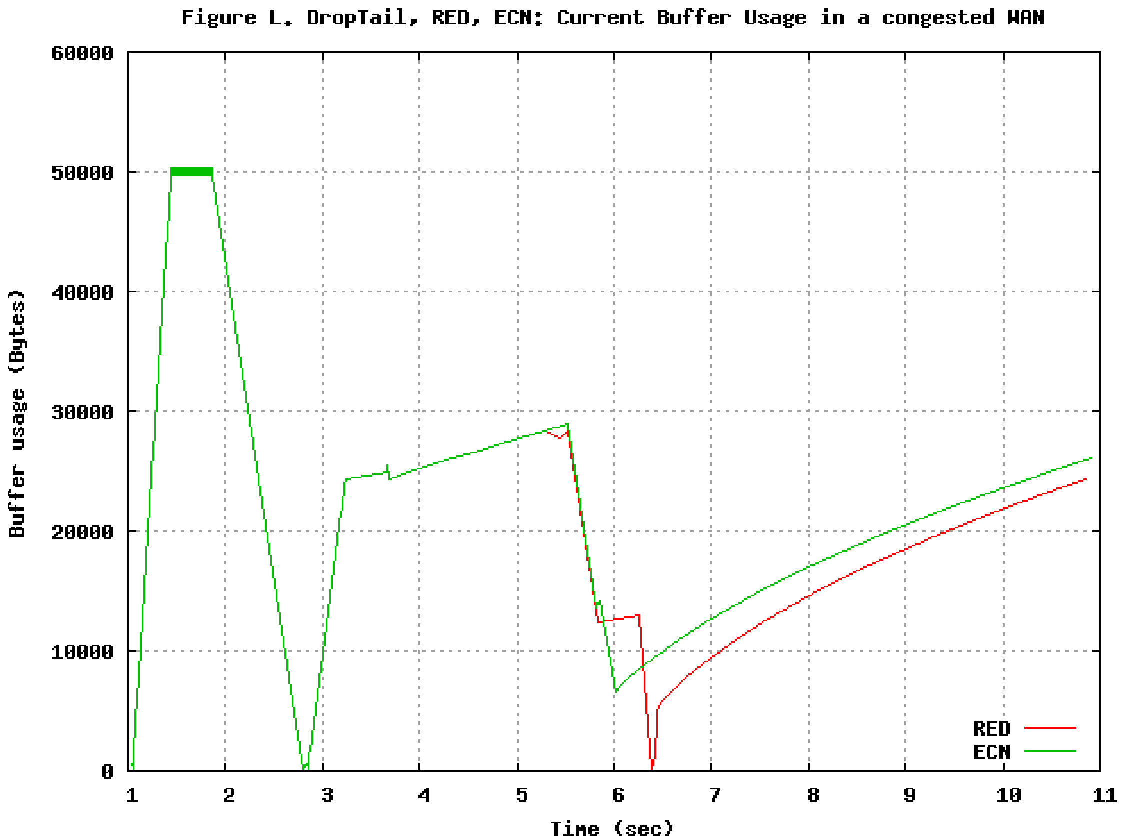

In the final experiment (Figures E,F,G,H) we keep the modified threshold sizes, but increase the number of nodes connected to each gateway from 1 to 4.

Lets start from Figures G and H, compared to figures K and L respectively. Observe that the shapes of the plots are similar, but in case of multiple nodes the curve is shifted up, and is more smooth. The reason is that while some nodes throttle their congestion window, other nodes continue to supply traffic to the gateway.

Notice that the more nodes are in the network, the less noticeable is the difference between ECN, RED and DropTail schemes in terms of throughput and buffer usage. However [1] demonstrated that DropTail has a bias towards burst traffic in such cases, while ECN and RED are fair to all the flows going through the gateway. This can be easily noticed by conducting the experiment when at some point of time during the experiment an additional source of bursty traffic is added to the gateway. Then the share of total throughput given to some connection will suddenly decrease.

Figure E. DropTail, RED, ECN: Average throughput behaviour in a congested WAN

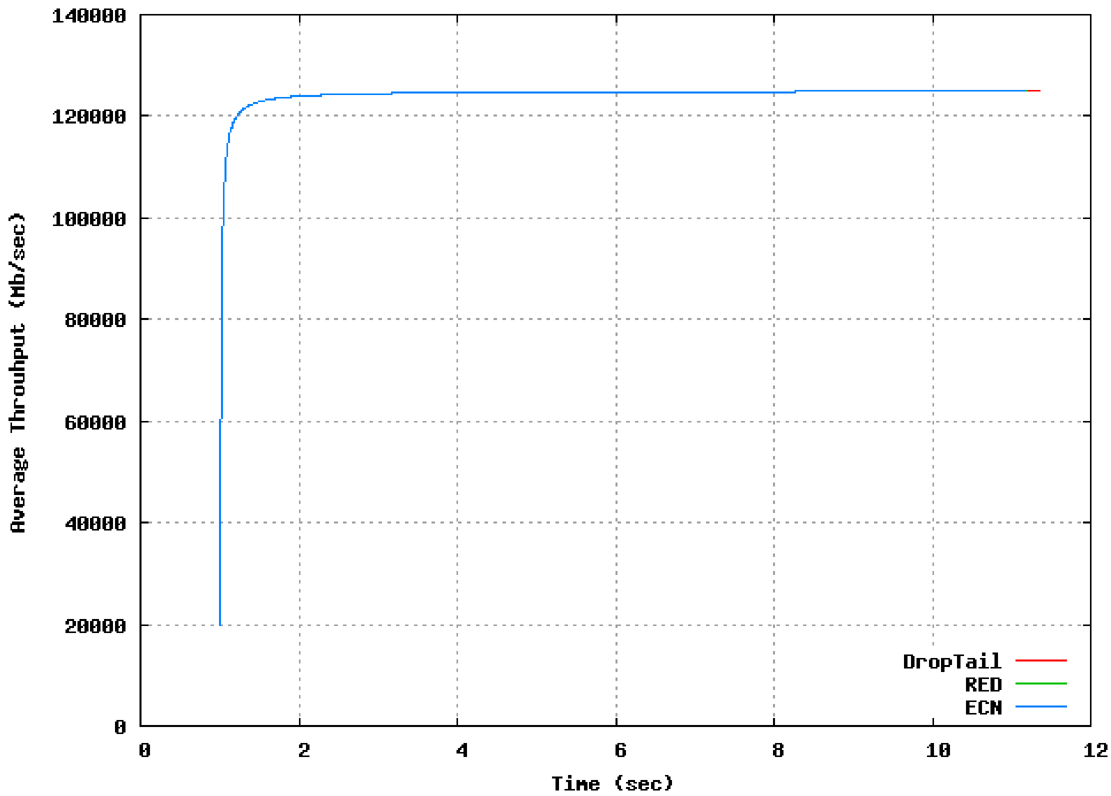

Figure F. DropTail, RED, ECN: Additive throughput behaviour in a congested WAN

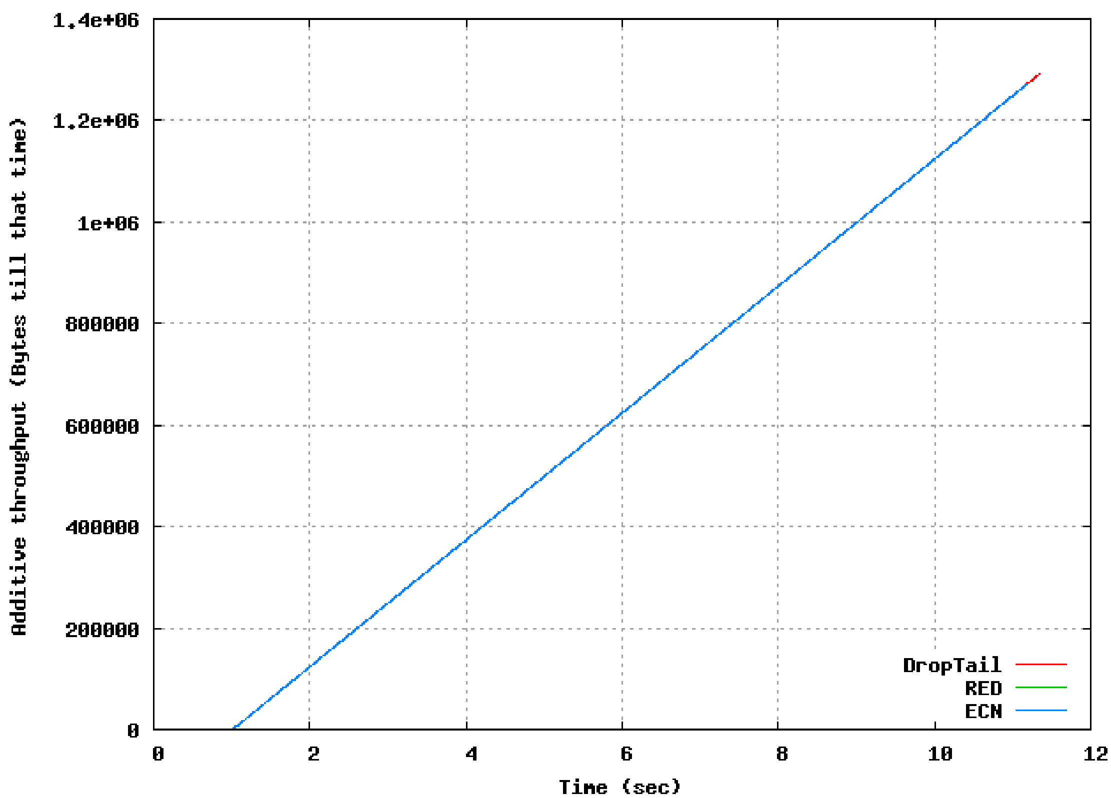

g

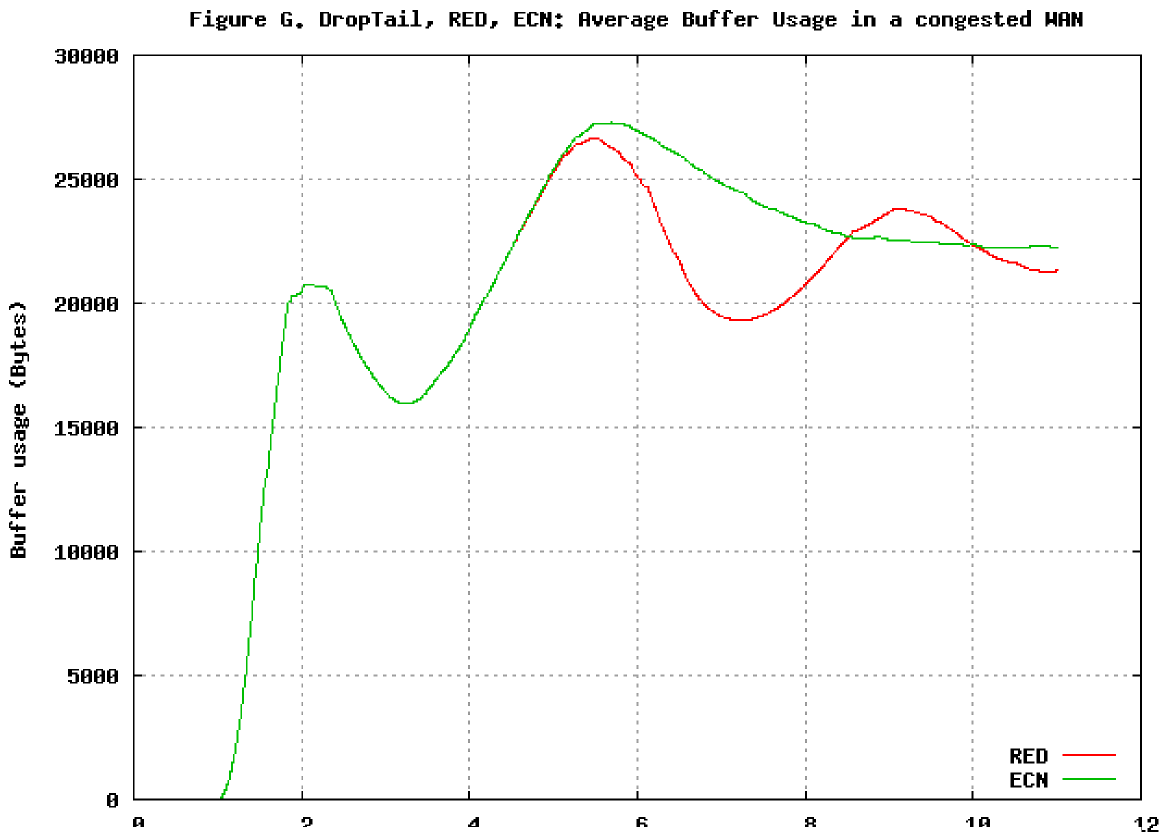
Figure G. DropTail, RED, ECN: Average Buffer Usage in a congested WAN
Buffer usage (Bytes)
30000
25000
20000
15000
10000
5000
0
0
2
4
6
8
10
12
RED
ECN

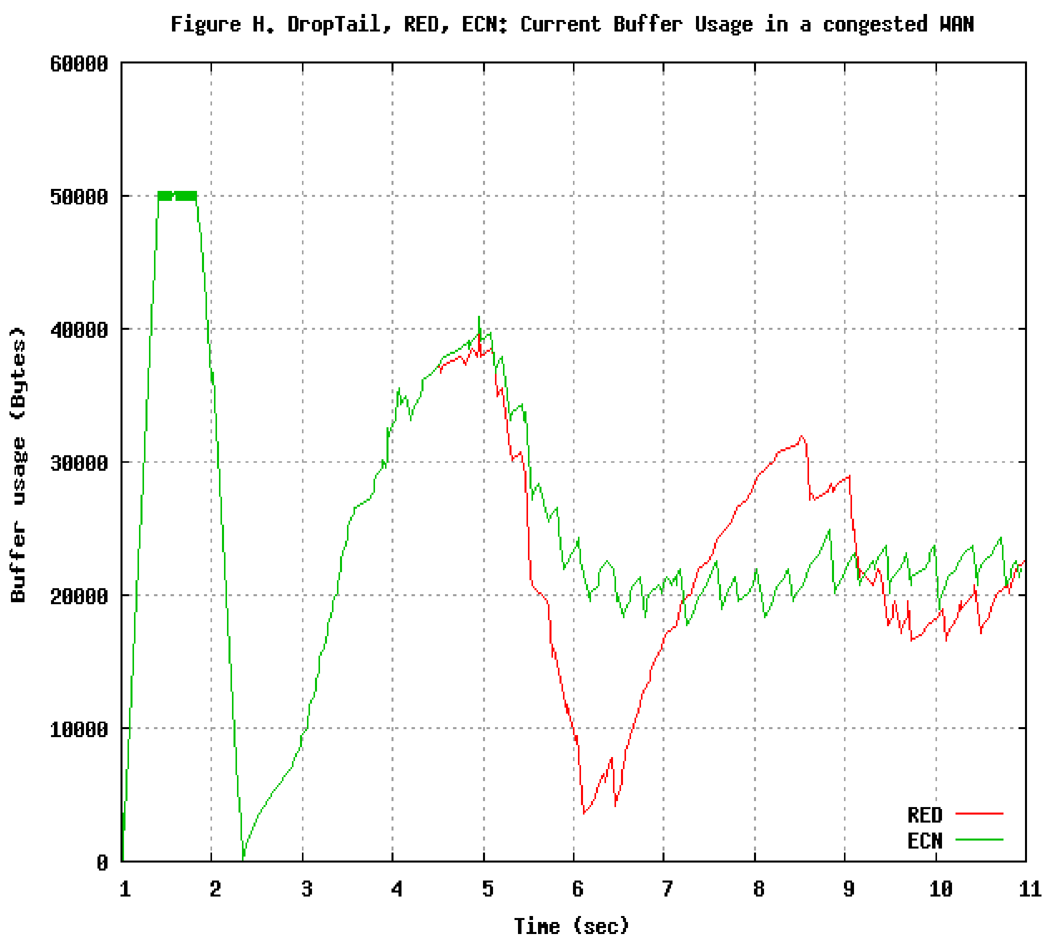
Figure H. DropTail, RED, ECN: Current Buffer Usage in a congested WAN
Buffer usage (Bytes)
60000
50000
40000
30000
20000
10000
0
1
2
3
4
5
6
7
8
9
10
11
Time (sec)
RED
ECN

4 Conclusion

As a result of experiments described in this report we were able to confirm several advantages of RED and ECN over regular TCP Reno with DropTail routers. Namely the throughput is close to the level of DropTail (especially in case of multiple nodes per gateway) but the buffers are used more efficiently, resulting in better latency and fairness. The value of this report is that unlike similar works this area we simulated WAN environment.

5 References

[1] Floyd, S., and Jacobson, V., Random Early Detection gateways for Congestion Avoidance, IEEE/ACM Transactions on Networking, V.1 N.4, August 1993, p. 397-413

[2] Floyd, S., TCP and explicit congestion notification, SIGCOMM Comput. Commun. Rev., V.24 N.5, 1994, p. 8-23