Preliminary Analysis of the TCP Behavior in 802.16 Networks

Jesus A. Perez[1], Benoit Donnet[2], and Olivier Bonaventure[2]

[1] ITESM – Campus Cuernavaca, Department of Electronics, Temixco, Mexico
[2] Université Catholique de Louvain, CSE Department, Louvain-la-Neuve, Belgium

Abstract. Currently, broadband wireless access is gaining a great deal of interest from the networking research community. Particularly, the recently standardized WiMAX presents interesting perspectives, notably due to its capacity to offer consistent bandwidth and therefore consistent QoS. However, the behavior of network protocols, such as TCP, has not been studied in detail in a WiMAX environment. This is a problem that could slow down the widespread deployment of WiMAX. In this paper, we present preliminary results of the performance of TCP in a pre-WiMAX network. We are interested in the RTT and the relationship between the delay and packet loss rate. We find that TCP presents an acceptable cyclic behavior but with a high percentage of packet loss, 6.2%. We further notice some burst packet losses that vary in size and duration. Some of these bursts correspond to a packet loss rate of 100% during periods of up to one second.

Key words: WiMAX, TCP, measurement

1 Introduction

During the past few years several wireless technologies have been deployed. In particular, 802.11 [1] has been positioned as the de-facto standard for WLANs and has been widely deployed throughout the world. The third generation of mobile telephony is also growing everyday. However in the last few years, the research and business communities have increased their interest in the new broadband wireless technologies. In particular WiMAX [2] has become the most important since it is the moniker used for the IEEE 802.16 wireless interface specification.

WiMAX is a standard-based technology which will serve as a wireless extension or alternative to cable and DSL for broadband access. Particularly for end users in rural, sparsely populated areas or in areas where laying cable is difficult or expensive [3]. WiMAX will provide a new broadband access path to Internet. But companies and communities along will benefit from WiMAX as well, if they require mobile networks that cover a wider area than Wi-Fi.

WiMAX is built focusing on data services instead of voice services. It was designed predominantly for home and business users who do not have fixed-line access to broadband Internet. Since WiMAX has been developed primarily for the transmission of larger data volumes at high speeds, it provides a logical supplement to UMTS which is a technology that offers voice and multimedia services even when users are moving at high speed. Similarly, Wi-Fi and WiMAX are complementary technologies since WiMAX is able to cover wider areas than WiFi.

Most of the wireless technologies have incorporated TCP/IP into their protocol stacks. However, TCP was designed for wired networks. Its sliding window and congestion avoidance mechanisms were designed to avoid routers congestion.

In the past years, TCP has been extensively studied and extensions, such as BIC TCP [4] or Fast TCP [5], have been proposed to improve the congestion control mechanisms, especially in high-speed long distance networks. The TCP behavior has also been evaluated in wireless, i.e., 802.11, environment. It resulted, from this evaluation, several extensions to TCP [6], [7], [8], [9] that make TCP more robust to wireless specific conditions. Indeed, while network congestion is an acceptable assumption for packet losses in many networks, wireless networks might encounter two additional causes of packet losses [10]. The first one is the *random packet loss* that manifests itself through bit corruption. Such packets are discarded by the routers or the end-hosts. Second, a *disconnection packet loss* might occur when a mobile host completely disconnects from the wireless network. Note that this latter loss is a characteristic of infrastructured networks. Finally, the weather conditions might also affect the signal link quality [11].

For all of these reasons and considering that nowadays WiMAX is the most important and promising technology for broadband wireless access, it is very important to study TCP real behavior in one of the pre-WiMAX implementations.

This paper represents the first step towards a better understanding of TCP behavior in a WiMAX environment. Using a WiMAX commercial connection in Belgium, we present preliminary results of TCP measurements. To the best of our knowledge, this is the first paper that presents real measurements in a pre-WiMAX network and discusses how TCP behaves in such a network.

The remainder of this paper is organized as follows: Sec. 2 describes the basic features of WiMAX and Expedience technologies; Sec. 3 presents our preliminary TCP evaluation; Sec. 4 discusses some related work to our research; finally, Sec. 5 summarizes the preliminary results we have obtained from our tests.

2 WiMAX and Expedience Technology

The WiMax standard supports adaptive modulation, effectively balancing different data rates and link quality. The modulation method may be adjusted almost instantaneously for optimum data transfer. WiMAX is able to dynamically shift modulations from 64-QAM to QPSK via 16-QAM, displaying its ability to overcome Quality of Service (QoS) issues with dynamic bandwidth allocation over the distance between the *Base Station* (BS) and the *Subscriber Station* (SS).

Modulation schemes ensure that a quality signal is delivered over the distance by decreasing throughput. So throughput declines with distance, for example 12 Mbps to 2 miles, 6 Mbps to 3 miles, 3 Mbps to 4 miles Non-line-of-sight (NLOS). Generally, the Signal-to-Noise Ratio (SNR) requirements of an environment determine the modulation method to be used in the environment.

WiMAX incorporates a number of time-proven mechanisms to ensure good QoS. Most notable are Time Division Duplex (TDD), Frequency Division Duplex (FDD), and Orthogonal Frequency Division Multiplexing (OFDM). The WiMAX standard provides flexibility in spectrum usage by supporting both FDD and TDD. Thus, it can operate in both FDD/OFDM and TDD/OFDM modes.

WiMAX MAC ensures a number of QoS measures not seen in other wireless standards. Perhaps its greatest value is providing for dynamic bandwidth allocation. In order to support different kind of services (data and voice), WiMAX MAC accommodates both continuous and burstly traffic and assigns different QoS parameters to each service.

Currently, there are few WiMAX implementations in the market. For our measurements, we used the Expedience technology [12] which was the first commercial NLOS product for the broadband wireless market. The Expedience technology comes with an RSU-3510 modem that includes the following technologies:

– OFDM - The OFDM's implementation utilizes hundreds of individual carriers and a process for mapping a user's data to those carriers, to actually leverage the presence of multi-path to transmit and receive robustly in the NLOS service environment.
– TDD - The upstream and downstream links are on the same 6 MHz RF channel. So, it provides the highest flexibility in frequency utilization.
– Adaptive Modulation (AMOD) - It enables higher capacity per sector and a robust RF link all the way to the edge of the wireless cell. AMOD can double the capacity per sector over a single modulation level.
– Direct Burst Detection (DBD) - The equalization and detection scheme eliminates the need for equalizer training sequences before sending actual data. DBD minimizes latency of user packets and optimizes the efficiency of the MAC protocol.
– Expedience MAC - It is designed specifically for last mile broadband wireless data networks to optimize the scheduling and delivery of data and voice packets over a multiuser NLOS airlink.

From the features above, we can consider that the Expedience technology is very similar to WiMAX in the physical layer, since they both support OFDM, TDD and AMOD with the same modulation techniques. Unfortunately, regarding the link layer, we have not yet been able to get the technical details of the Expedience MAC. Nevertheless, since it was designed and optimized specifically for the last mile broadband access, we consider that it must have a very close behavior to WiMAX's MAC layer.

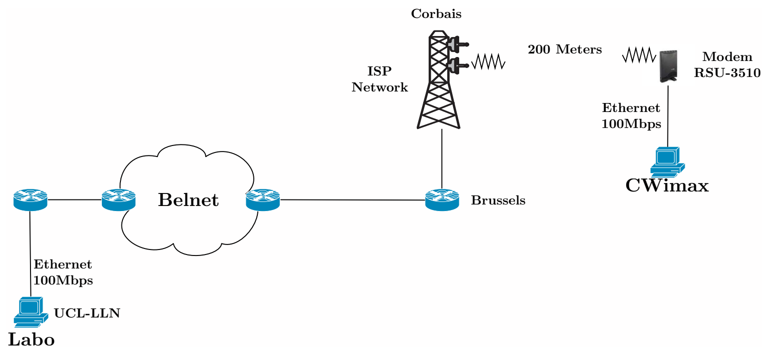

Fig. 1. WiMAX testbed for TCP measurements

		LABO	**CWIMAX**
OS	Type	Linux	Linux
	Kernel	2.6.8-2-386	2.6.15-23-386
CPU	Model	Intel Pentium	Intel Pentium
	Frequency	333.045 Mhz	2.40 Ghz
Network card		100 Mbps	100 Mbps

Table 1. Computers configuration

3 Preliminary Evaluation

3.1 Methodology

Our experimental testbed was very simple as shown in Fig. 1. We installed two computers, one (LABO) placed in our laboratory (Université Catholique de Louvain – Belgium) and the other (CWIMAX) in Corbais (Belgium), where the wireless ISP has coverage for the WiMAX service. Table 1 describes in detail the characteristics of both machines. LABO was connected to a fast Ethernet (100 Mbps) switched LAN. The traffic sent between the two computers crossed the wired Internet between the lab and the ISP's base station antenna. The traffic went then throughout the wireless segment between ISP's antenna and CWIMAX. The approximate distance between CWIMAX and the ISP's antenna was 200 meters. In order to send and receive the wireless signal, CWIMAX was connected through a fast Ethernet interface to the RSU-3510 Broadband wireless access modem which operates in the 3.4 – 3.6 Ghz band with a channel bandwidth from 3 to 7 MHz. The modem supports also OFDM modulation 4/16/64 QAM, operates in TDD mode with an output power of 2 Watts. The theoretical bandwidth of the commercial service was 3 Mbps for downlink and 256 Kbps for uplink.

The wired segment was mainly composed of the Belnet National Research Network (NREN) with uncongested links from 1 Gbps to 10 Gbps. The number of routers that are crossed in order to reach the destination was seven. Since the WAN links were fast enough and uncongested, we can considered that most of the delay on our connection took place in the ISP wireless segment. We evaluated the RTT difference between a ping between CWIMAX and LABO and CWIMAX and its default gateway, i.e., the first router just crossing the wireless segment. The difference was only of 2 ms.

We synchronized the clocks of both computers using NTP 4.2, and we ran `TCPdump` 3.8.3 in the background in order to capture packets sent and received by both computers and later analyze the packets from the `TCPdump` files.

We used the BIC TCP congestion avoidance mechanism, implemented by default in the Linux kernel that we used. BIC TCP includes *Selective Acknowledgements* (SACK) [13] and *Duplicate SACKs* (D-SACK) [14]. This is especially good on very lossy connections since SACKs allows one to only retransmit specific parts of the TCP window which lost data and not the whole TCP window. D-SACK is an extension to standard SACK and is used to tell the sender when a packet was received twice (i.e., it was duplicated). BIC TCP also includes the *Forward Acknowledgment* system (FACK) [15] in Linux, a special algorithm that works on top of the SACK options. The main idea of the FACK algorithm is to consider the most forward selective acknowledgment sequence number as a sign that all the previous un(selectively) acknowledged segments were lost. This observation allows one to improve recovery of losses significantly. In this way, we tried to improve the TCP performance as much as possible.

Basically, we ran a two-way test: firstly, from LABO to CWIMAX (i.e., downlink test) and secondly, from CWIMAX to LABO (i.e., uplink test), each test lasted ten minutes. Both parts of the two-way test were alternatively performed, with a break of one hour between each part. This two-way test was repeated every two hours during a whole day. It results, consequently, in a collection of twelve data sets in each way. We used Iperf 2.0.2 for generating the ten minutes required TCP traffic. The first two-way test was run on October 17, 2006 at 18:00. The last measurement was performed on October 18, 2006 at 16:00.

In order to obtain the RTT, the delay and packet loss ratio in each direction, we analyzed the traces we logged using `TCPtrace` and `TCPdump`.[3]

3.2 Results

In this section, we present preliminary results for RTT, delay and packet loss ratio. Because of space constraints, we choose to restrict our discussion to the downlink test. Interested readers might find further results about uplink tests, RTT and radio condition influence in an extended version of this paper [16].

Fig. 2 shows the RTT average, the maximum RTT, the minimum RTT and the standard deviation values found in each test during the whole day, the timing is indicated by the horizontal axis.

[3] Our data set is is freely available at `http://inl.info.ucl.ac.be/files/data-18.11.2006.tar.gz`.

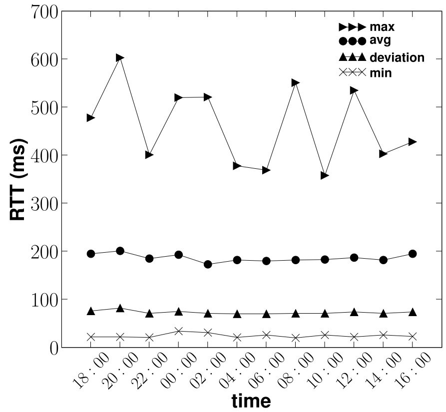

Fig. 2. Downlink - min, max, avg RTT values and deviation (whole day)

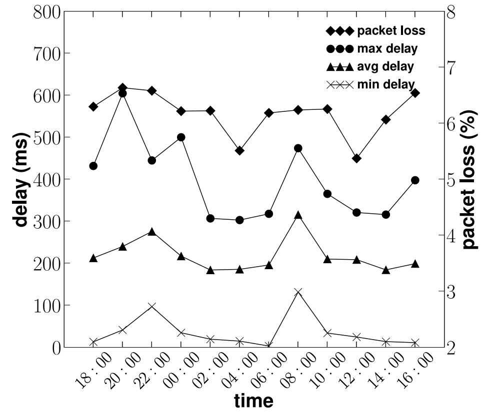

Fig. 3. Downlink – delay and packet loss (whole day)

The mean RTT over the whole day is 186.5 ms. We believe that this value is acceptable for most of applications. However, if we consider the theoretical bandwidth offered by the service provider (i.e., 3 Mbps), the RTT value should be lower. All the values plotted on Fig. 2 show certain stability during the day. The biggest values for the average (201 ms) and maximum RTT (603 ms) are found at 20:00. These values were expected since the Internet service is mainly focused on home users and, at this moment of the day, a higher number of customers is supposed to be connected. Also, RTT values are high because of the low bandwidth of the upstream link. This leads to delay the reception of the ACKs that are sent back to the sender. We could also notice that packets are lost more frequently in the upstream link (see [16] for further details). Potentially, ACKs might be lost too, impacting therefore the downlink performance. Finally, the maximum RTT is, however, unacceptable for real-time applications, such as on-line games.

Fig. 3 shows the downlink delay and packet loss average over the whole day. The left-side vertical axis of Fig. 3 displays the delay (in ms) while the right-side gives the packet loss percentage. Regarding the delay, we were concerned by the minimum, maximum and average values for each downlink test. The horizontal axis gives the time, i.e., the hour at which each downlink test was carried on.

Looking first at the delay, we see that the average downlink delay fluctuates between 184.12 ms (02:00) and 315.47 ms (08:00). The maximum delay, 604.67 ms, is found during the test carried on at 20:00. This is quite expected as it corresponds to a rush hour for a domestic access network. Regarding the packet loss ratio, we see that it is pretty high and it fluctuates between 5.37% (12:00) and 6.63% (20:00). This high packet loss ratio cannot be explained only with the congestion caused by the TCP traffic injected by our measurements. Two other factors must be taken into considerations. First, the traffic caused by the ISP

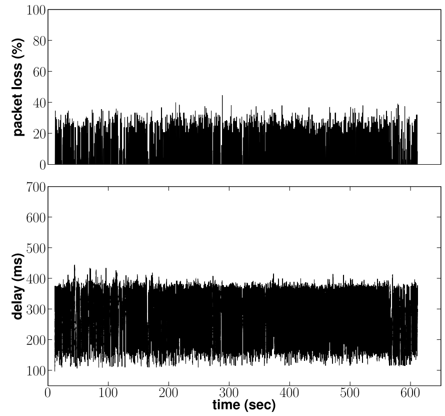

Fig. 4. Downlink – delay and packet loss (22:00 test)

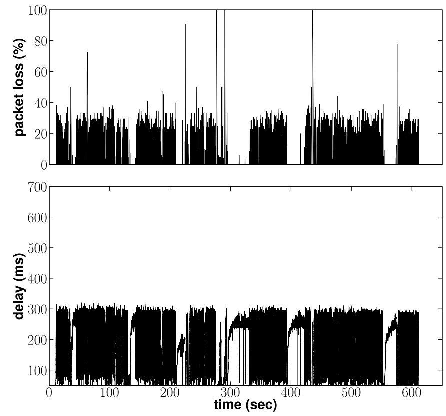

Fig. 5. Downlink – delay and packet loss (12:00 test)

clients can lead to congestion or an overflow in the WiMAX antenna placed in Corbais. Second, the WiMAX link can suffer from radio failures or temporary signal fades in the wireless modem, resulting in packet losses. From Fig. 3, we can observe a certain correlation between the delay and the percentage of packet loss. This is specially true if we consider the maximum delay values: the longer the delay, the larger the percentage of packet loss. In particular, the highest packet loss ratio is found when the maximum delay was observed.

Even though the delay and packet lost average are not so different during the twelve tests, we could find important differences when looking individually at each test. During six tests, the delay and packet lost were very stable and did not present any burst, like during the test carried on at 22:00 illustrated in Fig. 4.

On the contrary, the other six tests presented, at least, one bursty packet loss behavior. By bursty, we mean a packet loss ratio higher than 50% during some period of time, typically 100 ms. Within four tests, we found seven bursts with 100% of packet lost. In three tests, we found packet loss ratio higher than 80% but lower than 100%. Finally, in six tests, we found 29 bursts higher than 50% of packet loss but lower than 80%.

From our data set, we noticed that the test that presented the largest number of burst losses was the 12:00. The delay and packet loss for the 12:00 test are shown in Fig.5. The horizontal axis gives the time, in seconds, starting from 0 to 600 as each tests lasted ten minutes. The vertical axis is divided in two parts. The upper part shows the packet loss percentage and the lower part gives the delay (in ms).

The burst durations for the 12:00 test were very different, as depicted in Fig. 6. The horizontal axis shows a burst duration interval (in ms). Each interval is incremented by 100 ms. The vertical axis shows the number of burst losses.

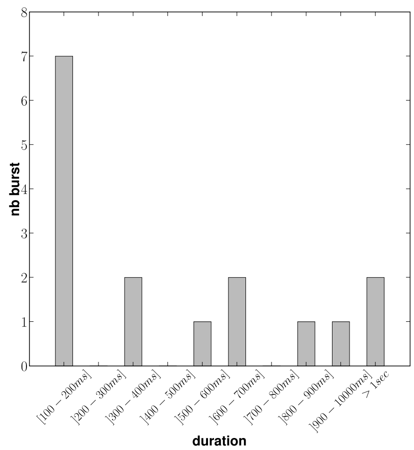

Fig. 6. Downlink – loss bursts duration distribution (12:00 test)

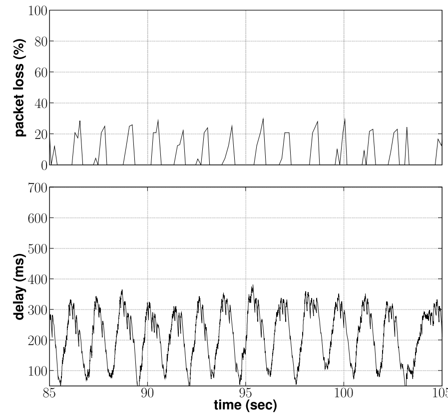

Fig. 7. Downlink – delay and packet loss (18:00 test)

As shown in Fig. 6, a large portion of bursts were short: between 100 and 200 ms. However, we notice that eight burst losses lasted more than 0.5 seconds. In particular, two bursts were longer than one second (1.06 and 1.36 seconds).

After this general analysis, we zoom in a smaller period of time to analyze the behavior of a representative TCP connection. During this test, the average goodput was 1.92 Mbps, the total amount of packets sent was 106,534 with an average rate of 177.55 packets/sec. Fig. 7 shows the delay and packet loss, for the 18:00 test, from the second 85 to the second 105. During that time, TCP presented a cyclic behavior expanding and contracting its congestion window, producing the delay of the packets increased and decreased. TCP takes around 700 ms to fully expand its congestion window and makes use of the total bandwidth available. However, the amount of packets losses was considerably high. Packet loss started and was detected by TCP after the delay was around 300 ms or above, normally when the delay reached the peaks in the graph. At that moment, the congestion window size was decreased, resulting in a delay reduction. As a consequence, the packet loss started to decrease also. However, it is important to point out that we have 20% or even 30% of packets lost during some moments which represents a pretty high ratio. The total percentage of packets lost during the 10 minutes test was around 6%. It is obvious to guess that this high percentage of packets lost has a strong impact on the network goodput.

Finally, we conclude this section by saying that Fig. 7 shows a correlation between packet losses and delay. This is a strong indicator for buffer overflow. Further, the test was carried on at 18:00, which is supposed to be a rush hour for Internet connections. If these losses were caused by wireless failures, the packet losses would be uncorrelated with the delay.

4 Related Work

The most similar work to our proposal is Chakravorty and Cartwright's work [17] on TCP over GPRS networks. They found, in GPRS networks, that the RTT was larger than 1,000 ms and can be highly variable. As a consequence, the available bandwidth could be quite variable too. However, the packet losses were relatively rare. These network characteristics do not interact well with current TCP implementations. Chakravorty and Cartwright showed that it takes many seconds before a new TCP connection could expand its congestion window to make full use of the available bandwidth. After that, TCP continues to expand the window needlessly, resulting in excessive queuing in the GPRS router. Both situations lead to very poor performance of protocols like HTTP (inflated RTT of 10 seconds). They finally showed that a simple transparent proxy interposed between the fixed and the GPRS networks improves the TCP connection performance.

Ramachandran and Bostian [18] modeled IEEE 802.16 and evaluated the performance of its MAC layer over several physical layer options using OPNET. They demonstrated the need for an algorithm to dynamically switch between different PHY burst profiles in order to improve the protocol performance. They also showed that the link layer delay does not affect significantly the TCP delay.

Xylomenos and Polyzos [19] measured TCP over WLANs with different wireless adapters. They found that PCMCIA adapters are slower than ISA due to less aggressive timing and buffer limitations. Faster senders can overwhelm slower receivers, leading to semi-periodic packet loss. When two ISA hosts communicate, many collisions occur in TCP, leading to a much degraded performance. These collisions are between one data and one acknowledgment packet. As long as these synchronization problems are avoided, CSMA/CA performs well with bidirectional traffic as it was tested with one ISA and one PCMCIA interface which have different timing mechanisms. MAC layer retransmissions could be beneficial for TCP but problematic for protocols or applications that prefer sending new packets than retransmitting old ones.

Vacirca and Cuomo [20] studied the TCP uplink/downlink unfairness in WLAN connections. Normally, in a common scenario the downlink system performance decreases inversely with the number of uplink competing flows. The performed test showed that this unfairness depends on the number of transmission attempts and access point (AP) buffers size. They also find that the uplink packet lost probability is mainly due to packet collisions on the wireless channel, while the downlink packets lost is due to collisions in the channel and to congestion in the AP buffer. To solve these anomalies, they proposed to include in AP a simple packet scheduling policy (software module) that succeeds in alleviating the uplink/downlink unfairness.

5 Conclusion

Nowadays, broadband wireless access are more and more used. In particular, the recently standardized WiMAX has raised attention due to its high potential for providing network community to a large set of users. However, due to its youth, the impact of WiMAX on network protocols, such as TCP, is unknown.

In this paper, we presented what is, to the best of our knowledge, the first measurement study of the TCP behavior in a WiMAX environment. Using a WiMAX commercial connectivity in Belgium, we performed a set of active measurements in order to better understand how TCP works in a WiMAX network.

In particular, in this paper, we focused on the RTT and the relationship between delay and packet loss ratio. We discovered that the RTT can be very high. We found that TCP presents an acceptable cyclic behavior: it took around 700 ms to expand its congestion window to make use of the full bandwidth available. The delay was 219.176 ms on average for downlink which is acceptable for most web applications. Nevertheless, the percentage of packet lost was high, 6.177% on average. During some TCP tests, we found some packet lost bursts that vary in size and duration. Some of these bursts were of 100% of packets lost and lasted more than one second.

Nevertheless, this paper represents our first step towards a good comprehension of TCP within WiMAX. Our future work is to reproduce TCP behavior in our lab with a model implemented in an emulator. Further investigations and measurements are needed to achieve this goal.

Acknowledgements

Jesus A. Perez was supported by a grant provided by the Coimbra Group and the Université Catholique de Louvain. Benoit Donnet and Olivier Bonaventure were supported by the European Commission-founded IST034819 OneLab Project. Finally, authors would like to wish to thank Tanguy Guillaume for helping them in processing the huge amount of packet's information.

References

1. IEEE 802 Standard Working Group: Wireless Lan Medium Access Control (MAC) and Physical Layer (PHY) Specifications. Standard 802.11-1999, Institute of Electrical and Electronics Engineers (Jun. 1999)
2. IEEE 802 Standard Working Group: IEEE standard for local and metropolitan area networks–part 16: Air interface for fixed broadband wireless access systems–amendment 2: Medium access control modifications. Standard 802.16a-2003, Institute of Electrical and Electronics Engineers (Apr. 2003)
3. Forsman, J., Keene, I., Tratz-Ryan, B., Simpson, R.: Market opportunities for WiMAX take shape (Dec. 2004) See `http://www.gartner.com`.
4. Lisong, X., Khaled, H., Injong, R.: Binary increase congestion control for fast long-distance networks. In: Proc. IEEE INFOCOM. (Mar. 2004)
5. Cheng, J., David, W., Steven, L.: Fast TCP: motivation, architecture, algorithm and performance. In: Proc. IEEE INFOCOM. (Mar. 2004)
6. Casetti, C., Gerla, M., Mascolo, S., Sanadidi, M.Y., Wang, R.: TCP Westwood: Bandwidth estimation for enhanced transport over wireless links. In: Proc. ACM SIGCOMM. (Jul. 2001)

7. Xu, K., Tian, Y., Ansari, N.: TCP-Jersey for wireless IP communications. IEEE Journal on Selected Areas in Communication (JSAC) **22**(4) (May 2004) 747–756
8. Fu, C.P., Liew, S.C.: TCP Veno: TCP enhancement for transmission over wireless access networks. IEEE Journal on Selected Areas in Communication (JSAC) **2**(21) (Feb. 2003) 216–228
9. Wu, E.K., Chen, M.Z.: JTCP: Jitter-based TCP for heterogeneous wireless networks. IEEE Journal on Selected Areas in Communications (JSAC) **22**(4) (May 2004) 757–766
10. Todorovic, M., Lopez-Benitez, N.: Efficient study of tcp protocols in infrastructured wireless networks. In: Proc. Internet Conference on Networking and Services (ICNS). (Jul. 2006)
11. Goense, D., Thelen, J., Langendoen, K.: Wireless sensor networks for precise phytophthora decision support. In: Proc. 5th European Conference on Precision Agriculture (5ECPA). (Jun. 2005)
12. NextNet: Expedience technology See `http://www.nextnetwireless.com/products_tech.asp`.
13. Mathis, M., Madhavi, J., Floyd, S., Romanow, A.: TCP selective acknowledgment options. RFC 2018, Internet Engineering Task Force (Dec. 1996)
14. Floyd, S., Madhavi, J., Mathis, M., Podolsky, M.: An extension to the selective acknowledgement (SACK) option for TCP. RFC 2883, Internet Engineering Task Force (Jul. 2000)
15. Mathis, M., Mahdavi, J.: Forward acknowledgment: Refining TCP congesstion control. In: Proc. ACM SIGCOMM. (Aug. 1996)
16. Perez, J.A., Donnet, B., Bonaventure, O.: Preliminary analysis of the tcp behavior in 802.16 networks. Technical Report 2006-09, Université Catholique de Louvain (Dec. 2006) See `http://www.info.ucl.ac.be/~donnet/wimax/wimax-TechRep.pdf`.
17. Ravij, C., Cartwright, J., Pratt, I.: Practical experience of TCP over GPRS. In: Proc. Global Telecommunications Conference (GLOBECOM). (Nov. 2002)
18. Ramachandran, S., Bostian, C.W., Midkiff, S.F.: Performance evaluation of IEEE 802.16 for broadband wireless access. In: Proc. OPNETWORK. (Aug. 2002)
19. Xylomenos, G., Polyzos, G.: TCP and UDP performance over a wireless LAN. In: Proc. IEEE INFOCOM. (Mar. 1999)
20. Vacirca, F., Cuomo, F.: Experimental results on the support of TCP over 802.11b: an insight into fairness issues. In: Proc. 3rd Conference on Wireless On-Demand Network Systems and Services (WONS). (Jan. 2006)