TCP/IP Performance over 3G Wireless Links with Rate and Delay Variation

Mun Choon Chan and Ramachandran Ramjee
Bell Labs, Lucent Technologies
munchoon,ramjee@bell-labs.com

ABSTRACT

Wireless link losses result in poor TCP throughput since losses are perceived as congestion by TCP, resulting in source throttling. In order to mitigate this effect, 3G wireless link designers have augmented their system with extensive local retransmission mechanisms. In addition, in order to increase throughput, intelligent channel state based scheduling have also been introduced. While these mechanisms have reduced the impact of losses on TCP throughput and improved the channel utilization, these gains have come at the expense of increased delay and rate variability. In this paper, we comprehensively evaluate the impact of variable rate and variable delay on long-lived TCP performance. We propose a model to explain and predict TCP's throughput over a link with variable rate and/or delay. We also propose a network-based solution called *Ack Regulator* that mitigates the effect of variable rate and/or delay without significantly increasing the round trip time, while improving TCP performance by up to 40%.

Categories and Subject Descriptors

C.2.1 [**Network Architecture and Design**]: Wireless communication

General Terms

Algorithms, Performance, Design

Keywords

TCP, 3G Wireless, Link and Rate Variation

1. INTRODUCTION

Third generation wide-area wireless networks are currently being deployed in the United States in the form of 3G1X technology [10] with speeds up to 144Kbps. Data-only enhancements to 3G1X have already been standardized in the 3G1X-EVDO standard (also called High Data Rate or HDR) with speeds up to 2Mbps [6]. UMTS [24] is the third generation wireless technology in Europe and Asia with deployments planned this year. As these 3G networks provide pervasive internet access, good performance of TCP over these wireless links will be critical for end user satisfaction.

While the performance of TCP has been studied extensively over wireless links [3, 4, 15, 20], most attention has been paid to the impact of wireless channel losses on TCP. Losses are perceived as congestion by TCP, resulting in source throttling and very low net throughput.

In order to mitigate the effects of losses, 3G wireless link designers have augmented their system with extensive local retransmission mechanisms. For example, link layer retransmission protocols such as RLP and RLC are used in 3G1X [22] and UMTS [21], respectively. These mechanisms ensure packet loss probability of less than 1% on the wireless link, thereby mitigating the adverse impact of loss on TCP. While these mechanisms mitigate losses, they also increase delay variability. For example, as we shall see in Section 3, ping latencies vary between 179ms to over 1 second in a 3G1X system.

In addition, in order to increase throughput, intelligent channel state based scheduling have also been introduced. Channel state based scheduling [7] refers to scheduling techniques which take the quality of wireless channel into account while scheduling data packets of different users at the base station. The intuition behind this approach is that since the channel quality varies asynchronously with time due to fading, it is preferable to give priority to a user with better channel quality at each scheduling epoch. While strict priority could lead to starvation of users with inferior channel quality, a scheduling algorithm such as proportional fair [6] can provide long-term fairness among different users. However, while channel-state based scheduling improves overall throughput, it also increases rate variability.

Thus, while the impact of losses on TCP throughput have been significantly reduced by local link layer mechanisms and higher raw throughput achieved by channel-state based scheduling mechanisms, these gains have come at the expense of increased delay and rate variability. This rate and delay variability translates to bursty ack arrivals (also called ack compression) at the TCP source. Since TCP uses ack clocking to probe for more bandwidth, bursty ack arrival leads to release of a burst of packets from the TCP source. When this burst arrives at a link with variable rate or delay, it could result in multiple packet losses. These multiple losses significantly degrade TCPs throughput.

In this paper, we make three main contributions. First,

MOBICOM'02 September 23–28, Atlanta, Georgia, USA

we comprehensively evaluate the impact of variable rate and variable delay on long-lived TCP performance. Second, we propose a model to explain and predict TCP's throughput over a link with variable rate and/or delay. Third, we propose a network-based solution called *Ack Regulator* that mitigates the effect of variable rate and/or delay without significantly increasing the round trip time, thereby improving TCP performance.

The remaining sections are organized as follows. In Section 2, we discuss related work. In Section 3, we present the motivation for our work using traces from a 3G1X system. In Section 4, we describe a model for computing the throughput of a long-lived TCP flow over links with variable rate and variable delay. We then present a simple network-based solution, called *Ack Regulator*, to mitigate the effect of variable rate/delay in Section 5. In Section 6, we present extensive simulation results that compare TCP performance with and without Ack Regulator, highlighting the performance gains using the Ack Regulator when TCP is subjected to variable rate and delay. Finally, in Section 7, we present our conclusions.

2. RELATED WORK

In this section, we review prior work on improving TCP performance over wireless networks. Related work on the modeling of TCP performance is presented in Section 4.

A lot of prior work has focused on avoiding the case of a TCP source misinterpreting packet losses in the wireless link as congestion signals. In [4], a snoop agent is introduced inside the network to perform duplicate ack suppression and local retransmissions on the wireless link to enhance TCP performance. In [3], the TCP connection is split into two separate connections, one over the fixed network and the second over the wireless link. The second connection can recover from losses quickly, resulting in better throughput. Link-layer enhancements for reducing wireless link losses including retransmission and forward error correction have been proposed in [20]. Link layer retransmission is now part of both the CDMA2000 and UMTS standards [10, 24]. In order to handle disconnections (a case of long-lived loss), M-TCP has been proposed [8]. The idea is to send the last ack with a zero-sized receiver window so that the sender can be in persist mode during the disconnection. Link failures are also common in Ad Hoc networks and techniques to improve TCP performance in the presence of link failures have been proposed in [11]. Note that none of these approaches address specifically the impact of delay and rate variation on TCP, which is the focus of this paper.

Several generic TCP enhancements with special applicability to wireless links are detailed in [12, 13]. These include enabling the Time Stamp option, use of large window size and window scale option, disabling Van Jacobson header compression, and the use of Selective Acknowledgments (Sack). Large window size and window scaling are necessary because of the large delay of wireless link while Sack could help TCP recover from multiple losses without the expensive timeout recovery mechanism.

Another issue with large delay variation in wireless links is spurious timeouts where TCP unnecessarily retransmits a packet (and lowers its congestion window to a minimum) after a timeout, when the packet is merely delayed. In [13], the authors refer to rate variability due to periodic allocation and de-allocation of high-speed channels in 3G networks as Bandwidth Oscillation. Bandwidth Oscillation can also lead to spurious timeouts in TCP because as the rate changes from high to low, the rtt value increases and a low Retransmission Timeout (RTO) value causes a spurious retransmission and unnecessarily forces TCP into slow start. In [15], the authors conduct experiments of TCP over GSM circuit channels and show that spurious timeouts are extremely rare. However, 3G wireless links can have larger variations than GSM due to processing delays and rate variations due to channel state based scheduling. Given the increased variability on 3G packet channels, the use of TCP time stamp option for finer tracking of TCP round trip times and possibly the use of Eifel retransmission timer [16] instead of the conventional TCP timer can help avoid spurious timeouts.

As mentioned earlier, the effect of delay and rate variability is ack compression and this results in increased burstiness at the source. Ack compression can also be caused by bidirectional flows over regular wired networks or single flow over networks with large asymmetry. This phenomenon has been studied and several techniques have been proposed to tackle the burstiness of ack compression. In order to tackle burstiness, the authors in [18] propose several schemes that withholds acks such that there is no packet loss at the bottleneck router, resulting in full throughput. However, the round trip time is unbounded and can be very large. In [23], the authors implement an ack pacing technique at the bottleneck router to reduce burstiness and ensure fairness among different flows. In the case of asymmetric channels, solutions proposed [5] include ack congestion control and ack filtering (dropping of acks), reducing source burstiness by sender adaptation and giving priority to acks when scheduling inside the network. However, the magnitude of asymmetry in 3G networks is not large enough and can be tolerated by TCP without ack congestion control or ack filtering according to [12].

Note that, in our case, ack compression occurs because of link variation and not due to asymmetry or bidirectional flows. Thus, we require a solution that specifically adapts to link variation. Moreover, the node at the edge of the 3G wireless access network is very likely to be the bottleneck router (given rates of 144Kbps to 2Mbps on the wireless link) and is the element that is exposed to varying delays and service rates. Thus, this node is the ideal place to regulate the acks in order to improve TCP performance. This is discussed in more detail in the next section.

3. MOTIVATION

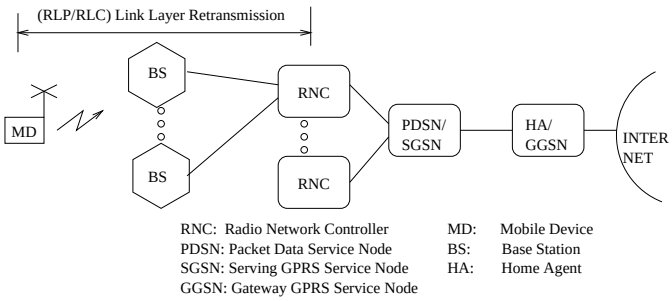

Figure 1: 3G network architecture

A simplified architecture of a 3G wireless network is shown

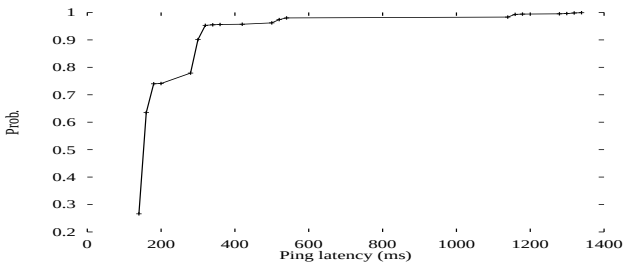

Figure 2: CDF of Ping Latencies

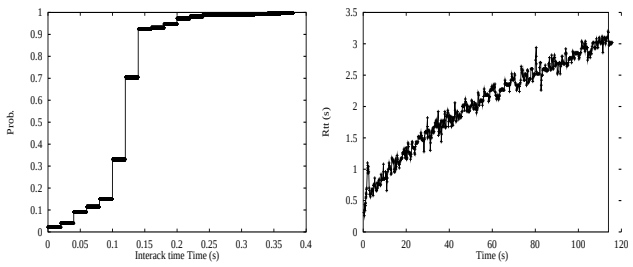

(a) TCP Ack Inter-arrival (b) TCP rtt value

Figure 3: 3G Link Delay Variability

in Figure 1. The base stations are connected to a node called the Radio Network Controller (RNC). The RNC performs CDMA specific functions such as soft handoffs, encryption, power control etc. It also performs link layer retransmission using RLP(RLC) in 3G1X(UMTS) system. In the 3G1X system, the RNC is connected to a PDSN using a GRE tunnel (one form of IP in IP tunnel) and the PDSN terminates PPP with the mobile device. If Mobile IP service is enabled, the PDSN also acts as a Foreign Agent and connects to a Home Agent. In the UMTS system, the RNC is connected to a SGSN using a GTP tunnel (another form of IP in IP tunnel); the SGSN is connected to a GGSN, again through a GTP tunnel. Note that the tunneling between the various nodes allows for these nodes to be connected directly or through IP/ATM networks.

In this architecture, the RNC receives a PPP/IP packet through the GRE/GTP tunnel from the PDSN/SGSN. The RNC fragments this packet into a number of radio frames and then performs transmission and local retransmission of these radio frames using the RLP(RLC) protocol. The base station (BS) receives the radio frames from the RNC and then schedules the transmission of the radio frames on the wireless link using a scheduling algorithm that takes the wireless channel state into account. The mobile device receives the radio frames and if it discovers loss of radio frames, it requests local retransmission using the RLP(RLC) protocol. Note that, in order to implement RLP(RLC), the RNC needs to keep a per-user queue of radio frames. The RNC can typically scale up to tens of base stations and thousands of active users.

In order to illustrate the variability seen in a 3G system, we obtained some traces from a 3G1X system. The system consisted of an integrated BS/RNC, a server connected to the RNC using a 10Mbps Ethernet and a mobile device connected to the BS using a 3G1X link with 144Kbps downlink in infinite burst mode and 8Kbps uplink. The infinite burst mode implies that the rate is fixed and so the system only had delay variability.

Figure 2 plots the cumulative distribution function (cdf) of ping latencies from a set of 1000 pings from the server to the mobile device (with no observed loss). While about 75% of the latency values are below 200ms, the latency values go all the way to over 1s with about 3% of the values higher than 500ms.

In the second experiment, a TCP source at the server using Sack with timestamp option transferred a 2MB file to the mobile device. The MTU was 1500 bytes with user data size of 1448 bytes. The buffer at the RNC was larger than the TCP window size[1]. and thus, the transfer resulted in no TCP packet loss and a maximal throughput of about 135Kb/s. The transmission time at the bottleneck link is $1.448 * 8/135 = 86$ms. If the wireless link delay were constant, the TCP acks arriving at the source would be evenly spaced with a duration of 172ms because of the delayed ack feature of TCP (every 2 packets are acked rather than every packet). Figure 3(a) plots the cdf of TCP ack inter-arrival time (time between two consecutive acks) at the server. As can be seen, there is significant ack compression with over 10% of the acks arriving within 50ms of the previous ack. Note that the ack packet size is 52 bytes (40 + timestamp) and ack transmission time on the uplink is $52 * 8/8$=52ms; an interack spacing of less then 52ms is a result of uplink delay variation.

Note that the delay variability and the resulting ack compression did not cause any throughput degradation in our system. This was due to the fact that the buffering in the system was greater than the TCP window size resulting in no buffer overflow loss. Figure 3(b) depicts the TCP round trip time (rtt) values over time. Since the buffer at the RNC is able to accommodate the whole TCP window, the rtt increases to over 3s representing a case of over 30 packets in the buffer at the RNC ($30 * 0.086 = 2.5$s). Given an average ping latency of $215ms$ and a transmission time of $86ms$ for a 1500 byte packet, the bandwidth delay product of the link is approximately $(0.215 + 0.86) * 135$=$5KB$ or about 3 packets. Thus, the system had a buffer of over 10 times the bandwidth delay product. Given that we had only one TCP flow in the system, a buffer of over 64KB is not a problem. But, if every TCP flow is allocated a buffer of 64KB, the buffer requirements at the RNC would be very expensive, since the RNC supports thousands of active users.

Even discounting the cost of large buffers, the inflated rtt value due to the excessive buffering has several negative consequences as identified in [15]. First, an inflated rtt implies a large retransmission timeout value (rto). In the case of multiple packet losses (either on the wireless link or in a router elsewhere in the network), a timeout-based recovery would cause excessive delay, especially if exponential backoff gets invoked. Second, if the timestamp option is not used, the rtt sampling rate is reduced and this can cause spurious timeouts. Third, there is a higher probability that the data in the queue becomes obsolete (for e.g., due to user aborting the transmission), but the queue will still have to be drained resulting in wasted bandwidth.

Thus, while excessive buffering at the RNC can absorb the variability of the wireless links without causing TCP throughput degradation, it has significant negative side effects, making it an undesirable solution.

[1] We did not have control over the buffer size at the RNC in our system.

4. MODEL

In this section, we model the performance of a single long-lived TCP flow over a network with a single bottleneck server that exhibits rate variation based on a given general distribution and a single wireless link attached to the bottleneck server that exhibits delay variation based on another given distribution.

We use a general distribution of rate and delay values for the discussion in this section since we would like to capture the inherent variation in rate and delay that is a characteristic of the 3G wireless data environment. Given that the wireless standards are constantly evolving, the actual rate and delay distribution will vary from one standard or implementation to another and is outside the scope of this paper. Later, in Section 6, we will evaluate TCP performance over a specific wireless link, the 3G1X-EVDO (HDR) system, using simulation.

We would like to model TCP performance in the case of variable rate and delay for two reasons. One, we would like to understand the dynamics so that we can design an appropriate mechanism to improve TCP performance. Two, we would like to have a more accurate model that specifically takes the burstiness caused by ack compression due to rate/delay variability into account.

TCP performance modeling has been extensively studied in the literature [1, 2, 9, 14, 17, 19]. Most of these models assume constant delay and service rate at the bottleneck router and calculate TCP throughput in terms of packet loss probability and round trip time. In [19], the authors model TCP performance assuming deterministic time between congestion events [1]. In [17], the authors improve the throughput prediction of [19] assuming exponential time between congestion events (loss indications as Poisson). In our case, ack compressions and link variation causes bursty losses and the deterministic or Poisson loss models are not likely to be as accurate. In [9], the authors model an UMTS wireless network by extending the model from [19] and inflating the rtt value to account for the average additional delay incurred on the wireless link. However, we believe this will not result in an accurate model because 1) the rtt value in [19] is already an end-to-end measured value and 2) the loss process is much more bursty than the deterministic loss assumption in [19]. In [2], the authors observe that mean values are not sufficient to predict the throughput when routers have varying bandwidth and show that increasing variance for the same mean service rate decreases TCP throughput. However, the approach is numerical, and provides little intuition in the case of delay variance.

Our approach starts with the model in [14] which describes how TCP functions in an "ideal" environment with constant round trip time, constant service rate and suffers loss only through buffer overflow. A brief summary of the result from [14] is presented here before we proceed to our model, which can be seen as an extension. We chose to extend the model in [14] since it makes no assumption about the nature of loss event process (which is highly bursty in our case) and explicitly accounts for link delay and service rate (which are variable in our case). For simplicity, we will only discuss the analysis of TCP Reno. TCP Sack can be analyzed similarly. We also assume that the sender is not limited by the maximum receiver window; simple modifications can be made to the analysis for handling this case.

Figure 4(a) shows how the TCP congestion window varies

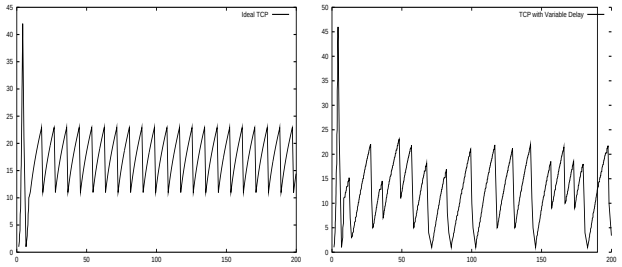

(a) Constant delay (b) Variable delay

Figure 4: TCP Congestion Window Evolution over time

in a constant rate and delay setting. The initial phase where TCP tries to probe for available bandwidth is the *Slow Start* phase. After slow start, TCP goes to *Congestion avoidance* phase. In the case of long-lived TCP flow, one can focus only on the congestion avoidance phase. Let μ be the constant service rate, τ the constant propagation delay, T the minimum round trip time $(\tau + 1/\mu)$ and B the buffer size. The congestion window follows a regular saw-tooth pattern, going from W_0 to W_{max}, where $W_0 = W_{max}/2$ and $W_{max} = \mu\tau + B + 1$. Due to the regularity of each of the saw-tooth, consider one such saw-tooth. Within a single saw-tooth, the congestion avoidance phase is divided into two epochs. In the first epoch, say epoch A, the congestion window increases from W_0 to μT, in time t_A with number of packets sent n_A. In the second epoch, say epoch B, the congestion window increases from μT to W_{max}, in time t_B with number of packets sent n_B. TCP throughput (ignoring slow start) is simply given by $(n_A + n_B)/(t_A + t_B)$ where

$$t_A = T(\mu T - W_0) \quad (1)$$

$$n_A = (W_0 t_A + t_A^2/(2T))/T \quad (2)$$

$$t_B = (W_{max}^2 - (\mu T)^2)/(2\mu) \quad (3)$$

$$n_B = \mu t_B \quad (4)$$

This model, while very accurate for constant μ and T, breaks down when the constant propagation and service rate assumptions are not valid. Figure 4(b) shows how the congestion window becomes much more *irregular* when there is substantial variation in the wireless link delay. This is because the delay variation and ack compression result in multiple packet losses.

There are three main differences in the TCP congestion window behavior under variable rate/delay from the traditional saw-tooth behavior. First, while the traditional saw-tooth behavior always results in one packet loss due to buffer overflow, we have possibilities for multiple packet losses due to link variation. To account for this, we augment our model with parameters $p1, p2, p3$ representing respectively the conditional probability of a single packet loss, double packet loss, and three or more packet losses. Note that, $p1 + p2 + p3 = 1$ by this definition. Second, while the loss in the traditional saw-tooth model always occurs when window size reaches $W_{max} = \mu\tau + B + 1$, in our model losses can occur at different values of window size, since μ and τ are now both variables instead of constants. We capture this by a parameter $W_f = \sqrt{\Sigma_{i=1}^{N} W_{max_i}^2/N}$, that is the square root of the second moment of the W_{max} values of each cycle. The reason we do this instead of obtaining a simple mean of W_{max} values is because throughput is related to W_f quadratically (since it is the area under the curve in the

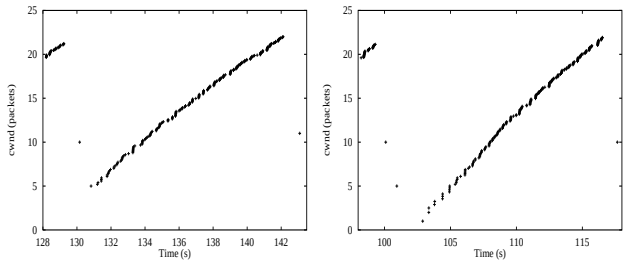

(a) Two packet loss (b) Three packet loss

Figure 5: Congestion Window with multiple losses

congestion window graph). Third, due to the fact that we have multiple packet losses in our model, we need to consider timeouts and slow starts in our throughput calculation. We represent the timeout duration by the T_0 parameter which represents the average timeout value, similar to the timeout parameter in [19].

We now model the highly variable congestion window behavior of a TCP source under rate/delay variation. We first use W_f instead of W_{max}. We approximate τ (the propagation delay) by $\hat{\tau}$, the average link delay in the presence of delay variability. We replace μ (the service rate) by $\hat{\mu}$, the average service rate in the presence of rate variability. Thus, T becomes $\hat{T} = (\hat{\tau} + 1 / \hat{\mu})$. Now consider three different congestion window patterns: with probability $p1$, single loss followed by congestion avoidance, with probability $p2$, double loss followed by congestion avoidance, and with probability $p3$, triple loss and timeout followed by slow start and congestion avoidance[2].

First, consider the single loss event in the congestion avoidance phase. This is the classic saw-tooth pattern with two epochs as identified in [14]. Lets call these A1 and B1 epochs. In epoch A1, window size grows from W_{01} to $\hat{\mu}\hat{T}$ in time, t_{A1}, with number of packets transmitted, n_{A1}. In epoch B1, window size grows from $\hat{\mu}\hat{T}$ to W_f in time, t_{B1}, with number of packets transmitted, n_{B1}. Thus, *with probability* $p1$, $n_{A1}+n_{B1}$ *packets are transmitted in time* $t_{A1}+t_{B1}$ where

$$W_{01} = (int)W_f/2 \quad (5)$$
$$t_{A1} = \hat{T}(\hat{\mu}\hat{T} - W_{01}) \quad (6)$$
$$n_{A1} = (W_{01}t_{A1} + t_{A1}^2/(2\hat{T}))/\hat{T} \quad (7)$$
$$t_{B1} = (W_f^2 - (\hat{\mu}\hat{T})^2)/(2\hat{\mu}) \quad (8)$$
$$n_{B1} = \hat{\mu}t_{B1} \quad (9)$$

Next, consider the two loss event. An example of this event is shown in Figure 5(a). The trace is obtained using ns-2 simulation described in Section 6. In this case, after the first fast retransmit (around 130s), the source receives another set of duplicate acks to trigger the second fast retransmit (around 131s). This fixes the two losses and the congestion window starts growing from W_{02}. The second retransmit is triggered by the new set of duplicate acks in response to the first retransmission. Thus, the duration between the first and second fast retransmit is the time required for the first retransmission to reach the receiver (with a full buffer) plus the time for the duplicate ack to return to the sender. In other words, this duration can be approximated by the average link delay with a full buffer, $\hat{T} + B/\hat{\mu}=t_R$. We have three epochs now, epoch t_R (time 130-131s)with one retransmission and zero new packet, epoch A2 (131-137s) with window size growing from W_{02} to $\hat{\mu}\hat{T}$ in time, t_{A2}, with number of packets transmitted, n_{A2}, and epoch B1 (137-143s) as before. Thus, *with probability* $p2$, $n_{A2}+n_{B1}$ *packets are transmitted in time* $t_R+t_{A2}+t_{B1}$ where

$$W_{02} = (int)W_{01}/2 \quad (10)$$
$$t_R = \hat{T} + B/\hat{\mu} \quad (11)$$
$$t_{A2} = \hat{T}(\hat{\mu}\hat{T} - W_{02}) \quad (12)$$
$$n_{A2} = (W_{02}t_{A2} + t_{A2}^2/(2\hat{T}))/\hat{T} \quad (13)$$

Finally, consider the three loss event. An example of this event is shown in Figure 5(b). In this case, after the first fast retransmit, we receive another set of duplicate acks to trigger the second fast retransmit. This does not fix the three losses and TCP times out. Thus, we now have five epochs: first is the retransmission epoch (100-101s) with time t_R and zero new packet, second is the timeout epoch (101-103s) with time T_0 and zero new packet, third is the slow start epoch (103-106s) where the window grows exponentially up to previous ssthresh value of W_{03} in time t_{ss}(Eqn. 15) with number of packets transmitted n_{ss} (Eqn. 16)[3], fourth is epoch A3 (106-111s) where the window size grows from W_{03} to $\hat{\mu}\hat{T}$ in time t_{A3}(Eqn. 17) with number of packets transmitted n_{A3} (Eqn. 18), and fifth is epoch B1 (111-118s) as before. Thus, *with probability* $p3$, $n_{ss}+n_{A3}+n_{B1}$ *packets are transmitted in time* $t_R+T_0+t_{ss}+t_{A2}+t_{B1}$ where

$$W_{03} = (int)W_{02}/2 \quad (14)$$
$$t_{ss} = \hat{T}log_2(W_{03}) \quad (15)$$
$$n_{ss} = W_{03}/\hat{T} \quad (16)$$
$$t_{A3} = \hat{T}(\hat{\mu}\hat{T} - W_{03}) \quad (17)$$
$$n_{A3} = (W_{03}t_{A3} + t_{A3}^2/(2\hat{T}))/\hat{T} \quad (18)$$

Given that the different types of packet loss events are independent and using $p1+p2+p3=1$, the average TCP throughput can now be approximated by a weighted combination of the three types of loss events to be

$$\frac{p3*(n_{ss}+n_{A3}) + p2*n_{A2} + p1*n_{A1} + n_{B1}}{p3*(t_R+T_0+t_{ss}+t_{A3}) + p2*(t_R+t_{A2}) + p1*t_{A1} + t_{B1}} \quad (19)$$

If any of t_* are less than 0, those respective epochs do not occur and we can use the above equation while setting the respective n_*, t_* to zero.

In this paper, we infer parameters such as $p1$, $p2$, $p3$, W_f, and T_0 from the traces. Models such as [19] also infer the loss probability, round trip time, and timeout durations from traces.

Table 1 lists the various parameters used by the different models for simulations with rate and delay variability. We use a packet size of 1000 bytes, a buffer of 10 which represents the product of the average bandwidth times average delay and we ensure that the source is not window limited. TD and TO denote the number of loss events that are of the triple duplicate and timeout type respectively and these

[2]We assume that three or more packet losses result in a timeout; this is almost always true if the source is TCP reno.

[3]using analysis similar to [14] and assuming adequate buffer so that there is no loss in slow start.

Item	Rate(Kb/s)	Delay(ms)	pkts	TD	TO	T_0	rtt	p1	p2	W_f	$\hat{T}$	$\hat{\mu}$
1	200	400	89713	401	1	1.76	616.2	0.998	0.000	22.00	440	25.0
2	200	380+e(20)	83426	498	1	1.71	579.3	0.639	0.357	21.38	442	25.0
3	200	350+e(50)	78827	489	12	1.79	595.8	0.599	0.367	21.24	461	25.0
4	200	300+e(100)	58348	496	114	1.92	606.0	0.339	0.279	18.95	517	25.0
5	u(200,20)	400	82180	504	1	1.75	578.1	0.535	0.460	21.61	400	24.74
6	u(200,50)	400	74840	517	29	1.80	579.9	0.510	0.403	20.52	400	23.34
7	u(200,75)	400	62674	516	81	1.86	585.9	0.398	0.348	19.05	400	20.93
8	u(200,50)	350+e(50)	70489	507	43	1.81	595.7	0.496	0.377	20.15	459	23.34
9	u(200,75)	300+e(100)	53357	497	93	2.03	635.7	0.404	0.298	17.78	511	20.93

Table 1: Simulation and Model parameters

Item	Simulator Goodput	Model 1 [19] (accu.)	Model 2 [17] (accu.)	Model 3[Eqn. 19] (accu.)
1	199.8	228.5(0.86)	201.9(0.99)	199.8(1.0)
2	185.4	208.0(0.88)	186.0(1.0)	186.0(1.0)
3	175.1	195.5(0.88)	177.2(0.99)	180.9(0.97)
4	129.4	145.3(0.88)	153.7(0.81)	137.0(0.94)
5	182.5	205.2(0.88)	184.6(0.99)	181.3(0.99)
6	166.2	186.0(0.88)	174.6(0.95)	165.2(0.99)
7	139.2	158.4(0.86)	163.4(0.83)	137.2(0.99)
8	156.5	174.6(0.88)	166.5(0.94)	160.2(0.97)
9	118.4	134.0(0.87)	142.6(0.80)	125.0(0.94)

Table 2: Simulation and Model throughput values

values are used by models in [19] and [17]. The simulation is run for 3600 seconds. We simulate delay and rate variability with exponential and uniform distributions respectively ($u(a,b)$ in the table represents uniform distribution with mean a and standard deviation b while $e(a)$ represents an exponential distribution with mean a). The details of the simulation are presented in Section 6.

Table 2 compares the throughput of simulation of different distributions for rate and delay variability at the server and the throughput predicted by the exact equation of the model in [19], the Poisson model in [17] and by equation 19. The accuracy of the prediction, defined as 1 minus the ratio of the difference between the model and simulation throughput value over the simulation throughput value, is listed in the parenthesis. As the last column shows, the match between our model and simulation is extremely accurate when the delay/rate variation is small and the match is still well over 90% even when the variation is large. The Poisson loss model used in [17] performs very well when the variability is low but, understandably, does not predict well when variability increases. The deterministic loss model seems to consistently overestimate the throughput.

From Table 1, one can clearly see the impact of delay and rate variability. As the variability increases, the probability of double loss, $p2$, and three or more losses, $p3$=(1-$p2$-$p1$), start increasing while the goodput of the TCP flow starts decreasing. For example, comparing case 1 to case 4, $p1$ decreases from 0.998 to 0.339 while $p3$ increases. Increases in $p2$ and $p3$ come about because when the product $\hat{T}\hat{\mu}$ decreases, a pipe that used to accommodate more packets suddenly becomes smaller causing additional packet losses. Given that $n_{A1}/t_{A1} > n_{A2}/(t_R + t_{A2}) > (n_{ss} + n_{A3})/(t_R + T_0 + t_{ss} + t_{A3})$, any solution that improves TCP performance must reduce the occurrence of multiple packet losses, $p2$ and $p3$. We present a solution that tries to achieve this in the next section.

5. ACK REGULATOR

In this section, we present our network-based solution for improving TCP performance in the presence of varying bandwidth and delay. The solution is designed for improving the performance of TCP flows towards the mobile host (for downloading-type applications) since links like HDR are designed for such applications. The solution is implemented at the wireless edge, specifically at the RNC, at the layer just above RLP/RLC. Note that, in order to implement the standard-based RLP/RLC, the RNC already needs to maintain a per-user queue. Our solution requires a per-TCP-flow queue, which should not result in significant additional overhead given the low bandwidth nature of the wireless environment. We also assume that the data and ack packets go through the same RNC; this is true in the case of 3G networks where the TCP flow is anchored at the RNC because of the presence of soft handoff and RLP.

We desire a solution that is simple to implement and remains robust across different implementations of TCP. To this end, we focus only on the congestion avoidance phase of TCP and aim to achieve the classic saw-tooth congestion window behavior even in the presence of varying rates and delays by controlling the buffer overflow process in the bottleneck link. We also assume for this discussion that every packet is acknowledged (the discussion can be easily modified to account for delayed acks where single ack packets acknowledge multiple data packets).

Our solution is called the Ack Regulator since it regulates the flow of acks back to the TCP source. The intuition behind the regulation algorithm is to avoid any buffer overflow loss until the congestion window at the TCP source reaches a pre-determined threshold and beyond that, allow only a single buffer overflow loss. This ensures that the TCP source operates mainly in the congestion avoidance phase with congestion window exhibiting the classic saw-tooth behavior. Before we present our solution, we describe two variables that will aid in the presentation of our solution.

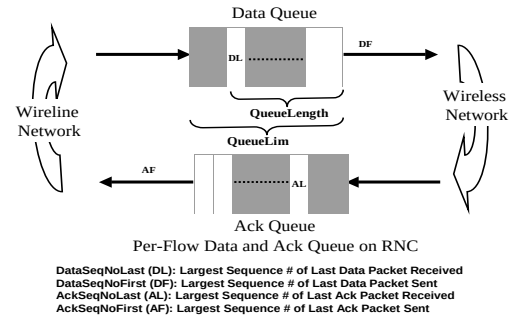

Figure 6: Ack Regulator Implementation

ConservativeMode: Mode of operation during which each time an ack is sent back towards the TCP source, there is buffer space for at least two data packets from the source.

Note that if TCP operates in the congestion avoidance phase, there would be no buffer overflow loss as long as the algorithm operates in conservative mode. This follows from the fact that, during congestion avoidance phase, TCP increases its window size by at most one on reception of an ack. This implies that on reception of an ack, TCP source sends either one packet (no window increase) or two packets (window increase). Therefore, if there is space for at least two packets in the buffer at the time of an ack being sent back, there can be no packet loss.

AckReleaseCount: The sum of total number of acks sent back towards the source and the total number of data packets from the source in transit towards the RNC due to previous acks released, assuming TCP source window is constant.

AckReleaseCount represents the number of packets that can be expected to arrive in the buffer at the RNC assuming that the source window size remains constant. Thus, buffer space equal to AckReleaseCount must be reserved whenever a new ack is sent back to the source if buffer overflow is to be avoided.

```
On Enque of Ack/Deque of data packet:
1.  AcksSent=0;
2.  BufferAvail=QueueLim-QueueLength;
3.  BufferAvail-=(AckReleaseCount+ConservativeMode);
4.  if (BufferAvail>=1)
5.    if (AckSeqNoLast-AckSeqNoFirst<BufferAvail)
5.1       AcksSent+=(AckSeqNoLast-AckSeqNoFirst);
5.2       AckSeqNoFirst=AckSeqNoLast;
      else
5.3       AckSeqNoFirst+=BufferAvail;
5.4       AcksSent+=BufferAvail;
5.5   Send acks up to AckSeqNoFirst;
```

Figure 7: Ack Regulator processing at the RNC

Figure 6 shows the data and ack flow and the queue variables involved in the Ack Regulator algorithm, which is presented in Figure 7. We assume for now that the AckReleaseCount and ConservativeMode variables are as defined earlier. We later discuss how these variables are updated. The Ack Regulator algorithm runs on every transmission of a data packet (deque) and every arrival of an ack packet (enque). The instantaneous buffer availability in the data queue is maintained by the BufferAvail variable (line 2). BufferAvail is then reduced by the AckReleaseCount and the ConservativeMode variables (line 3).

Depending on the value of the ConservativeMode variable (1 or 0), the algorithm operates in two modes, a conservative mode or a non-conservative, respectively. In the conservative mode, an extra buffer space is reserved in the data queue to ensure that there is no loss even if TCP congestion window is increased by 1, while, in the non-conservative mode, a single packet loss occurs if TCP increases its congestion window by 1. Now, after taking AckReleaseCount and ConservativeMode variables into account, if there is at least one buffer space available (line 4) and, if the number of acks present in the ack queue (AckSeqNoLast - AckSeqNoFirst) is lesser than BufferAvail, all those acks are sent to the source (lines 5.1,5.2); otherwise only BufferAvail number of acks are sent to the source (lines 5.3,5.4).

Note that the actual transmission of acks (line 5.5) is not presented here. The transmission of AcksSent acks can be performed one ack at a time or acks can be bunched together due to the cumulative nature of TCP acks. However, care must be taken to preserve the duplicate acks since the TCP source relies on the number of duplicate acks to adjust its congestion window. Also, whenever three or more duplicate acks are sent back, it is important that one extra buffer space be reserved to account for the fast retransmission algorithm. Additional buffer reservations of two packets to account for the Limited Transmit algorithm [12] can also be provided for, if necessary.

```
1. Initialize ConservativeMode=1; α = 2
2. On Enque of ack packet:
   if ((DataSeqNoLast-AckSeqNoFirst)>α*QueueLim)
      ConservativeMode=0;
3. On Enque and Drop of data packet:
   Conservative Mode=1;
4. On Enque/Deque of data packet:
   if (((DataSeqNoLast-AckSeqNoFirst)<α*QueueLim/2)
      OR (DataQueueLength==0))
      ConservativeMode=1;
```

Figure 8: ConservativeMode updates

We now present the algorithm (Figure 8) for updating the ConservativeMode variable which controls the switching of the Ack Regulator algorithm between the conservative and the non-conservative modes. The algorithm starts in conservative mode (line 1). Whenever a targeted TCP window size is reached (in this case, 2*QueueLim) , the algorithm is switched into non-conservative mode (line 2). TCP Window Size is approximated here by the difference between the largest sequence number in the data queue and the sequence number in the ack queue. This is a reasonable approximation in our case since the wireless link is likely the bottleneck and most (if not all) of the queuing is done at the RNC. When operating in the non-conservative mode, no additional buffer space is reserved. This implies that there will be single loss the next time the TCP source increases it window size. At the detection of the packet loss, the algorithm again switches back to the *conservative* mode (line 3). This ensures that losses are of the single loss variety as long as the estimate of AckReleaseCount is conservative. Line 4 in the

algorithm results in a switch back into conservative mode whenever the data queue length goes to zero or whenever the TCP window size is halved. This handles the case when TCP reacts to losses elsewhere in the network and the Ack Regulator can go back to being conservative. Note that, if the TCP source is ECN capable, instead of switching to non-conservative mode, the Ack Regulator can simply mark the ECN bit to signal the source to reduce its congestion window, resulting in no packet loss.

```
1.Initialize AckReleaseCount=0;
2. On Enque of Ack/Deque of data packet:
   (after processing in Fig 7)
   AckReleaseCount+=AcksSent;
3. On Enque of data packet:
   if (AckReleaseCount>0)
      AckReleaseCount–;
4. On Deque of data packet:
   if (DataQueueLength==0)
      AckReleaseCount=0;
```

Figure 9: AckReleaseCount updates

We finally present the algorithm for updating the AckReleaseCount variable in Figure 9. Since AckReleaseCount estimates the expected number of data packets that are arriving and reserves buffer space for them, it is important to get an accurate estimate. An overestimate of AckReleaseCount would result in unnecessary reservation of buffers that won't be occupied, while an underestimate of AckReleaseCount can lead to buffer overflow loss(es) even in conservative mode due to inadequate reservation.

With the knowledge of the exact version of the TCP source and the round trip time from the RNC to the source, it is possible to compute an exact estimate of AckReleaseCount. However, since we would like to be agnostic to TCP version as far as possible and also be robust against varying round trip times on the wired network, our algorithm tries to maintain a conservative estimate of AckReleaseCount. Whenever we send acks back to the source, we update AckReleaseCount by that many acks (line 2). Likewise, whenever a data packet arrives into the RNC from the source, we decrement the variable while ensuring that it does not go below zero (line 3).

While maintaining a non-negative AckReleaseCount in this manner avoids underestimation, it also can result in unbounded growth of AckReleaseCount leading to significant overestimation as errors accumulate. For example, we increase AckReleaseCount whenever we send acks back to the source; however, if TCP is reducing its window size due to loss, we cannot expect any data packets in response to the acks being released. Thus, over time, AckReleaseCount can grow in an unbounded manner. In order to avoid this scenario, we reset AckReleaseCount to zero (line 4) whenever the data queue is empty. Thus, while this reset operation is necessary for synchronizing the real and estimated AckReleaseCount after a loss, it is not a conservative mechanism in general since a AckReleaseCount of zero implies that no buffer space is currently reserved for any incoming data packets that are unaccounted for. However, by doing the reset only when the data queue is empty, we significantly reduce the chance of the unaccounted data packets causing a buffer overflow loss. We discuss the impact of this estimation algorithm of AckReleaseCount in Section 6.6.

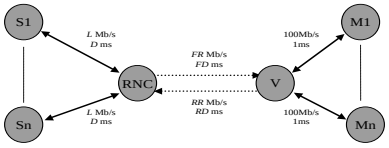

Figure 10: Simulation Topology

Finally, we assume that there is enough buffer space for the ack packets in the RNC. The maximum number of ack packets is the maximum window size achieved by the TCP flow (α*QueueLim in our algorithm). Ack packets do not have to be buffered as is, since storing the sequence numbers is sufficient (however, care should be taken to preserve duplicate ack sequence numbers as is). Thus, memory requirement for ack storage is very minimal.

6. SIMULATION RESULTS

In this section, we present detailed simulation results comparing the performance of TCP Reno and TCP Sack, in the presence and absence of the Ack Regulator. First, we study the effect of variable bandwidth and variable delay using different distributions on the throughput of a single long-lived TCP flow. Next, we present a model for 3G1X-EVDO (HDR) system (which exhibits both variable rate and variable delay), and evaluate the performance of a single TCP flow in the HDR environment. Then, we present the performance of multiple TCP flows sharing a single HDR wireless link. Finally, we briefly discuss the impact of different parameters affecting the behavior of Ack Regulator.

All simulations are performed using ns-2. The simulation topology used is shown in Figure 10. $S_i, i = 1..n$ corresponds to the set of TCP source nodes sending packets to a set of the TCP sink nodes $M_i, i = 1..n$. Each set of S_i, M_i nodes form a TCP pair. The RNC is connected to the M_i nodes through a V (virtual) node for simulation purposes. L, the bandwidth between S_i and the RNC, is set to 100Mb/s and D is set to 1ms except in cases where D is explicitly varied. The forward wireless channel is simulated as having rate FR and delay FD, and the reverse wireless channel has rate RR and delay RD.

Each simulation run lasts for 3600s (1hr) unless otherwise specified and all simulations use packet size of 1KB. TCP maximum window size is set to 500KB. Using such a large window size ensures that TCP is never window limited in all experiments except in cases where the window size is explicitly varied.

6.1 Variable Delay

In this section, the effect of delay variation is illustrated by varying FD, the forward link delay. Without modification, the use of a random link delay in the simulation will result in out-of-order packets since packet transmitted later with lower delay can overtake packets transmitted earlier with higher delay. However, since delay variability in our model is caused by factors that will not result in packet reordering (e.g. processing time variation) and RLP delivers packet in sequence, the simulation code is modified such that packets cannot reach the next hop until the packet transmitted earlier has arrived. This modification applies to all simulations with variable link delay.

Figure 11(a) shows throughput for a single TCP flow ($n =$

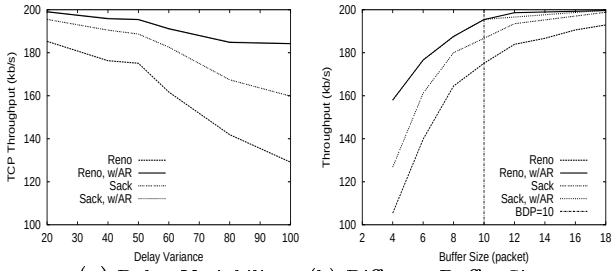

(a) Delay Variability (b) Different Buffer Size

Figure 11: Throughput with Variable Delay e(x)+400-x

1) for $FR = 200Kb/s$ and $RR = 64Kb/s$. FD has an exponential distribution with a mean that varies from 20ms to 100ms, and $RD = 400ms$ - mean(FD) so that average FD+RD is maintained at 400ms. The buffer size on the bottleneck link for each run is set to 10, the product of the mean throughput of (200Kb/s or 25pkt/s) and mean link delay (0.4s). This product will be referred to as the bandwidth-delay product (BDP) in later sections. Additional delay distributions like uniform, normal, lognormal, and Poisson were also experimented with. Since the results are similar, only plots for an exponential delay distribution are shown.

As expected, when the delay variation increases, throughput decreases for both TCP Reno and TCP Sack. By increasing the delay variance from 20 to 100, throughput of TCP Reno decreases by 30% and TCP Sack decreases by 19%. On the other hand, TCP Reno and TCP Sack flows which are Ack Regulated are much more robust and its throughput decreases by only 8%. Relatively to one another, *Ack Regulator performs up to* 43% *better than TCP Reno and* 19% *better than TCP Sack.* Another interesting result is that *Ack Regulator delivers the same throughput irrespective of whether the TCP source is Reno or Sack.* This is understandable given the fact that the Ack Regulator tries to ensure that only single buffer overflow loss occurs and in this regime, Reno and Sack are known to behave similarly. This property of Ack Regulator is extremely useful since for a flow to use TCP Sack, both the sender and receiver needs to be upgraded. Given that there are still significant number of web servers that have not yet been upgraded to TCP Sack [28], deployment of Ack Regulator would ensure excellent performance irrespective of the TCP version running.

Figure 11(b) shows how throughput varies with buffer size with the same set of parameters except for FD, which is now fixed with a mean of 50ms (exponentially distributed). Even with a very small buffer of 5 packets (0.5 BDP), Ack Regulator is able to maintain a throughput of over 80% of the maximum throughput of 200Kb/s. Thus, *Ack Regulator delivers robust throughput performance across different buffer sizes.* This property is very important in a varying rate and delay environment of a wireless system, since it is difficult to size the system with an optimal buffer size, given that the BDP also varies with time. For a buffer of 4 packets, the improvement over TCP Reno and Sack is about 50% and 24% respectively. As buffer size increases, the throughput difference decreases. With buffer size close to 20 packets (2 BDP), TCP Sack performs close to Ack Regulated flows, while improvement over TCP Reno is about 4%.

Finally, in Table 3, we list parameter values from the simulation for delay variance of 100. First, consider Reno and Reno with Ack Regulator (first two rows). It is clear that Ack Regulator is able to significantly reduce the conditional probability of multiple losses $p2$ and $p3$ as well as absolute number of loss events (TD and TO) resulting in substantial gains over Reno. Next, consider Sack and Sack with Ack Regulator (last two rows). In this case, we can see that Sack is very effective in eliminating most of the timeout occurrences. However, Ack Regulator is still able to reduce the absolute number of loss events by allowing the congestion window to grow to higher values (24 vs 19), resulting in throughput gains.

Item	Rate, Kb/s	TD	TO	p1	p2	p3	W_f
Reno	129	496	114	0.34	0.3	0.38	19
Reno+AR	184	302	8	0.98	0.0	0.02	24
Sack	160	434	4	0.99	0.0	0.01	19
Sack+AR	184	302	8	0.97	0.0	0.03	24

Table 3: Parameters from simulation for variance=100

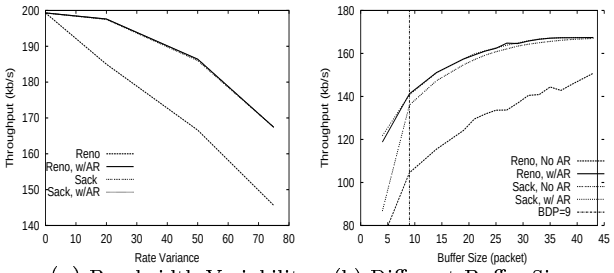

(a) Bandwidth Variability (b) Different Buffer Size

Figure 12: Throughput with Variable Bandwidth u(200,x)

6.2 Variable Bandwidth

In this section, we vary the link bandwidth, FR. Figure 12(a) shows throughput for a single TCP flow. FR is uniformly distributed with a mean of 200 Kb/s and the variance is varied from 20 to 75. $FD = 200ms$, $RR = 64Kb/s$ and $RD = 200ms$. The buffer size on the bottleneck link for each run is 10. Again, we have experimented with other bandwidth distributions, but, due to lack of space, only uniform distribution is shown. Note that, with variable rate, the maximum throughput achievable is different from the mean rate. For uniform distribution, a simple closed form formula for the throughput is simply $1/\int_a^b 1/xdx = 1/(\ln b - \ln a)$ where b is the maximum rate and a is the minimum rate.

When the rate variance increases, throughput of TCP Reno decreases as expected. Compared to TCP Reno, Ack Regulator improves the throughput by up to 15%. However, TCP Sack performs very well and has almost the same throughput as Ack Regulated flows. Based on the calculations for maximum throughput discussed before, it can be shown that all flows except Reno achieve maximum throughput. This shows that if rate variation is not large enough, TCP Sack is able to handle the variability. However, for very large rate variations (e.g. rate with lognormal distribution and a large variance), the performance of TCP Sack is worse than when Ack Regulator is present.

Figure 12(b) shows how the throughput varies with buffer size. Note that with a lower throughput, bandwidth delay

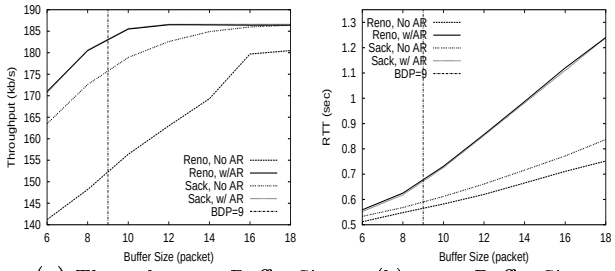

(a) Throughput vs Buffer Size (b) rtt vs Buffer Size

Figure 13: Throughput and rtt for u(200,50),350+e(50)

product is smaller than 10 packets. Again, Ack Regulated TCP flows perform particularly well when the buffer size is small. With buffer size of 5, the improvement over TCP Sack is 40%.

6.3 Variable Delay and Bandwidth

In this section, we vary both the bandwidth and delay of the wireless link. FR is uniformly distributed with a mean of 200 Kb/s and variance of 50, DR is exponentially distributed with a mean of 50ms, $RR = 64Kb/s$ and $RD = 350ms$. The maximum achievable throughput is 186.7 Kb/s. The BDP is therefore about 9 packets.

Figure 13(a) shows the throughput for a single TCP flow with the buffer size ranging from 7 to 20. The combination of variable rate and delay has a large negative impact on the performance of TCP Reno and it is only able to achieve 70% to 80% of the bandwidth of Ack Regulated flows when the buffer size is 6 packets. Even with a buffer size of 18 packets, the throughput difference is more than 5%. Throughput of TCP Sack is about 5% to 10% lower than Ack Regulator, until the buffer size reaches 18 packets (about 2 BDP).

One of the cost of using the Ack Regulator is the increase in average round trip time (rtt). The average rtt values for all 4 types of flows are shown in 13(b) for different buffer sizes. TCP Reno has the lowest rtt followed by TCP Sack and the rate of rtt increase with buffer size is comparable. With Ack Regulator, rtt increase is comparable with unregulated flows for buffer size less than 9 (1 BDP). For larger buffer sizes, since Ack Regulator uses $\alpha = 2$ times buffer size to regulate the acks in conservative mode, rtt increases faster with buffer size than regular TCP, where only the data packet buffer size contributes to rtt. For example, with buffer size of 9, Ack Regulated flows have a rtt 15% larger and with buffer size of 18, the rtt is 48% larger compared to TCP Sack. This effect can be controlled by varying the α parameter of the Ack Regulator.

6.4 Simulation with High Data Rate

High Data Rate (HDR) [6] is a Qualcomm proposed CDMA air interface standard (3G1x-EVDO) for supporting high speed asymmetrical data services. One of the main ideas behind HDR is the use of channel-state based scheduling which transmits packets to the user with the best signal-to-noise ratio. The actual rate available to the selected user depends on the current signal-to-noise ratio experienced by the user. The higher the ratio, the higher the rate available to the user. In addition, in order to provide some form of fairness, a Proportional Fair scheduler is used which provides long-term fairness to flows from different users. We use Qualcomm's Proportional Fair scheduler in our simulation

Rate(Kb/s)	Prob.	Rate(Kb/s)	Prob.
38.4	0.033	614.4	0.172
76.8	0.015	921.6	0.145
102.6	0.043	1228.8	0.260
153.6	0.023	1843.2	0.042
204.8	0.060	2457.6	0.011
307.2	0.168		

Table 4: HDR Data Rates for a one user system

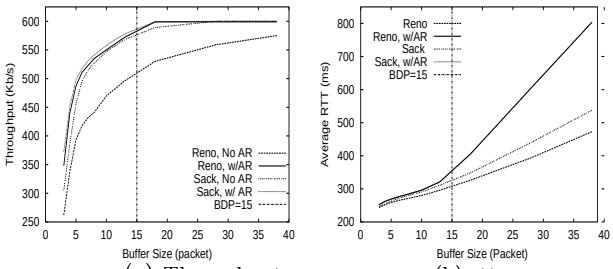

(a) Throughput (b) rtt

Figure 14: Throughput/rtt with HDR, Single Flow

with an averaging window of 1000 time slots, where each slot is 1.67 ms. While the HDR system results in higher raw throughput, the rate and delay variation seen is substantial.

In this section, we model a simplified HDR environment in ns-2, focusing on the layer 3 scheduling and packet fragmentation. The fading model for the wireless link used is based on Jake's Rayleigh fading channel model [25]. This gives us the instantaneous signal-to-noise ratio. Using Table 2 in [6] which lists the rate achievable for a given signal-to-noise ratio assuming a frame error rate of less than 1%, the achievable bandwidth distribution (with one user) for our simulation is shown in Table 4.

The simulation settings are as follows. FR is a variable that has a bandwidth distribution of Table 4, due to the variations of the fading conditions of the channel. Based on the guidelines from [26], FD is modeled as having a uniform distribution with mean 75ms and variance 30 and RD is modeled as having a uniform distribution with mean 125ms and variance 15. These are conservative estimates. We expect delay variations in actual systems to be higher (for example, note the ping latencies from our experiment in Section 3). The uplink in a HDR system is circuit-based and RR is set to be 64Kb/s.

Figure 14(a) shows how throughput for a single TCP flow varies with buffer size. Assuming an average bandwidth of 600Kb/s and a link delay of 200ms, BDP is 15 packets.

Again, the performance of TCP Reno flows that are Ack Regulated is significantly better than plain TCP Reno over the range of buffer size experimented, with improvements from 4% to 25%. TCP Sack flows also performs worse than Ack Regulated flows up to buffer size of 20. The improvement of Ack Regulator over TCP Sack ranges from 0.5% to 18%.

As mentioned earlier, one of the costs of using the Ack Regulator is increase in average rtt. The average rtt for all 4 types of flows are shown in 14(b) with buffer size varying from 5 to 40. The effect is similar to the rtt variation with buffer size seen in Section 6.3.

6.5 Multiple TCP Flows

In this simulation, the number of flows (n) sharing the

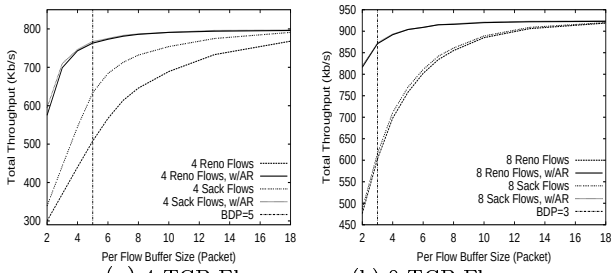

(a) 4 TCP Flows (b) 8 TCP Flows

Figure 15: Throughput with HDR, Multiple Flows

bottleneck link is increased to 4 and 8. Per-flow buffering is provided for each TCP flow. For 4 flows, using mean rate of 200Kb/s, 1KB packet and rtt of 0.2s, BDP is 5 packets per flow. For 8 flows, using mean rate of 120Kb/s, 1KB packet and rtt of 0.2s, BDP is 3 packets per flow.

As the number of TCP flows increases, the expected rate and delay variation seen by individual flows also increases. Thus, even though the total throughout of the system increases with more users due to channel-state based scheduling, the improvement is reduced by the channel variability.

Figure 15(a) shows the throughput for 4 TCP flows. The improvement of Ack Regulator over TCP Sack increases compared to the single TCP case. For example, the gain is 17% with per-flow buffer size of 5 (BDP). For Reno the gain is even greater. With per-flow buffer size of 5, the improvement is 33%. Similar result can also be observed for the case of 8 TCP flows as shown in Figure 15(b). For both TCP Reno and Sack, the gain is about 31% and 29% respectively for per-flow buffer size of 3. From the figure, it can seen that, *for TCP Sack and Reno to achieve close to maximum throughput without Ack Regulator, at least three times the buffer requirements of Ack Regulator is necessary* (buffer requirements for acks in the Ack Regulator is negligible compared to the 1KB packet buffer since only the sequence number needs to be stored for the acks). This not only increases the cost of the RNC, which needs to support thousands of active flows, it also has the undesirable side-effects of large rtt's that was noted in Section 3.

With multiple TCP flows, the issue of throughput fairness naturally arises. One way to quantify how bandwidth is shared among flows is to use the fairness index described in [27]. This index is computed as the ratio of the square of the total throughput to n times the square of the individual flow throughput. If all flows get the same allocation, then the fairness index is 1. As the differences in allocation increases, fairness decreases. A scheme which allocates bandwidth to only a few selected users has a fairness index near 0.

Computation of this index is performed for all multiple flows simulation and the index is greater than 0.99 in all cases. This result is expected since with per flow buffering, and proportional fair scheduling, the long term throughput of many TCP long-lived flows sharing the same link should be fair.

6.6 Parameters affecting the performance of Ack Regulator

Due to lack of space, we will only briefly present the results of varying parameters such as wired network latency and α.

As the network latency is varied from 20ms to 100ms, throughput decreases by 1.63% and 2.62% for Reno and Reno with Ack Regulator flows, respectively. Most of the decrease can be accounted for by the impact of increase in latency on TCP throughput. The result shows that the $AckReleaseCount$ estimation algorithm is effective and hence the Ack Regulator is able to reserve the appropriate amount of buffer for expected packet arrivals even with substantial wireline delay.

In another experiment, the parameter α in an Ack Regulated TCP flow is varied from 1 to 4. When α is increased from 1 to 3, the TCP flow is able to achieve its maximum throughput at a smaller buffer size. As α increases, the rtt also increases and when α is increased to 4, throughput decreases for larger buffer sizes (> 15). The decrease in throughput is caused by the accumulation of sufficiently large amount of duplicate acks that are sent to the TCP sender. A value of $\alpha = 2$ appears to be a good choice, balancing throughput and rtt for reasonable buffer sizes.

6.7 Summary of Results and Discussion

In this section, we first summarize the results from the simulation experiments and then briefly touch upon other issues.

We first started with experiments using a wireless link with variable delay. We showed that Ack Regulator delivers performance up to 43% better than TCP Reno and 19% better than TCP Sack when the buffer size was set to one BDP. We then examined the impact of a wireless link with variable rate. We saw that when the rate variance increases, throughput of TCP Reno decreases as expected. Compared to TCP Reno, Ack Regulator improves the throughput by up to 15%. However, TCP Sack performs very well and has almost the same throughput as Ack Regulated flows as long as the rate variation is not extremely large.

We next considered the impact of a wireless link with variable delay and variable rate. We found that this combination had a large negative impact on the performance of both TCP Reno and Sack (up to 22% and 10% improvement respectively for Ack Regulated flows). We then considered a specific wireless link standard called HDR which exhibits both variable delay and variable rate. The results were as expected, with Ack Regulator improving TCP Reno performance by 5% to 33% and TCP Sack by 0.5% to 24%. We then evaluated the impact of multiple TCP flows sharing the HDR link. The gains of Ack Regulator over normal TCP flows were even greater in this case (with 32% to 36% improvements) when the buffer size is set to one BDP.

In general, we showed that Ack Regulator delivers the same high throughput irrespective of whether the TCP flow is Reno or Sack. We further showed that Ack Regulator delivers robust throughput performance across different buffer sizes with the performance improvement of Ack Regulator increasing as buffer size is reduced.

We only considered TCP flows towards the mobile host (for downloading-type applications) since links like HDR are designed for such applications. In the case of TCP flows in the other direction (from the mobile host), Ack Regulator can be implemented, if necessary, at the mobile host to optimize the use of buffer on the wireless interface card.

Finally, Ack Regulator cannot be used if the flow uses end-to-end IPSEC. This is also true for all performance enhancing proxies. However, we believe that proxies for performance improvement are critical in current wireless networks. In order to allow for these proxies without compromising

security, a split security model can be adopted where the RNC, under the control of the network provider, becomes a trusted element. In this model, a VPN approach to security (say, using IPSEC) is used on the wireline network between the RNC and the correspondent host and 3G authentication and link-layer encryption mechanisms are used between the RNC and mobile host. This allows the RNC to support proxies such as the Ack Regulator to improve performance without compromising security.

7. CONCLUSION

In this paper, we comprehensively evaluated the impact of variable rate and variable delay on TCP performance. We first proposed a model to explain and predict TCP's throughput over a link with variable rate and delay. Our model was able to accurately (better than 90%) predict throughput of TCP flows even in the case of large delay and rate variation. Based on our TCP model, we proposed a network based solution called *Ack Regulator* to mitigate the effect of rate and delay variability. The performance of Ack Regulator was evaluated extensively using both general models for rate and delay variability as well as a simplified model of a 3^{rd} Generation high speed wireless data air interface. Ack Regulator was able to improve the performance of TCP Reno and TCP Sack by up to 40% without significantly increasingly the round trip time. We also showed that Ack Regulator delivers the same high throughput irrespective of whether the TCP source is Reno or Sack. Furthermore, Ack Regulator also delivered robust throughput performance across different buffer sizes. Given the difficulties in knowing in advance the achievable throughput and delay (and hence the correct BDP value), a scheme, like Ack Regulator, which works well for both large and small buffers is essential. In summary, Ack Regulator is an effective network-based solution that significantly improves TCP performance over wireless links with variable rate and delay.

Acknowledgements

The authors would like to thank Lijun Qian for providing the fading code used in the HDR simulation, Clement Lee and Girish Chandranmenon for providing the 3G1x trace and Sandy Thuel for comments on earlier versions of this paper.

8. REFERENCES

[1] E. Altman, K. Avrachenkov and C. Barakat, "A Stochastic Model of TCP/IP with Stationary Random Loss," in Proceedings of SIGCOMM 2000.

[2] F. Baccelli and D. Hong,"TCP is Max-Plus Linear," in Proceedings of SIGCOMM 2000.

[3] A. Bakre and B.R. Badrinath, "Handoff and System Support for Indirect TCP/IP," in proceedings of Second Usenix Symposium on Mobile and Location-Independent Computing, Apr 1995.

[4] H. Balakrishnan *et al.*, "Improving TCP/IP Performance over Wireless Networks," in proceedings of ACM Mobicom, Nov 1995.

[5] H. Balakrishnan, V.N. Padmanabhan, R.H. Katz, "The Effects of Asymmetry on TCP Performance," Proc. ACM/IEEE Mobicom, Sep. 1997.

[6] P. Bender *et al.*, "A Bandwidth Efficient High Speed Wireless Data Service for Nomadic Users," in IEEE Communications Magazine, Jul 2000.

[7] P. Bhagwat *at al*, "Enhancing Throughput over Wireless LANs Using Channel State Dependent Packet Scheduling," in Proc. IEEE INFOCOM'96.

[8] K. Brown and S.Singh, "M-TCP: TCP for Mobile Cellular Networks," ACM Computer Communications Review Vol. 27(5), 1997.

[9] A. Canton and T. Chahed, "End-to-end reliability in UMTS: TCP over ARQ," in proceedings of Globecomm 2001.

[10] TIA/EIA/cdma2000, "Mobile Station - Base Station Compatibility Standard for Dual-Mode Wideband Spread Spectrum Cellular Systems", Washington: Telecommunication Industry Association, 1999.

[11] G. Holland and N. H. Vaidya, "Analysis of TCP Performance over Mobile Ad Hoc Networks," in Proceedings of ACM Mobicom'99.

[12] H. Inamura *et al.*, "TCP over 2.5G and 3G Wireless Networks," draft-ietf-pilc-2.5g3g-07, Aug. 2002.

[13] F. Khafizov and M. Yavuz, "TCP over CDMA2000 networks," Internet Draft, draft-khafizov-pilc-cdma2000-00.txt.

[14] T. V. Lakshman and U. Madhow, "The Performance of Networks with High Bandwidth-delay Products and Random Loss," in IEEE/ACM Transactions on Networking, Jun. 1997.

[15] R. Ludwig *et al.*, "Multi-layer Tracing of TCP over a Reliable Wireless Link," in Proceedings of ACM SIGMETRICS 1999.

[16] Reiner Ludwig and Randy H. Katz "The Eifel Algorithm: Making TCP Robust Against Spurious Retransmissions," in ACM Computer Communications Review, Vol. 30, No. 1, 2000.

[17] V. Misra, W. Gong and D. Towsley, "Stochastic Differential Equation Modeling and Analysis of TCP Windowsize Behavior," in Proceedings of Performance'99.

[18] P. Narvaez and K.-Y. Siu, "New Techniques for Regulating TCP Flow over Heterogeneous Networks," in LCN'98.

[19] "Modeling TCP Throughput: a Simple Model and its Empirical Validation," in Proceedings of SIGCOMM 1998.

[20] S. Paul *et al.*, "An Asymmetric Link-Layer Protocol for Digital Cellular Communications," in proceedings of INFOCOM 1995.

[21] Third Generation Partnership Project, "RLC Protocol Specification (3G TS 25.322:)", 1999.

[22] TIA/EIA/IS-707-A-2.10, "Data Service Options for Spread Spectrum Systems: Radio Link Protocol Type 3", January 2000.

[23] S. Karandikar *et al.*, "TCP rate control," in ACM Computer Communication Review, Jan 2000.

[24] 3G Partnership Project, Release 99.

[25] "Microwave mobile communications," edited by W. C. Jakes, Wiley, 1974.

[26] "Delays in the HDR System," QUALCOMM, Jun. 2000.

[27] R. Jain, "The Art of Computer Systems Performance Analysis," Wiley, 1991.

[28] J. Padhye and S. Floyd, "On Inferring TCP Behavior," in Proceedings of SIGCOMM'2001.