

SENDER SIDE INTELLIGENCE FOR TCP THROUGHPUT ENHANCEMENT IN WIRED-CUM-WIRELESS NETWORK

Anup K Ghosh

IBM India Pvt Ltd
Calcutta 700 091, India

Sudipto Das*

Dept. of Computer Science
University of California, Santa Barbara, USA

Rajesh Roy†

Dept. of Computer Science
University of California, Davis, USA

Amitava Mukherjee, Senior Member IEEE

Royal Institute of Technology, Stockholm 16440, Sweden
IBM India Pvt Ltd, Calcutta 700 091, India

ABSTRACT

Performance of the TCP Congestion Control Algorithm has been the focus of research over the last decade. In this paper we propose modifications to TCP Congestion Control to improve its performance in wired-cum-wireless networks. The key idea is to determine the Optimal Congestion Window for a TCP Sender, in a particular network scenario (that corresponds to the fair share of that connection) and keep this congestion window a constant to a point where the fair share in the network has changed considerably from the instance of the calculation of the size of the last window. At this point, the TCP Congestion Window is recalculated according to the nature of new scenario. The proposed mechanism is particularly effective over wireless links, which have an inherently loss-prone nature, as Modified TCP's congestion window being independent of packet losses (be it corruption losses or it congestion losses), keeps transmitting at the same rate as before.

I. INTRODUCTION

A well-known problem of the TCP Congestion Control Algorithm [1–3] in wired – cum – wireless environment is that it relies on packet loss as an indicator of network congestion. In order to ease the congestion scenario and to avoid a congestion collapse, a TCP Reno Sender reduces the congestion window ($cwnd$ expressed in number of segments) and refrains from sending packets. In the wired portion of the network a congested router is invariably the likely reason of packet loss, whereas in the wireless portion a noisy, fading radio channel is the more likely cause of loss. This results in degraded performance of TCP Reno since it lacks the ability to distinguish and isolate congestion loss from wireless loss. Various approaches to address this problem have been discussed and compared in the work by Balakrishnan et al. [4, 5] where the authors carefully contrast three approaches: end-to-end (E2E), Split Connection, and Localized Link Layer methods were carefully contrasted. The split-connection approach [6, 7] violates the semantics of E2E reliability and requires a lot of state maintenance at the base station. On the other hand, TCP Westwood (TCPW) is a new congestion control algorithm that is based on E2E bandwidth estimation [8–11].

*This work was carried out when the author was an undergraduate student at Jadavpur University, Kolkata

†This work was carried out when the author was an undergraduate student at Jadavpur University, Kolkata

In this paper, we propose a TCP Sender side modification of the TCP congestion control algorithm [12] that takes the E2E approach. The crux of idea is that for a given network scenario, the Modified TCP Sender determines its optimal fair share of bandwidth in the link setting its Congestion Window ($cwnd$) in a way that it can effectively transmit with a rate that utilizes the fair share of bandwidth in that scenario. After the $cwnd$ is set to a value optimal for a given network scenario, it is kept constant to the point where the network scenario has changed by a extent significantly altering the connection's fair share. Since the value of $cwnd$ is not decreased at any packet loss indication like retransmission on receipt of a triple DUPACK, or a coarse timeout caused by the expiration of the Retransmission Timer, hence it is not susceptible to performance degradation and $cwnd$ reduction on the occurrences of stray packet losses. This leads to an enhanced performance in the wireless scenarios, as the losses are never an indication of congestion, rather they are caused due to the inherent loss-prone nature of the radio propagation medium. We provide simulation results in support to our claim that constant $cwnd$ can outperform TCP Reno in static (i.e., certain time interval) network scenarios.

The rest of this paper is organized as follows: section II. summarizes some related work; section III. gives the analytical approach; section IV. describes the algorithm used by the sender; section V. summarizes the results obtained by the simulations; section VI. gives an idea of the challenges faced while implementing such a strategy; and finally, section VII. concludes the paper.

II. RELATED WORK

NCPLD [13] compares the measured Round Trip Time (rtt) with the lowest rtt (or that at the knee of the goodput – load curve). If the former is close to the latter, then the cause of a packet loss is assumed to be wireless errors. TCPW [8–11] measures goodput (or reception rate) and uses that rate to set the congestion window whenever a packet is detected lost. If the current goodput is below a certain band around the mean, then the cause of a packet loss is assumed to be congestion, otherwise the cause of loss is attributed to wireless errors.

This paper uses the TCPW bandwidth estimation scheme and also compares the performance of the Modified sender with the TCPW sender. The TCPW sender monitors ACKs to estimate the bandwidth currently used by, and thus available to the connection. More precisely, the sender uses (1) the ACK reception rate and (2) the information an ACK conveys regarding the

amount of data delivered to the destination. The Westwood algorithm is described briefly below. Let us assume that an ACK is received at the source at time t_k , notifying that d_k bytes have been received at the TCP receiver. We can measure the following sample bandwidth used by that connection as $b_k = d_k/\Delta_k$, where $\Delta_k = t_k - t_{k-1}$ and t_{k-1} is the time the previous ACK was received. The following discrete-time filter is used which is obtained by discretizing a continuous low-pass filter using the Tustin approximation

$$b_k = \alpha_k b_{k-1} + (1 - \alpha_k)(b_k + b_{k-1})/2 \quad (1)$$

where b_k is the filtered estimate of the available bandwidth at time $t = t_k$, $\alpha_k = (2\tau - \Delta_k)/(2\tau + \Delta_k)$, and $1/\tau$ is the cutoff frequency of the filter.

Algorithm after n duplicate ACKs

The pseudo code of the algorithm is the following:
 if (n DUPACKs are received)
 $ssthresh = (BWE * RTT_{min})/seg_size$;
 if ($cwin > ssthresh$) /* congestion avoid */
 $cwin = ssthresh$;
 endif
 endif

Here, seg_size identifies the length of the payload of a TCP segment in bits.

Algorithm after coarse timeout expiration

The pseudo code of the algorithm is:
 if (coarse timeout expires)
 $ssthresh = (BWE * RTT_{min})/seg_size$;
 if ($ssthresh < 2$)
 $ssthresh = 2$;
 endif;
 $cwin = 1$;
 endif

III. ANALYTICAL APPROACH

The logic for using a constant window would be summarized as under: As in [1] if we measure the network load by average queue length over fixed intervals of some appropriate length, and L_i be the load at instant i , then, for a congested network we have:

$$L_i = N + \gamma L_{i-1} \quad (2)$$

where N (a constant) accounts for the average arrival rate of the new traffic, and γL_{i-1} accounts for the traffic left from the last time interval. Evidently, the term γL_{i-1} arises when the sender is sending at a rate which is greater than its fair share leading to a fraction of packets from the previous round remaining in the network when the packets from the next round arrive in the network. But if the sender is sending at a rate that utilizes its fair share, the γL_{i-1} vanishes, equation (2) thereby reduces to

$$L_i = N \quad (3)$$

which is a constant, and this forms the basis for use of a constant congestion window.

IV. TCP MODIFICATIONS

The key idea here [12] is that we can divide the entire lifetime of a TCP connection into a finite number of slots such that the connection's fair share in the network remains almost same in a particular slot, i.e. we may assume that the network scenario remains almost static with such slot. A change in the available share of a network, due to some connections leaving the network or some new connections joining, ends a slot and marks the beginning of the next slot. Our proposal is to use a constant TCP Congestion Window during these slots where the network scenario is assumed to remain unchanged. The beginning of a new slot would trigger a window recalculation and the $cwnd$ would be set according to the connection's available share in that slot.

In the proposed mechanism, we use a bandwidth estimation algorithm similar to that of TCPW to obtain an estimate of the available fair share. The change in the rtt measurements is used as a trigger to move to the recalculation phase from the constant window phase (our model uses the knee region in the rtt curve as in [13] to detect a change in fair share and trigger recalculation). In our model, the Modified TCP Sender moves through three distinct phases during its lifetime: the *startup phase* followed by mutually interleaved *window recalculation phase* and *constant window phase*. The three phases are described with some detail as under.

- **The Startup Phase**

At connection setup, the sender has no inkling of the network scenario. In order to impart dynamic nature, the Sender refrains from using typical default values for these essential attributes of the connection. The sender uses a slow start mechanism as in [1]. The sender continues the slow start process for say k rounds, during which it acquires various vital information about the network such as the minimum rtt measurement, a measure of the network bandwidth that the connection etc. After the first k rounds, the sender has acquired enough information about the network and hence calculates $cwnd$ for the first time.

- **Window Recalculation Phase**

When a change in available fair share is detected by the trigger, the TCP sender enters this phase. This is the most crucial phase of the connection, as in this phase, the $cwnd$ is calculated which is kept a constant during the next phase. Hence the performance of the sender, how well it utilizes its share of the network, depends on the $cwnd$ calculated. Along with the window recalculation process, the current value of the smoothed rtt measurements, obtained by passing the coarse rtt measurements of the individual segment through a low pass filter as suggested by Jacobson [1], is also archived for future reference. An efficient Bandwidth Estimation Algorithm must be in place to determine the fair share of the connection in the network. The accuracy of this algorithm in determining the network share would determine the performance of the Modified TCP Sender.

- **The Constant Window Phase**

During this phase of the connection, the $cwnd$ is kept a constant irrespective of the number of ACKs received or any indications of packet loss like DUPACK or a coarse timeout. The sender keeps track of the rtt estimates from the segments that have been delivered. If the percentage change in the smoothed rtt measurements over the archives rtt measure is greater than a specified threshold, the sender exits the constant window phase and enters the Window Recalculation Phase i.e. if $|\frac{rtt_{arc} - rtt_{var}}{rtt_{arc}}| > \beta$, a window recalculation is made. The algorithm's pseudo code is as follows

```

Algorithm
if( slow_start_state)
slow_start(); /* open cwnd by one segment on
each ACK arrival */
else
if( $|\frac{rtt_{arc} - rtt_{var}}{rtt_{arc}}| > \beta$ )
/* Fractional increase greater than threshold */
/* recalculate window and archive the value of  $rtt_{var}$  */
 $cwnd\_ = (Estimated\_Bandwidth * RTT_{min}) / seg\_size\_;$ 
if( $cwnd\_ < 1$ )  $cwnd\_ = 1;$ 
 $rtt_{arc} = rtt_{var} ;$ 
    
```

In the pseudo code, $seg_size_$ identifies the length of the TCP segments in bytes; RTT_{min} is the estimated minimum value of rtt throughout the lifetime of the particular connection, and $Estimated_Bandwidth$ is the Bandwidth Estimate obtained by some Bandwidth Estimation Algorithm.

V. PERFORMANCE ANALYSIS

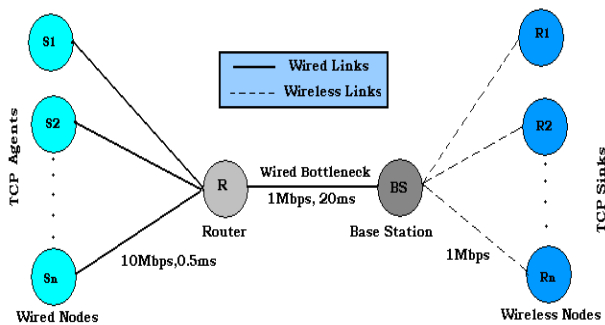


Figure 1: Network Scenario used for simulation

In this section, we report 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 and TCP Westwood sources operating 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 [13]. 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

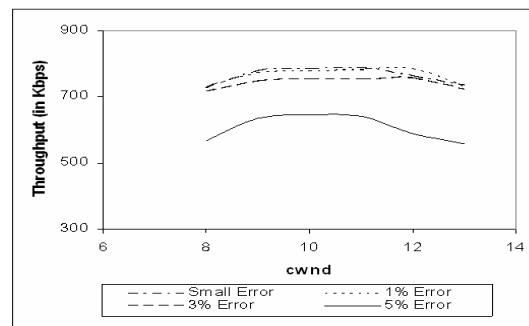


Figure 2: Variation of Throughput with varying $cwnd$ for various bit error rates for single S/R pair

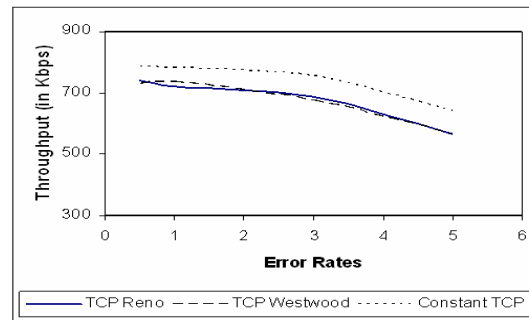


Figure 3: Variation of throughput with varying error rates in scenarios with constant bit error rates for single S/R pair




Figure 4: Variation of Throughput with number of TCP connections sharing the link in scenarios with 5% loss in constant bit error

bytes (1040 bytes with headers) in all experiments. The wireless subnet is error prone. 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 250 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. The wired subnet is error free while the wireless subnet is prone to varying error rates.

The performance of the Modified TCP Senders, TCP Westwood, and TCP Reno has been compared based on the through-

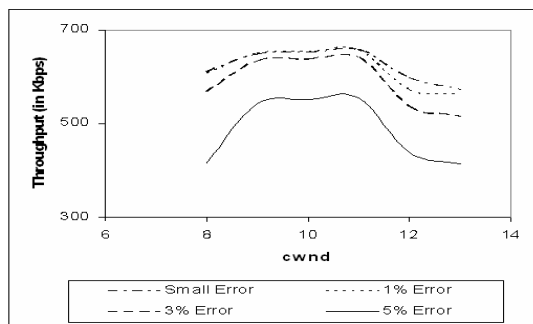


Figure 5: Variation of Throughput with varying cwnd error for various burst error for single S/R pair

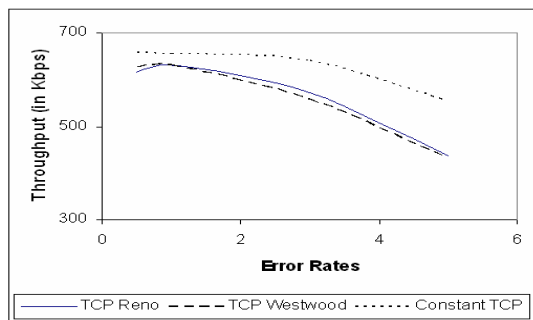


Figure 6: Variation of throughput with varying error rates in scenarios with burst error and for single S/R pair

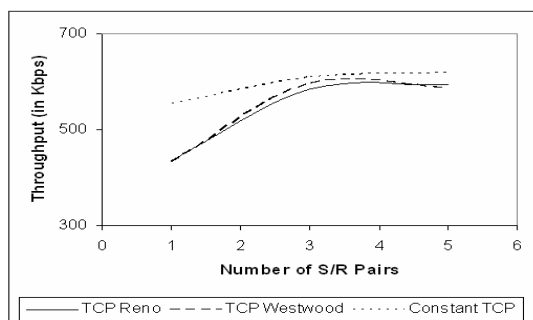


Figure 7: Variation of Throughput with number of TCP connections sharing the link in scenarios with 5% loss in burst error

put metric, i.e. the number of data packets received at the sender. We have analyzed the performance of the Constant Congestion Window aspect of the Modified TCP to assert that in the time slots when the share of a connection in the network remains unchanged, a Constant cwnd TCP outperforms the Reno and Westwood sources in situations with wireless errors. Our assumption is that the share of a connection remains unchanged during the entire period of simulation. The optimal cwnd for a given scenario has been evaluated the cwnd of the modified TCP Sender has been set accordingly. One aspect is to be noted that we are not simulating the entire lifetime of a modified TCP sender. Rather, our analysis is concentrated only on the Constant Congestion Window phase of the connection.

All simulations in this paper have been carried out using the LBL network simulator ns2 [14, 15] with appropriate modifications for implementation of the changes in the modified TCP sender. For comparison with TCP Westwood, the corresponding TCPW modules were used [16].

Figure 1 shows a schematic of the scenario used for simulation. A number of TCP connections share a common wired bottleneck that connects the intermediate router to the base station. When there is only one TCP connection in the network, there is no loss due to congestion. As a result, any packet loss is due to wireless errors. Hence, we can evaluate the performance of the Modified congestion control algorithm in scenarios where wireless loss is the only cause for packet loss. As the numbers of source/receiver pairs are increased, gradually the wired link between the router and the base station would become congested. Hence packets will also be lost both due to congestion as well as wireless errors. Hence, the performance of the Modified TCP sender in congested networks can also be evaluated using the same scenario.

A. Constant Bit Error rates

In the scenarios under consideration, the wireless subnet is prone to constant bit error rates. Figure 2 compares the performance of the Constant cwnd senders for different values of cwnd. As is evident, for every scenario, there exists a value of cwnd (in some cases more than one) for which the performance of the TCP Sender is maximum. This is the optimal cwnd for the given network scenario. In Figure 2, cwnd is expressed in segments.

As is evident from Figure 3, a Constant cwnd TCP outperforms the Reno and Westwood senders operating in similar network conditions. A 10-15% increase in throughput has been obtained as is evident from Figure 3. In this figure, the error rates are expressed as percentage. Figure 4 compares the performance of the TCP variants for multiple connections sharing the wired bottleneck and hence, packet is lost due to congestion as well. The Constant cwnd TCP sender outperforms Reno and Westwood in such scenarios as well

B. Burst Error

This subsection compares the performance of Modified TCP, TCP Reno and TCP Westwood based on the throughput metric. In the scenarios under consideration, the wireless subnet is prone to burst error. The burst error is modeled using a discrete time first order Markov Model. The pattern of errors is described by the transition matrix

$$\mathbf{M} = \begin{pmatrix} p_{BB} & p_{BG} \\ p_{GB} & p_{GG} \end{pmatrix}$$

where p_{BG} is the transition from bad to good, i.e., the conditional probability that successful transmission occurs in a slot given that a failure occurred in the previous slot, and the other entries in the matrix are defined similarly. It is to be noted that $1/(1 - p_{BB})$ represents the average length of a burst of errors, which is described by a geometric random variable.

Figure 5 compares the performance of the Constant cwnd senders for different values of cwnd. As is evident, for every

scenario, there exists a value of $cwnd$ (in some cases more than one) for which the performance of the TCP Sender is maximum. This is the optimal $cwnd$ for the given network scenario. In Figure 5, $cwnd$ is expressed in segments.

When comparing the performance of the Constant $cwnd$ TCP Sender, the t $cwnd$ is set to the optimal value for the given scenario. As is evident from Figure 6, a Constant $cwnd$ TCP outperforms the Reno and Westwood senders operating in similar network conditions. A 10-20% increase in throughput has been obtained as is evident from the figure. In Figure 6, the error rates are expressed as percentage.

Figure 7 compares the performance of the TCP variants for multiple connections sharing the wired bottleneck and hence, packet is lost due to congestion as well. The Constant $cwnd$ TCP sender outperforms Reno and Westwood in such scenarios as well.

VI. CHALLENGES IN IMPLEMENTING THE PROPOSED MECHANISM

In the earlier sections of the paper, we have proposed Sender side modification of the TCP congestion control algorithm. There are certain challenges which need to be met in order for this mechanism to work more efficiently. Firstly, the bandwidth determination algorithm would be precisely able to calculate the available fair share of the connection in the network. An incorrect estimation would negate the performance enhancement which would be gained by not reducing the window in case of wireless errors. Secondly, the triggering mechanism would be able to efficiently determine a change in the available fair share of the bandwidth in the network. Failure to do so would lead to potential over or under utilization of the available fair share in case the fair share of the connection decreases or increases respectively.

VII. CONCLUSION AND FUTURE WORK

In this paper, we propose a Sender side modification of the TCP congestion control algorithm. In addition to this proposal, we have evaluated and compared the performance of the modified TCP sender during a particular phase of its lifetime viz. the Constant Congestion Window Phase. The simulations performed has shown a throughput enhancement of 10-15% as compared to TCP Reno and Westwood in cases with constant bit error rates and about 10-20% in cases with burst error corresponding to a discrete time first-order Markov model. One important aspect of operation of this modified TCP is that the $cwnd$ should be set to a value optimal for a given connection. For this purpose, an efficient Bandwidth Estimation Algorithm would be designed that would dynamically determine a connection's fair share based on certain observed and measured parameters. We are working on to derive a function that would dynamically determine the $cwnd$ during the window recalculation phase.

REFERENCES

- [1] V. Jacobson and M. J. Karels, "Congestion Avoidance and Control," in *Proceedings of ACM SIGCOMM*, (Stanford, CA), pp. 314–329, August 1988.
- [2] V. Jacobson, "Modified TCP Congestion Avoidance Algorithm." Note sent to end2end-interest mailing list, 1990.
- [3] M. Allman, V. Paxson, and W. Stevens, "TCP Congestion Control." RFC 2581, April 1999.
- [4] H. B. et. al., "A comparison of mechanisms for improving TCP performance over wireless links," *IEEE/ACM Transactions on Networking*, vol. 5, pp. 756 – 769, December 1997.
- [5] H. Balakrishnan and R. Katz, "Explicit Loss Notification and Wireless Web Performance," in *Proceedings of IEEE GLOBECOM*, (Sydney, Australia), November 1998.
- [6] R. Yavatkar and N. Bhagwat, "Improving End - to - End Performance of TCP over Mobile Inter-networks," in *Proceedings of Workshop on Mobile Computing Systems and Applications*, December 1994.
- [7] A. Bakre and B. Badrinath, "I-TCP: Indirect TCP for Mobile Hosts," in *International Conference on Distributed Computing Systems*, pp. 136–143, 1995.
- [8] S. Mascolo, C. Casetti, M. Gerla, M. Sanadidi, and R. Wang, "TCP Westwood: End-to-End Bandwidth Estimation for Enhanced Transport over Wireless Links," *ACM Wireless Networks*, (Special issue with Selected papers from Mobicom 2001), vol. 8, pp. 467–479, Sept 2002.
- [9] M. Gerla, B. Ng, M. Sanadidi, M. Valla, and R. Wang, "TCP Westwood with Adaptive Bandwidth Estimation to Improve Efficiency/Friendliness Tradeoffs," *Computer Communications*, vol. 27, pp. 41–58, January 2004.
- [10] J. Chen, F. Paganini, R. Wang, M. Sanadidi, and M. Gerla, "Fluid-flow Analysis of TCP Westwood with RED," in *IEEE GLOBECOM*, (San Francisco, CA), pp. 4064– 4068, December 2003.
- [11] R. Wang, K. Yamada, M. Y. Sanadidi, and M. Gerla, "TCP With Sender-Side Intelligence to Handle Dynamic, Large, Leaky Pipes," *IEEE Journal on Selected Areas in Communications*, vol. 23, pp. 235–248, February 2005.
- [12] R. Roy, S. Das, A. Ghosh, and A. Mukherjee, "Modified TCP Congestion Control Algorithm for Throughput Enhancement in Wired-cum-Wireless Networks," in *Proceeding of SNCNW*, (Luella, Sweden), October 2006.
- [13] N. K. G. Samaraweera, "Non-Congestion Packet Loss Detection for TCP Error Recovery using Wireless Links," in *IEE Proceedings on Communications*, pp. 222 – 230, Aug 1999.
- [14] UCB/LBNL/VINT, "The ns-2 Network Simulator." <http://www.isi.edu/nsnam/ns/>, June 2004.
- [15] K. Fall and K. Varadhan, "The ns Manual." <http://www.isi.edu/nsnam/ns/ns-documentation.html>, December 2003.
- [16] Network Research Lab in Dept. of CS, UCLA, "TCP Westwood Modules for ns-2." http://www.cs.ucla.edu/NRL/hpi/tcpw/tcpw_ns2/tcp-westwood-ns2.html.