

Constant Congestion Window approach for TCP – effect on Fairness

Anup K Ghosh

IBM Global Services India, Salt Lake, Calcutta 700 091, India

anup.ghosh@in.ibm.com

Sudipto Das

Dept of CSE, Jadavpur University, Calcutta 700 032, India

Rajesh Roy

Dept of CSE, Jadavpur University, Calcutta 700 032, India

Amitava Mukherjee

Royal Institute of Technology, School of Electrical Engineering, Stockholm 10044, Sweden

amitava.mukherjee@in.ibm.com or amitava@imit.kth.se

(Author for Correspondence)

Abstract

This is an extension of the work done by Ghosh et al [1] whereby it was indicated that a fixed congestion window performs better than a TCP Reno implementation. This paper investigates the impact of this model on the fairness of connections and comes up with suggestion to improve this.

Keywords: congestion, congestion window (cwnd), wireless losses, mixed wireless network

1. Introduction

TCP [3], [4] was designed to work well in networks with low channel error rates. Wireless networks on the other hand are characterized by frequent transmission losses. As a result, when TCP is used in wired/wireless internetworks, the losses due to channel errors are mistaken as congestion losses and the sending rate is unnecessarily reduced in an attempt to relieve the congestion, resulting in a degraded performance [2]. Various approaches to address this problem had been discussed and compared in the work by Balakrishnan et al. [9]-[11]. Three alternative approaches – (i) link layer enhancements (ii) network layer enhancements and (iii) end-to-end (E2E) solutions – were carefully contrasted, of which the E2E solutions are the most versatile and efficient.

In [1] Ghosh et al proposed an E2E solution by making TCP Congestion Window (henceforth referred to as *cwnd* and expressed in number of segments) a constant. In such a situation, the congestion window for a sender is set to a value which is optimal for a given network scenario. The results of the simulations done by Ghosh et al [1] with their new model of TCP show that for a given network load condition, as

determined by the number of source-receiver pairs, a fixed window TCP performs better than variable window TCP under large and small wireless errors while performing equally well under no wireless loss condition.

Since a constant *cwnd* TCP is more aggressive in sending packets unlike Reno, there is a high probability of some connections grabbing a major share of the network resources while others are starved. In this paper we investigate this fairness issue and propose some improvements to the situation.

The rest of the paper is organized as follows: section 2 provides a complete summary of the simulation results comparing the performance of the modified TCP Sender and TCP Reno in various simulation scenarios and section 3 concludes the paper.

2. TCP performance and fairness

In this section, we discuss on the basic performance behavior of the modified TCP senders and its fairness among a number of connections sharing a bottleneck link. A performance comparison is made with the TCP Reno sources in similar network scenarios. Intermediate node buffer capacity is always set equal to the bandwidth delay product for the bottleneck link based on literature studied. Increasing the buffer capacity further does not have any impact on the performance; in fact it may have a negative effect as it may delay some packets by such an extent that the TCP senders time out and duplication of data within the network further congest it. The traffic model used is FTP with infinite data to send so that the sender has data to send whenever the network permits, and the packet size is set to 1000 bytes (1040 bytes with headers) in all experiments. The wireless

subnet is error prone, susceptible to constant error rates. In our simulations we have used the conventional TCP Sink which responds with an ACK for every packet received. There is no congestion or error in the ACK path. All simulations have been carried out for a period of 450 seconds with the TCP senders transmitting data for the entire period of simulation. All the simulations have been carried out with 802.11 MAC with a maximum available bandwidth of 1Mbps. A Two Ray Ground propagation model is used with an Omni-directional antenna.

In this section the performance of the Modified TCP Senders and TCP Reno has been compared based on the Goodput metric, i.e. the number of useful data packets received. It is obtained by subtracting the number of retransmitted packets from the number of packets transmitted.

All simulations in this paper have been carried out using the LBL network simulator ns2 [6], [7] with appropriate modifications for implementation of the changes in the modified TCP sender.

2.1. Performance evaluation in wireless network

In this section, we evaluate the performance of the modified TCP senders over the wireless networks. A performance comparison is made with existing TCP Reno working in similar network scenarios. This scenario has been studied to evaluate the performance of the modified TCP in networks where the only reason for packet loss is due to errors in the wireless subnet.

2.1.1. Simulation Set up




Figure 1: Simulation set up and configuration parameters for wireless scenarios

Figure 1 provides a schematic diagram of the simulation set up used for the wireless scenario. In our simulations, we have TCP senders sending data over peer to peer wired links linking the senders to a base station. The base station relays traffic to the corresponding wireless hosts that acts as the receiver.

2.1.2. Optimal Window determination

The idea proposed in this paper is the constancy of the TCP congestion window. Hence, for maximized performance of the sender in a given network scenario, the determination of an optimal cwnd becomes imperative.

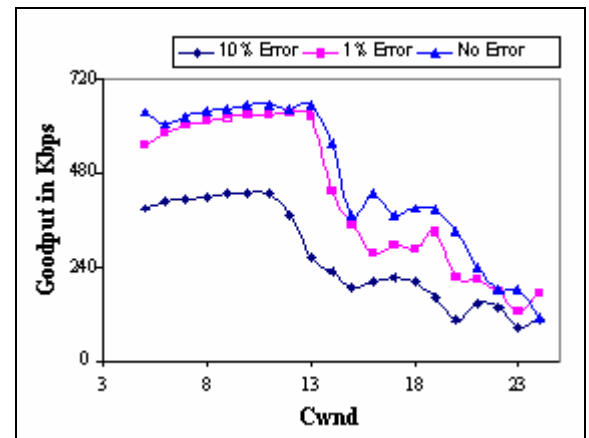


Figure 2: Goodput vs. cwnd curve for wireless scenarios for single source receiver pair at various error rates

We observed in figure 2 that there exists certain value of cwnd for which the system performance is a maximum. Since the peak is somewhat flattened, the performance is maximum, or very near to the maximum for a number of cwnd values, we can use any value from this to set the cwnd.

2.1.3. TCP performance evaluation in terms of goodput

In this subsection, we evaluate the performance of the modified senders, in terms of goodput, in comparison to the performance of the TCP Reno sources operating in similar network scenario.

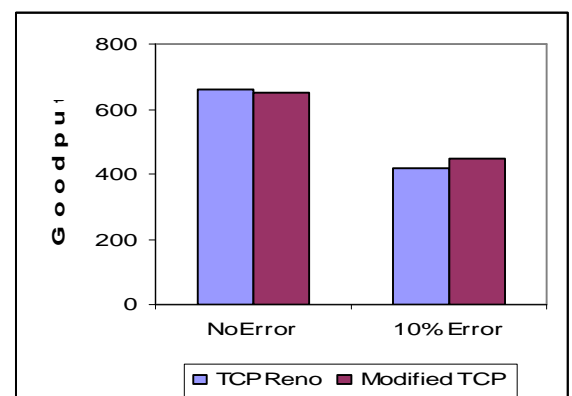


Figure 3: Performance comparison of TCP Reno and Modified TCP under various error rates

Clearly, from the figure 3, the modified TCP sender performs at par with to the TCP Reno senders in the case of no wireless errors consider, and outperforms it with the case of wireless errors in 10% constant bit error in the wireless links. The 10% improvement in goodput has been observed.

2.2. Performance in mixed network

In this section, we evaluate the performance of the modified TCP senders over the mixed (wired cum wireless) network. A performance comparison is made with the existing TCP Reno running in the similar scenario. Having established, in the previous subsection 2.1, that the Constant cwnd TCP Senders perform better in the situations where wireless loss is the only source of packet loss, this scenario is used to evaluate the performance of TCP Senders in the networks where packet losses can be caused not only by wireless losses but also congestion in the network as well.

2.2.1. Simulation Set up

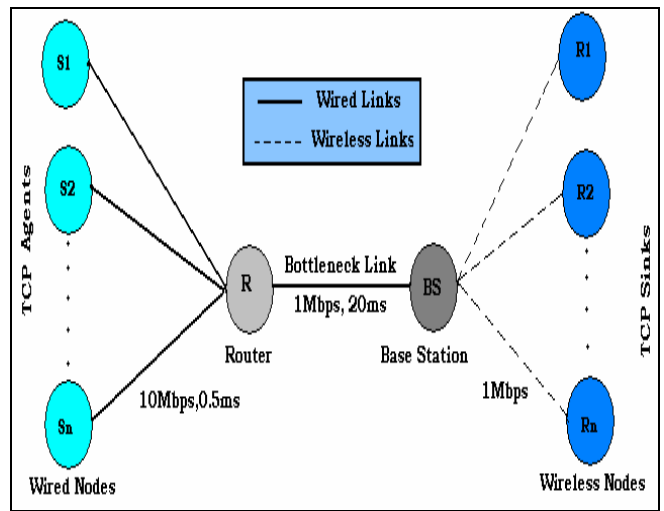


Figure 4: Simulation set up and configuration parameters for mixed network

Figure 4 provides a schematic window of the simulation set up used. All the sources are connected to a router sending data over a wired link that proves to be the bottleneck link in the network and connects the wired subnet to the base station. The base station relays traffic to the corresponding wireless hosts that act as the receivers. The objective behind selecting such a scenario is the evaluation of the performance of the modified TCP sources in congested wired network.

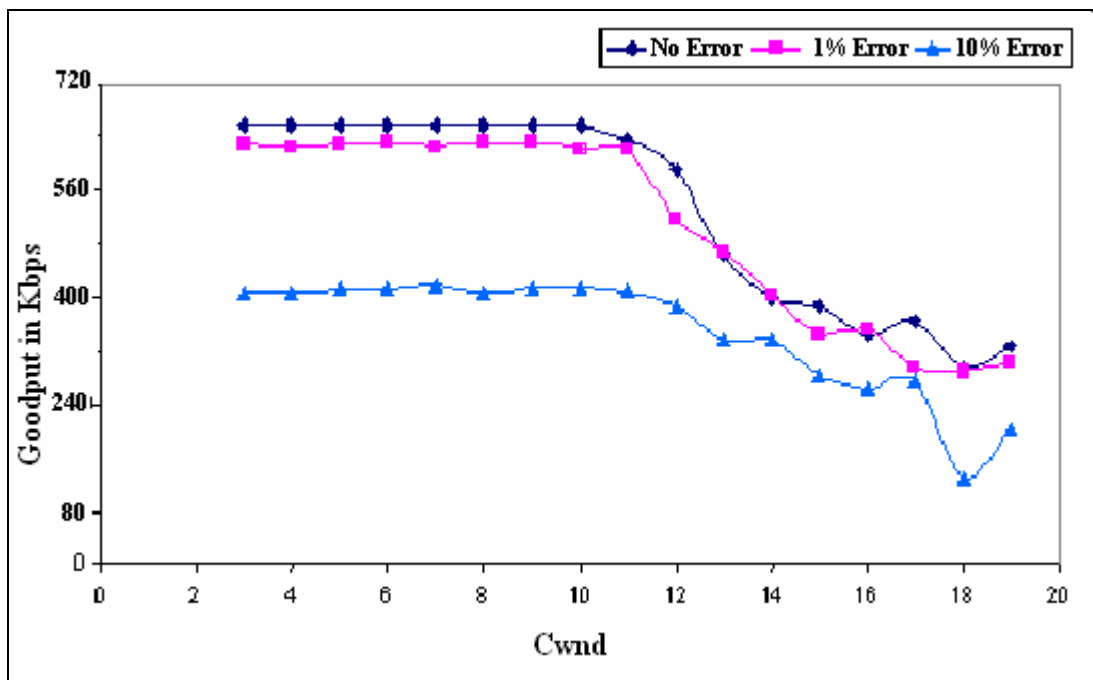


Figure 5: Goodput vs. cwnd curve for mixed wireless scenarios for 20 Source Receiver pairs at various error rates




Figure 6: Goodput vs. cwnd curve for mixed wireless scenarios for different number of Source Receiver pairs and 10% error

2.2.2. Optimal Window Determination

The following figures 5 and 6 show the performances obtained from the set up described in 4.2.1 that are used to set the appropriate cwnd for the connections.

It is observed from the graphs that the performance of the system, as a whole, reaches a maximum for a particular value (or set of values) of cwnd and the cwnd would be set accordingly for obtaining the maximum performance. In the cases, we have obtained a somewhat flattened peak such that the system performance is maximum, or very close to being the maximum, for a number of values of cwnd.

2.2.3. TCP performance evaluation in terms of goodput

In this subsection, we discuss on the basic performance of the modified TCP in comparison with the conventional TCP Reno sources with the various percentages of channel errors.




Figure 7: Comparing the TCP Reno and Modified TCP Goodput under various error scenarios

From figure 7, we have seen that the modified TCP performs almost same with Reno sources having no wireless errors, while outperforming them in the scenarios with wireless errors. An 8% increase in goodput has been obtained for the scenario explained in the previous subsection. A further enhancement is obtained with the introduction of RED gateways [7] at the intermediate routers which leads to a better sharing of the bottleneck as is depicted in figure 8.

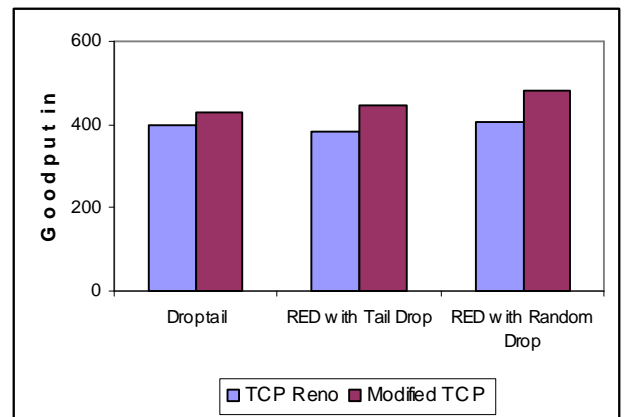


Figure 8: Comparing the TCP Reno and Modified TCP Goodput under various intermediate router queues

2.2.4. TCP fairness for multiple source– destination pairs

In this subsection, we report on the fairness of the modified TCP senders among a number of TCP connections sharing a bottleneck link. Fairness is an important performance metric for a TCP Sender and determines how a number of connections sharing a

bottleneck equally share the network resources amongst themselves.

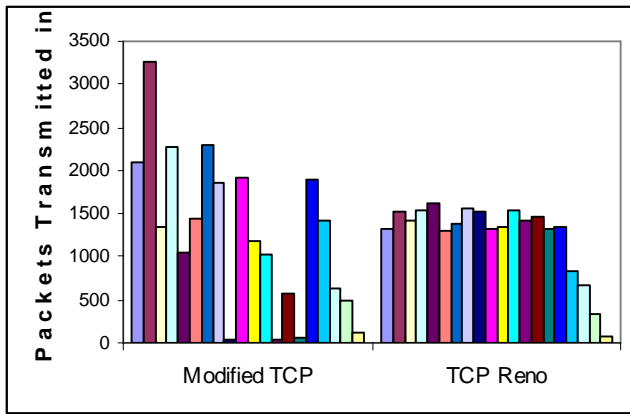


Figure 9: Packets transmitted when the intermediate router queue is DropTail




Figure10: Packets transmitted when the intermediate router queue is a RED queue

In figures 9 and 10 shown above, each colored bar represents the number of packets or segments that have been transmitted over the period of 450 secs. Hence, equal height of the bar would imply better fairness. Or in other words, the smaller the standard deviation, the greater is the fairness achieved by the flows. In the Modified TCP, certain senders do tend to be aggressive, but introduction of active queue management (AQM) techniques (like RED [7] gateways), that tends to curb the aggressive instincts of the misbehaving sources, the problem is somewhat alleviated and the fairness is enhanced to a certain extent. One important thing to be noted in this regard is that the AQM techniques does not slow down the aggressive TCP senders, instead by dropping more packets from these sources provides other sources, starved of resources, to grab some available bandwidth at the bottleneck link.

3. Conclusion and Future work

In this paper, we have proposed a constant congestion window for a sender during its entire lifetime. The window is set appropriately in order that, for a given network scenario, the connection utilizes its fair share in the network. With cwnd set to such an optimal window for a given scenario, we have obtained a comparable performance as TCP Reno in scenarios with no wireless errors, and a minimum 8% increase in goodput, in comparison to TCP Reno, when used with drop tail FIFO queues. This performance has further been improved with the use of RED queues at the intermediate router and 17–18% increase in goodput has been obtained in comparison to TCP Reno sources operating under similar scenario. The fairness, in the scenario when multiple TCP sources send data over a bottleneck link, of this modified sender is almost comparable to the TCP Reno sources operating over similar networks. The fairness is further enhanced with the introduction of AQM technique like RED gateways.

The determination of the optimal value of cwnd is done statically and the entire process has to be repeated for every new scenario in which this modified TCP sender is placed. But ideally, in a real life network scenario, the source would be capable of dynamically determining its optimal window while transmitting data over a particular network and set its cwnd accordingly. Also, in dynamic network scenario, the connection’s fair share can change considerably as compared with that when the cwnd is set. In such a scenario, the source would be equipped to identify such a change in the fair share and accordingly modify its cwnd. We are exploring these aspects currently. Also the effect of introduction of other AQM techniques like CHOCe [8] etc, on the performance of the modified senders – in terms of goodput and fairness – needs further study.

References

- [1] A. K. Ghosh, A. Mukherjee and D. Saha, “Some Simulation Studies to Characterize TCP Window Control Behavior in Wired/Wireless Internetworks”, Proceedings of IEEE International Conference on Personal Wireless Communications (ICPWC) 2005, Jan 23-25, 2005, India
- [2] S. Biaz and N. Vaidya. Discriminating congestion losses from wireless losses using inter-arrival times at receiver. IEEE Symposium ASSET ’99, Texas, USA, 1999
- [3] V. Jacobson and M. J. Karels, Congestion Avoidance and Control, in Proceedings of ACM SIGCOMM, August 1988, vol. 18, pp 314-329
- [4] V. Jacobson, Modified TCP congestion avoidance algorithm, Note sent to end2end-interest mailing list, 1990
- [5] UCB/LBNL/VINT. The ns-2 network simulator, June 2004, <http://www.isi.edu/nsnam/ns/>
- [6] K. Fall and K. Varadhan, The ns Manual, December 13, 2003, <http://www.isi.edu/nsnam/ns/ns-documentation.html>

[7] S. Floyd and V. Jacobson, Random Early Detection Gateways for Congestion Avoidance, IEEE/ACM Transactions on Networking, Aug. 1993

[8] Ao Tang, Jiantao Wang, and Steven H. Low, Understanding CHOKe: Throughput and Spatial Characteristics, IEEE/ACM Transactions on Networking, Vol. 12, No. 4, August 2004

[9] H. Balakrishnan, V.N. Padmanabhan, S. Seshan and R.H. Katz, A comparison of mechanisms for improving

TCP performance over wireless links, IEEE/ACM Transactions on Networking (December 1997).

[10] H. Balakrishnan, S. Seshan, E. Amir and R.H. Katz, Improving TCP/IP performance over wireless networks, in: MOBICOM'95, Berkeley, CA (November 1995).

[11] H. Balakrishnan and R.H. Katz, Explicit loss notification and wireless web performance, in: Proceedings of IEEE GLOBECOM'98 Internet Mini-Conference, Sydney, Australia (November 1998).