

# The Impact of UMTS/WCDMA Channel Round-Trip Time Variations on TCP Vegas Throughput

Anthony Lo<sup>†</sup>, Geert Heijenk<sup>‡</sup> and Ignas Niemegeers<sup>†</sup>

<sup>†</sup>Delft University of Technology

P O Box 5031, 2600 GA Delft

The Netherlands

e-mail: [A.Lo, I.Niemegeers]@ewi.tudelft.nl

<sup>‡</sup>University of Twente

P O Box 217, 7500 AE Enschede

The Netherlands

e-mail: Geert.Heijenk@utwente.nl

**Abstract**—Universal Mobile Telecommunications System (UMTS) is a third-generation cellular network that enables high-speed mobile Internet access. This paper evaluates the performance of the Transmission Control Protocol (TCP) Vegas and compares its performance with TCP Reno in a UMTS environment. Bulk data transfer was considered in the simulation with varying radio channel conditions. Data losses were assumed to be only caused by the radio channel. Simulation results show that the performance of Vegas is unexpectedly worse than Reno even though the data losses were completely recovered by the UMTS radio link control layer. This has led us to conduct a detailed investigation on the behavior of Vegas in order to identify the cause of performance degradation in Vegas. The poor performance of Vegas is attributed to the UMTS radio interface characteristics which characterize by large and highly variable TCP round-trip times. The round-trip time variations activated Vegas' congestion avoidance and the sender would interpret the round-trip time variation as a sign of network congestion. Consequently, the sender shrinks its window size which reduces the transmission rate even though no TCP packet losses occur. Furthermore, a sudden increase in the instantaneous round-trip time can trigger spurious timeouts at the TCP sender using Vegas which performs unnecessary retransmissions. Spurious timeouts can lead to significant throughput reduction. Reno, on the other hand, does not show any abnormality and delivers the expected performance.

**Keywords**—UMTS; WCDMA; TCP Vegas

## I. INTRODUCTION

Universal Mobile Telecommunications System (UMTS) [1,2] is a third-generation cellular mobile network where the radio interface is based on code division multiple access, known as Wideband Code Division Multiple Access (WCDMA). To date, a number of mobile operators in Europe and Asia have launched their UMTS commercial service and some plan to roll out the UMTS networks. In addition to the legacy voice service, UMTS enables mobile users access to the Internet in a seamless fashion (i.e., always on) at data rates up to 2 Mb/s in indoor/small-cell environments, and wide-area coverage of up to 384 kb/s.

Today, the Internet is the most popular and widely used packet-switched network that supports applications like File Transfer Protocol (FTP), Email, etc. These Internet applications rely on two commonly used protocols, namely,

Transmission Control Protocol and Internet Protocol (TCP/IP) [3], to reliably transport data across heterogeneous networks. IP is concerned with routing data from source to destination host through one or more networks connected by routers, while TCP provides reliable end-to-end data transfer and congestion control. TCP Reno [3] is the most extensively used variant of TCP, while TCP Vegas [4] is a newer variant with improved congestion avoidance and retransmission mechanisms. Unlike Reno, Vegas constantly tries to detect congestion in the network before packet loss occurs and lower the rate linearly when sign of congestion is detected. On the contrary, Reno only reacts when packet losses are detected.

The performance of TCP over wireless networks has been extensively studied [5-8]. All these studies show that TCP performance is significantly degraded since TCP interprets packet losses due to the radio channel as signs of network congestion, which resulting in sender throttling and causes significant throughput reduction. Various solutions were proposed in the literature to combat TCP performance degradation, and in general, can be classified into three major categories: link-layer [7-10], split-connection [11] and proxy [6]. However, all these studies including the proposed solutions were purely targeted at Reno rather than Vegas. We anticipated that Vegas will be widely spread in the near future since it has been supported by some operating system (e.g. [12]) in addition to Reno. But, it is currently unclear if Vegas will obtain the same performance gain as Reno using the proposed solutions. The solution employed by UMTS falls under the link-layer category. The paper aims at assessing the performance of TCP Vegas and compares its performance with Reno in a UMTS environment, in particular, if Vegas can achieve the same performance improvement as Reno over the UMTS radio interface. We employ a simulation-based approach to analyze the performance of TCP Vegas and Reno over UMTS for FTP traffic with varying channel conditions.

## II. UNIVERSAL MOBILE TELECOMMUNICATIONS SYSTEM (UMTS)

### A. System Architecture

Fig. 1 shows a simplified architecture of UMTS for packet-switched operation [2,13], which consists of one or several User Equipments (UEs), the UMTS Terrestrial Radio Access

Network (UTRAN) and the core network. The UTRAN is composed of several Node Bs connected to a Radio Network Controller (RNC). The core network, which is the backbone of UMTS, comprises the Serving GPRS Support Node (SGSN) and the Gateway GPRS Support Node (GGSN). The SGSNs route packets to and from UTRAN, while GGSNs interface with external IP networks. UE, which is a mobile station, is connected to Node B over the UMTS radio interface.




Figure 1. UMTS Reference Architecture.

### B. Protocol Architecture

Fig. 2 depicts the UMTS protocol architecture for the transmission of user data which is generated by TCP-based applications. The applications as well as the TCP/IP protocol suite are located at the end-nodes, namely, the UE and a host.

The Packet Data Convergence Protocol (PDCP) provides header compression functionality. The Radio Link Control (RLC) layer can operate in three different modes: acknowledged, unacknowledged and transparent. The acknowledged mode provides reliable data transfer over the error-prone radio interface. Both the unacknowledged and transparent modes do not guarantee data delivery. The Medium Access Control (MAC) layer can operate in either dedicated or common mode. In the dedicated mode, dedicated physical channels are allocated and used exclusively by one user (or UE), whereas in the common mode, users share common physical channels for transmitting and receiving data. The Physical (PHY) layer contains, besides all radio frequency functionality, spreading, and the signal processing including, power control, forward error-correction and interleaving.

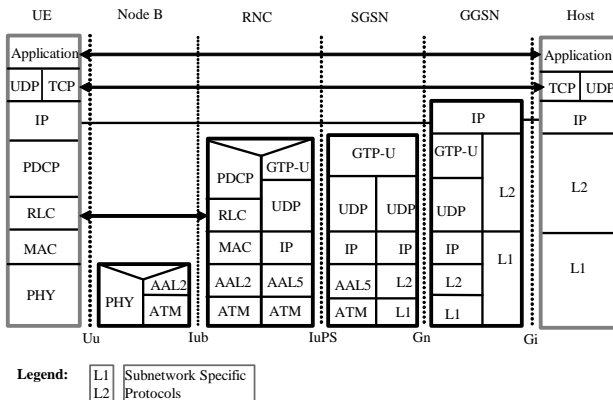


Figure 2. UMTS Protocol Architecture – U-Plane.

### III. TCP VEGAS

Both TCP Vegas and Reno perceive packet losses as a sign of network congestion. Reno reacts to network congestion when packet losses are detected via timeout or three duplicate acknowledgements. On the other hand, Vegas continuously monitors the state of the network and increment or decrement the current window size in order to prevent packet drops due to buffer overflowing in the intermediate routers. Vegas detects signs of incipient congestion by comparing the expected throughput to the measured throughput, which is given as follows [4]:

$$\Delta = (\text{expected} - \text{actual}) \times RTT_{base} \quad (1)$$

where  $\text{expected} = \text{windowSize}/RTT_{base}$  and  $\text{actual} = \text{windowSize}/RTT$ .  $\text{windowSize}$  is the current window size which is the number of segments in transit;  $RTT_{base}$  is the minimum of all the instantaneous round-trip times; and  $RTT$  is the average round-trip time measured for each individual segment transmitted in the  $\text{windowSize}$ . Vegas also defines two thresholds  $\alpha$  and  $\beta$ , which are normally set to 1 and 3, respectively. When  $\Delta < \alpha$ , Vegas increases the congestion window linearly in the next round-trip time; and when  $\Delta > \beta$ , Vegas decreases the congestion window linearly in the next round-trip time. The congestion window is unchanged when  $\alpha < \Delta < \beta$ .

In the case of Reno, packet losses are detected via the receipt of three duplicate acknowledgements or retransmission timeout expiration. The latter resets the congestion window size to one segment, while the former reduces the congestion window by one half of the current window size.

### IV. SIMULATION MODELS

In order to analyze the performance of Vegas and Reno over UMTS, network-level simulations were carried out using *ns-2* [14], which is an event-driven simulator. Several extensions were made to this simulator for modeling UMTS. The extensions were developed within the framework of the IST SEACORN project [15]. With the extensions, instances of UMTS nodes, viz., UE, Node B and RNC can be created. The model used for simulation analysis is illustrated in Fig. 3. The model is based on the system architecture discussed in the previous section (see Fig. 1). UE, Node B, RNC and host are modeled according to the aforementioned protocol stack illustrated in Fig. 2. The TCP/IP protocol stack of *ns-2* was used.

Since the primary aim of the simulation was to investigate the impact of the radio interface on end-to-end TCP performance, we assume that no packet losses, errors or congestion on either the Internet or the UMTS core network. Hence, the TCP performance is solely attributed to the radio interface. The links between two nodes are labeled with their bit rate (in bits per second) and delay (in seconds). Each link capacity was chosen so that the radio channel is the connection bottleneck. Hence, the TCP performance is solely attributed to the UMTS radio interface. Consequently, the functionality of SGSN and GGSN was abstracted out and modeled as traditional *ns* nodes since generally they are wired nodes and, in many ways, mimic the behavior of IP router. Currently, no

header compression technique is supported in the PDCP layer. In the following subsections, the UMTS model is described in detail.




Figure 3. Top-level Simulation Model.

#### A. RLC Model

The RLC model supports both acknowledged and unacknowledged modes. For TCP-based applications, the acknowledged mode was used in the simulation since the acknowledged mode was designed to hide losses due to radio channels from TCP. The retransmission strategy adopted by the acknowledged mode is the Selective-Repeat ARQ (Automatic Repeat reQuest) scheme. With Selective-Repeat ARQ, the only RLC blocks retransmitted are those that receive a negative acknowledgement. An RLC block consists of a header and a payload which carries higher layer data. A status message is used by the receiver for notifying loss or corruption of an RLC block. The status message is in bitmap format. That is,  $bit_j$  indicates whether the  $j$ th RLC block has been correctly received or not. The frequency of sending STATUS messages is not specified in the standard [16]. However, several mechanisms are defined, which can trigger a STATUS message. Either the sender or the receiver can trigger the STATUS message. Table I and Table II list the triggering mechanisms for sender and receiver, respectively. It is important to note that not all the triggering mechanisms are needed for the Selective-Repeat ARQ to operate. However, a combination of triggering mechanisms, which deliver optimum performance, is sought.

The advantage of receiver-initiated mechanisms is that the receiver has direct information about missing blocks. For the sender-initiated mechanisms, the sender has first to request a STATUS message by enabling the POLL flag in the RLC block and wait for a reply, which has longer turn around time. Therefore, receiver-initiated mechanisms are preferred. Nevertheless, sender-initiated mechanisms are required to prevent deadlocks and stall conditions. Periodic mechanisms might be more robust compared to others but may result in too frequent STATUS message. In addition, a timer is required at the sender and receiver for proper operation of the triggering mechanisms. At the sender, the timer is called *poll* timer, which is started when a request for status messages is sent to the receiver. If the status message from the receiver does not arrive before the timer expires, the sender repeats the same procedure again. The receiver is equipped with a timer called *status prohibit* timer, which controls the time interval between status messages if triggered consecutively. If the interval is too short, then bandwidth is wasted. On the other hand, if the interval is too long, bandwidth is preserved, but delay increases. The selected triggering mechanisms for the RLC model are labeled by the rows in dark grey.

TABLE I. SENDER-INITIATED MECHANISMS

Trigger	Explanation
Last Block in buffer or retransmission buffer	status report is requested by enabling the poll flag in the RLC header
Every $m$ blocks	poll flag is enabled for every $m$ blocks
Every $n$ service data units	poll flag is enabled for every $n$ service data units
Utilization of Send Window	poll flag is enabled when the Send Window is $x\%$ full
Periodic Poll	poll is generated periodically based on a timer

TABLE II. RECEIVER-INITIATED MECHANISMS

Trigger	Explanation
Detection of missing blocks	Status message is generated once a gap is detected in the RLC sequence number
Periodic status	Status message is sent periodically based on a timer
Estimated block Counter	Status message is generated if not all the retransmitted blocks are received within an estimated period

#### B. MAC Model

The MAC model implemented the dedicated mode. It requests the number of blocks buffered at the RLC layer, which are ready for transmission, and submits to the PHY layer as transport blocks. In this case, each transport block corresponds to an RLC block since no MAC header is required in the dedicated mode as depicted in Fig. 4. The frequency in which the PHY layer can accept transport blocks from MAC is defined by the Transmission Time Interval (TTI). In the UMTS standard, the values of TTI are 10 ms, 20 ms, 40 ms and 80 ms.

#### C. PHY Model

The PHY model is responsible for transmitting transport blocks over the physical channels. For the MAC dedicated mode, the transport blocks are sent over the Dedicated Physical Channel (DPCH). The channel bit rates and the TTI associated with the channel considered in the simulation are shown in Table III. Note that the bit rates exclude the RLC headers. The other PHY layer functionality is not implemented in the model. Since the PHY layer passes the transport block to the MAC layer together with the error indication from the Cyclic Redundancy Check (CRC), the output of the PHY layer can be characterized by the overall probability of transport block error – also called Transport Block Error Rate (TBLER) in this paper. Thus, an error model based on uniform distribution of transport block errors, was used in the simulation. It is valid to assume that the erroneous transport blocks perceived by the RLC is independent and uniformly distributed as a result of interleaving and forward error-correction mechanisms used by the PHY layer. The TBLER, in the range from 0 to 30%, was considered in the simulation.

The transmission of an IP packet over the radio interface is illustrated in Fig. 4. The RLC entity receives a PDCP packet which comprises an IP packet of 552 Bytes or an acknowledgement of 40 Bytes, and additionally the PDCP header of 1 Byte. This PDCP packet is segmented into multiple RLC blocks of fixed sizes. Each of these blocks fits into a transport block in which a CRC is attached at the PHY layer. In the simulation, the RLC header and the payload size was set to

2 Bytes and 40 Bytes, respectively. For this RLC payload size and, for instance, a bit rate of 384 kb/s, twelve transport blocks are transmitted within one TTI of 10 ms. The other simulation parameters are summarized in Table III.

TABLE III. SIMULATION PARAMETERS

Application	File Transfer Protocol (FTP)				
TCP	TCP Variants	Vegas and Reno			
	Window Size (Segments)	64			
	Maximum Segment Size (Bytes)	512			
	TCP Header Size (Bytes)	20			
IP	IP Header Size (Bytes)	20			
	IP Packet Loss Rate in the Internet	0%			
PDCP	TCP/IP Header compression	No			
RLC	RLC Mode	Acknowledged Mode with In-sequence delivery			
	Window Size (Blocks)	4096			
	Payload Size (Bits)	320			
	RLC Header (Bits)	16			
MAC	MAC Header (Bits)	0			
	MAC Multiplexing	Not required for DPCH			
PHY	Physical Channel Type	DPCH			
		Uplink		Downlink	
		Bit Rate (kb/s)	TTI (ms)	Bit Rate (kb/s)	TTI (ms)
		64	20	384	10
				128	20
				64	20
		Transport Block Size (Bits)	336		
	Transport BLER	0 – 30%			
	Error Model	Uniform Distribution			

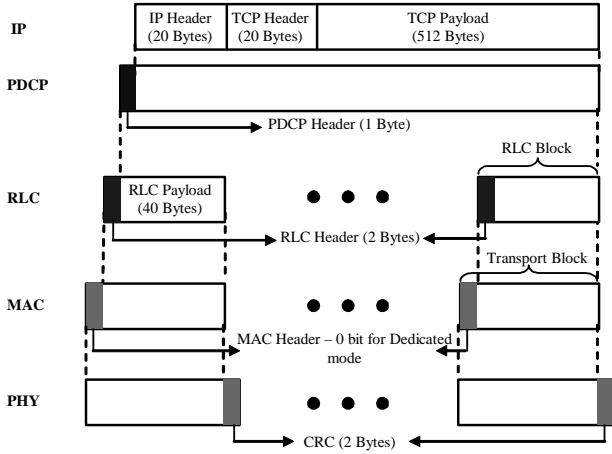


Figure 4. IP Packet Data Transfer.

## V. SIMULATION RESULTS

### A. TCP Throughput

End-to-end TCP *throughput* is used as the performance measure. The throughput (in bits per second) is defined as the amount of successfully received TCP segments by the receiver within the simulation duration. The TCP throughput was

obtained using a single FTP session between a UE and a host. Data is transferred from the host to the UE. That means, the only higher layer data going in the uplink channel are TCP acknowledgements. The FTP session was run for 1000s. First, the simulation was run using TCP Vegas, and then, the same set of simulation was repeated for Reno. Fig. 5 plots the TCP throughput as a function of TBLE for Vegas and Reno for the three different channel bit rates. For both Reno and Vegas, the throughput obtained for the different radio channel bit rates is normalized to the corresponding downlink channel bit rate.

Under ideal radio channel (i.e., 0% TBLE), both Reno and Vegas attain the maximum throughput of approximately 96% of the radio channel capacity, which is the maximum achievable throughput, after discounting the overhead of RLC control messages, PDCP and IP headers. As observed in Fig. 5, the performance of Vegas is worse than Reno and rapidly deteriorating when the TBLE increases even though the simulation traces showed that the RLC acknowledged mode successfully delivered every transmitted TCP segment to the receiver. In addition, Vegas performance is degrading for higher channel bit rates. In the case of 384 kb/s, when the TBLE is 30%, Vegas' throughput drops to 5%, which is 18% lesser than the 64 kb/s case and 60% lower than Reno. Conversely, Reno obtains a similar throughput over the different channel bit rates. The cause of poor TCP Vegas performance is analyzed and explained in detail in the next section, and for ease of explanation, the results obtained for the 384 kb/s case was used for illustration purposes since it is the worst case scenario.




Figure 5. Throughput versus Transport Block Error Rate.

### B. TCP Round-Trip Time Variability

Fig. 6 plots the Vegas and Reno congestion window sizes in segments over the simulation duration for 5% TBLE. The poor performance of Vegas is clearly evidenced by the small window size, while Reno's congestion window grows exponentially until the maximum window size is reached, and remains at this size throughout the simulation. The maximum achievable window size of the TCP connection for both Vegas and Reno is 64 segments, which also corresponds to the bandwidth-delay product. The congestion window size of Vegas, however, only managed to reach a maximum window

size of 7 segments for a short period (appeared as peaks in the Vegas' curve of Fig. 6) and then dropped to 6 segments. The first peak occurred at approximately 430s. The Vegas' congestion window exhibits instability and oscillates about the mean window size of 5 segments. This mean window size is relatively small as compared with Reno, which explains for the low throughput achieved by Vegas. The peculiar behavior of Vegas is due to the congestion avoidance which uses the TCP round-trip times to adjust the window size. The round-trip time for each TCP segment was obtained from the simulation traces for 5% TBLER. In total, there were approximately 26,000 round-trip time samples produced from the simulation traces and plotted in Fig. 7. The y-axis shows the round-trip time expressed in seconds for each individual segment and the x-axis shows the time in seconds when each round-trip time was recorded. The round-trip time is defined as the amount of time from the segment is transmitted until an acknowledgement for the corresponding segment is received by the TCP sender. The plot indicates that the TCP round-trip time exhibits relatively large variation between consecutive round-trip times with the maximum and minimum values of 0.56s and 0.12s, respectively. The minimum round-trip time is the basic time to send a TCP segment and receive the corresponding acknowledgement in the absence of transmission errors over the radio channel. Therefore, the maximum variation in the round-trip time is 0.44s. As shown in Fig. 7, the round-trip time values are quantized at discrete levels (with a step size of 20 ms) because the round-trip time variation is solely due to the transmission delay over the radio interface. The delays in the core network and the Internet are fixed in all the simulations. As a result, the instantaneous round-trip time is a multiple of TTI and can be expressed as:

$$RTT = RTT_{base} + n \times TTI, \quad n = 0, 1, 2, \dots \quad (2)$$

In the case of asymmetric channel, the maximum TTI is taken into consideration. For instance, in the case of 384 kb/s, the TTI is set to 20 ms in (2).

Since the delay in the Internet and the UMTS core network is fixed in all the simulations (see Fig. 3), the large round-trip time variation is due to the in-sequence delivery and retransmission strategy at the RLC layer. Without in-sequence delivery, the TCP receiver can generate spurious duplicate acknowledgements due to out of order segments. The duplicate acknowledgement triggers the fast retransmit of the TCP sender, which leads to redundant retransmission. As mentioned, the retransmission strategy of RLC is selective-repeat ARQ which is used to retransmit erroneous RLC blocks that composed a TCP segment or an acknowledgement. Note that, the uplink and downlink channels are subject to equal block error rates. A large variation in round-trip time causes instability and can result in a destructive effect to Vegas which becomes unstable once the  $RTT_{base}$  in (1) is locked to the minimum round-trip time. For the 5% TBLER case, the  $RTT_{base}$  is 0.12s. Therefore, if  $windowSize$  is equal to 6 segments, and if the average round-trip time,  $RTT \geq 0.3s$ , the criterion  $\Delta > \beta$  is satisfied, which linearly decreases the congestion window as observed in Fig. 6. For  $windowSize = 5$  and  $RTT$  at the mean value of 0.2s, the criterion,  $\alpha < \Delta < \beta$ , is satisfied and the

congestion window is unchanged, which correlates to the mean window size of 5 segments.

A similar phenomenon was also observed for higher TBLERs. The round-trip time variation, however, is even greater as the TBLER increases, which means, the likelihood of  $\Delta$  exceeding  $\beta$  is higher than small TBLERs. In the case of 30% TBLER, the maximum variation in round-trip time is 1.44s as compared with 0.44s for 5% TBLER. In addition to greater round-trip time variation, spurious timeouts were observed for TBLERs equal to or greater than 20%, which causes significant throughput degradation. Fig. 8 presents Vegas and Reno's congestion window size over simulation duration for 30% TBLER. Note that, the behavior of Reno is similar to the 5% TBLER case, where the maximum window size is reached. As for Vegas, the congestion window size is relatively smaller than the 5% TBLER case. The congestion window opened up to 4 segments during slow start, and then dropped to 3 segments at around 103.6s as depicted in Fig. 8, which is due to the congestion avoidance. Unlike Reno, Vegas' retransmission timeout is accurately reflecting the measured round-trip time of TCP segments. As a result, a sudden increase of the instantaneous round-trip delay beyond the sender's retransmission timeout value causes spurious timeouts.

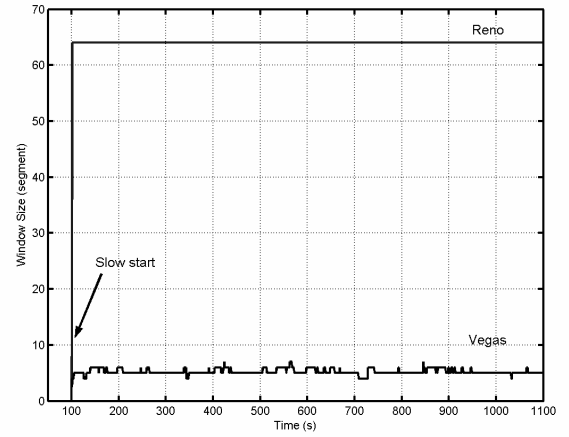


Figure 6. TCP Congestion Window versus Time for 5% TBLER.




Figure 7. TCP Vegas Round-Trip Time for 5% TBLER.

As shown in Fig. 8, several spurious timeouts occurred (marked by arrows) throughout the simulation, which prevents the congestion window from inflating. As a result, the congestion window size remained at 2 segments for more than fifty percent of the TCP connection time. Spurious timeouts can also cause redundant retransmissions which have contributed to the poor performance of Vegas. Fig. 9 shows the TCP trace for a spurious timeout event for 30% TBLER. The figure illustrates that the retransmission timeout value used by Vegas is vulnerable to the increased round-trip time and triggering unnecessary retransmissions.

In summary, Vegas would interpret an increase in round-trip delay as a sign of congestion in the network, and consequently, decrease the window size. However, this is not the desired behavior, but sender should perform the opposite.

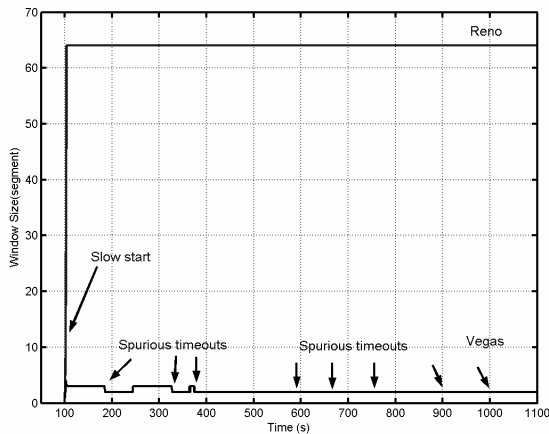


Figure 8. TCP Congestion Window versus Time for 30% TBLER.




Figure 9. TCP Vegas Trace shows Spurious Timeouts and Redundant Retransmissions for 30% TBLER.

## VI. CONCLUSIONS

The paper has evaluated and compared the performance of TCP Vegas with Reno over UMTS dedicated channels with different bit rates for bulk data transfers. Throughput was used as a performance measure. Throughput simulation results show that the performance of Vegas is worse than Reno under various transport block error rates. As a matter of fact, Vegas'

throughput collapses for transport block error rates greater than 30% even though the RLC acknowledged mode successfully delivered every transmitted TCP segment. Conversely, Reno's behavior does not show any abnormality and it achieves the expected throughput over the different transport block error rates. The interaction between Vegas and the UMTS radio interface protocols, in particular, the radio link control layer was examined in detail. The UMTS radio interface exhibits large and highly variable delay, which in turn, resulted in fluctuating TCP round-trip times. Vegas differs from Reno in the sense that it uses TCP round-trip times to detect congestion. Hence, the round-trip time variation is perceived as a sign of congestion and causes the sender to shrink its window size, which has a detrimental effect on the performance. Moreover, a sudden increase in instantaneous round-trip time can trigger spurious timeouts at the sender running Vegas which performs unnecessary retransmissions. Spurious timeouts occur at high transport block error rates. The UMTS radio interface characteristics pose performance issues to Vegas due to the new congestion avoidance and retransmission features. Reno is less intelligent but its congestion avoidance and retransmission mechanisms are more robust as compared with Vegas.

## REFERENCES

- [1] 3GPP, available from <http://www.3gpp.org>.
- [2] H. Kaaranen, S. Naghian, L. Laitinen, A. Ahtinen, and V. Niemi., UMTS Networks: Architecture, Mobility and Services, New York: John-Wiley & Sons, 2001.
- [3] W. Stevens, TCP/IP Illustrated, Vol 1, Reading, MA: Addison Wesley, 1994.
- [4] L. S. Brakmo and L. L. Peterson, "TCP Vegas: End to End Congestion Avoidance on a Global Internet", IEEE Journal on Selected Areas in Communications, vol. 13, no. 8, 1995.
- [5] R. Caceres and L. Iftode, "Improving the Performance of Reliable Transport Protocols in Mobile Computing Environments", IEEE Journal on Selected Areas in Communications, vol. 13, no. 5, 1995.
- [6] H. Balakrishnan, S. Seshan, and R. H. Katz, "Improving Reliable Transport and Handoff Performance in Cellular Wireless Networks", ACM Wireless Networks, vol. 1, no. 4, 1995.
- [7] N. M. Chaskar, T. V. Lakshman, and U. Madhow, "TCP over Wireless with Link Level Error Control: Analysis and Design Methodology", IEEE/ACM Transaction on Networking, vol. 7, no. 5, 1999.
- [8] Y. Bai, A. T. Ogielski, and G. Wu, "Interactions of TCP and Radio Link ARQ Protocol", Proceedings of the IEEE Vehicular Technology Conference (VTC'99 - Fall), Amsterdam, The Netherlands, 1999.
- [9] C. Parsa and J. J. Garcia-Luna-Aceves, "Improving TCP Performance over Wireless Networks at the Link Layer", Mobile Networks and Applications, vol. 5, no. 1, 2000.
- [10] E. Ayanoglu, S. Paul, T. F. LaPorta, K. K. Sabnani, and R. D. Gitlin, "AIRMAIL: A Link-Layer Protocol for Wireless Networks", Wireless Networks, vol. 1, no. 1, 1995.
- [11] A. Bakre and B. R. Badrinath, "Implementation and Performance Evaluation of Indirect-TCP", IEEE Transaction on Computers, vol. 46, no. 3, 1997.
- [12] N. Cardwell, "A TCP Vegas Implementation for Linux", <http://flophouse.com/~neal/uw/linux-vegas>
- [13] 3GPP, "Network architecture", TS 23.002, 2000.
- [14] K. Fall and K. Varadhan, "The ns Manual", <http://www.isi.edu/nsnam/ns/ns-documentation.html>.
- [15] SEACORN Home Page: <http://seacorn.ptinovacao.pt/>
- [16] 3GPP, "Radio Link control (RLC) Protocol Specification", TS 25.322, 2000.