Internet access performance in LTE TDD

Riikka Susitaival[1], Henning Wiemann[3], Jessica Östergaard[2], Anna Larmo[1]
Ericsson Research
[1]Finland, [2]Sweden, [3]Germany
[riikka.susitaival, henning.wiemann, jessica.ostergaard, anna.larmo]@ericsson.com

Abstract **The Time Division Duplex (TDD) uplink-downlink configuration of the 3GPP Long Term Evolution (LTE) determines how the ten subframes in a radio frame are divided between the downlink and the uplink. The specified configurations cover a wide range of allocations from a downlink-heavy resource distribution ratio (9:1) to an uplink-heavy ratio (2:3). In this paper, we compare the performance of Internet access using the TCP protocol in different downlink-heavy asymmetries. We find that the performance depends on many factors such as the transferred file size, the control channel errors and the downlink/uplink traffic mix. When the file size is small, TDD can not fully utilize its potentially higher downlink capacity because of longer uplink access delays as well as shortage of uplink resources in the chosen configurations. With an increased file size, this effect fades away and TDD provides higher download bit rates than FDD. The realized increase in bit rate is however not as high as the calculated increase in available downlink resources.**

Keywords- LTE, TDD, Internet access, TCP

I. INTRODUCTION

The first release (Release 8) of the Long Term Evolution (LTE) has recently been standardized by 3GPP. LTE provides high peak data rates up to 300 Mbps, improved spectrum efficiency, and reduced radio access delays. One key requirement in the development of LTE has been spectrum flexibility; LTE can be operated in different spectrum allocations from 1.4 to 20 MHz and in paired or unpaired spectrum [1]. The Frequency Division Duplex (FDD) mode uses paired spectrum where different carrier frequencies are used for downlink (DL) and uplink (UL). The Time Division Duplex (TDD) mode, on the other hand, uses unpaired spectrum and a single carrier frequency with separation of downlink and uplink in time. In TDD, a subframe in a radio frame can either be an uplink, a downlink or a special subframe. 3GPP has specified seven different uplink-downlink configurations.

The possibility to distribute TDD radio resources unevenly among uplink and downlink has many consequences. Depending on traffic characteristics and asymmetries, it is possible to dimension uplink/downlink capacity in such a way that the available spectrum is utilized most efficiently. This is done by selecting an appropriate TDD configuration. In order to avoid interference between uplink and downlink, such as base station to base station and mobile to mobile interference, the network is preferably synchronized and all macro cells within a geographical area use the same TDD configuration. Hence, the configuration should be set considering the average load in the system. Another characteristic of LTE TDD is that when the number of downlink subframes is different from the number of uplink subframes, there is no direct one-to-one mapping between uplink and downlink subframes. This has an effect on the Hybrid Automatic Repeat reQuest (HARQ) operation of LTE TDD as well as on the L1/L2 control signaling making it different compared to FDD.

This work concentrates on Internet access using the Transmission Control Protocol (TCP) [7]. In the case of a download, the data goes in the downlink and control data (TCP acknowledgements (ACKs)) in the uplink. The initial TCP congestion window is small and thus the radio link is typically not fully utilized in this first phase of the data transfer. The TCP performance suffers from high delays in this phase as they prevent the congestion window from increasing quickly, i.e., the link from being further utilized. Once the congestion window of the TCP transmitter exceeds the bandwidth-delay product of the link, TCP is rather insensitive to delays.

In this paper, we study downlink-heavy TDD configurations which aim at providing improved downlink performance. First, the HARQ aspects of LTE TDD are introduced and user level latencies are evaluated. Using advanced system simulations, we then study the actual performance of TDD and compare it to FDD. We use both pure download as well as mixed downlink/uplink traffic models. In addition, we study how L1/L2 control signaling errors affect the performance in TDD.

So far, specific features and lower-layer performance aspects have been studied and published for LTE TDD. For example, Astély et al. simulate TDD configuration 1 in [1], and Rahman and Astély study TDD ACK-NACK bundling with link level simulations in [2]. In contrast to previous studies, we consider also higher layer protocols aspects and two-directional traffic models using both uplink and downlink in this paper. TCP performance studies focusing on TDD asymmetries exist for WiMAX ([3], [4]), but due to inherent differences in the radio access network architectures, as well as TDD frame structures, these results cannot directly be applied to an LTE network.

This paper is organized as follows: in Section II, the TDD frame structure is described briefly. HARQ related issues such as user plane latency and ACK-NACK bundling, are evaluated in Sections III and IV. In Section V, the simulation model is presented and Section VI studies TDD performance by simulations. Finally, conclusions are drawn in Section VII.

978-1-4244-2519-8/10/$26.00 ©2010 IEEE

II. TDD Frame Structure

Similar to FDD, the TDD radio frame consists of 10 subframes, each having a length of 1 ms. In TDD, the subframe can be either DL, UL, or a special subframe between the DL and UL period. 3GPP has defined seven uplink-downlink configurations. The DL:UL patterns of the configurations are listed in Table I. For example, in configuration 1, there are always three consecutive DL subframes (including both the normal DL subframes and the special subframes), followed by two UL subframes. The same pattern is then repeated starting from DL. See Figure 1 for illustration of configuration 1 in more detail.

TABLE I. DOWNLINK-UPLINK CONFIGURATIONS SPECIFIED BY 3GPP.

Configuration	DL:UL
0	2:3
1	3:2
2	4:1
3	7:3
4	8:2
5	9:1
6	3:3:2:2

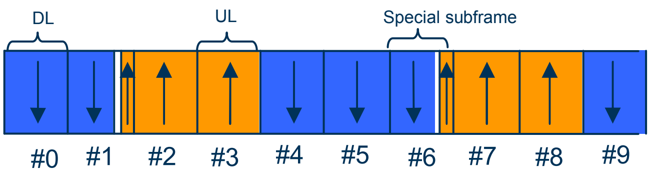

Figure 1. TDD uplink downlink Configuration 1, TDD DL:UL 3:2.

The special subframe consists of a downlink part (DwPTS), a guard period (GP), and an uplink part (UpPTS). The DL part of the special subframe can be considered as a normal downlink subframe for data and control but with a reduced number of data symbols. The guard period consists of a number of idle symbols where nothing is transmitted to protect against DL-UL interference. Finally, the UL part of the special subframe is considerably shorter than the downlink part and is only used for channel sounding and random access preamble transmission (no user data transmission). 3GPP has defined nine different configurations for how the 14 symbols of the special subframe are divided between the downlink part, the uplink part, and the guard period. The special subframe configurations are listed in the 3GPP specification [6].

III. HARQ Timing and User Plane Latencies in TDD Mode

The multiprocess stop-and-wait HARQ mechanism used in LTE to enable a low-error-rate data stream is based on a known delay between transmission of data and transmission of feedback. The feedback can be either a positive acknowledgement (ACK) or a negative acknowledgement (NACK). In FDD, the delay between the data and the feedback transmission is always 4 ms. In TDD, if there is a DL transmission in subframe *n,* the subframe *n+4* is not necessarily an UL subframe. Thus, the feedback has to be delayed. The same problem appears for uplink traffic and the corresponding downlink HARQ feedback. The exact HARQ timing relations for each TDD configuration have therefore been specified by 3GPP and can be found in [5].

Latencies play an important role in end-to-end performance of user applications. In the following, the transmission latency in uplink is studied and compared between FDD and TDD. It is assumed that the User Equipment (UE) is time-aligned, and that it has been configured with dedicated resources for sending a Scheduling Request (SR) to ask for an uplink transmission grant, but that it does not yet have the grant for transmission.

The procedure to transmit data, e.g., TCP ACK, in the uplink after data has arrived to the UE's buffer is as follows: 1) In the subframe where the UE has an SR resource available, the UE transmits an SR, which is a one-bit flag to indicate that the UE has new data. 2) The eNB receives the SR and after a processing delay, an initial grant is transmitted to the UE allocating frequency resources for UL transmission. The processing delay of the SR is selected to 3 ms which is similar to HARQ processing times specified for the UE. 3) Using the granted resources, the UE transmits data as well as a Buffer Status Report (BSR) to indicate to the eNodeB how much data it still has available in its buffer after the transmission. The delay from the grant transmission by the eNodeB to the uplink transmission by the UE is fixed to 4 ms in FDD but in TDD, it varies between uplink-downlink configurations. 4) When the eNB has received the BSR, it can continue allocating uplink resources to the UE and the UE can perform further uplink transmissions.

Figure 2 depicts the uplink data transmission procedure for FDD and the TDD 3:2 configuration. From the diagrams we can see that the delay from the SR transmission to the second scheduled transmission from which the UE can be scheduled continuously based on the buffer information, is 16 ms in FDD whereas in the studied TDD configuration, it is 20 ms. It can be calculated that this delay is 20 ms also for most of the other TDD configurations.

IV. ACK-NACK Bundling

In addition to HARQ timing and user plane latencies, there is another major difference between TDD and FDD: with downlink-heavy configurations, the UE must transmit multiple HARQ feedbacks in one uplink subframe. 3GPP has standardized two mechanisms to achieve that: ACK-NACK bundling and ACK-NACK multiplexing. In ACK-NACK multiplexing, all HARQ feedbacks are explicitly signaled in one uplink subframe and since more bits are transmitted, the required SINR is higher. In this study, ACK/NACK bundling is used and thus introduced in more detail.

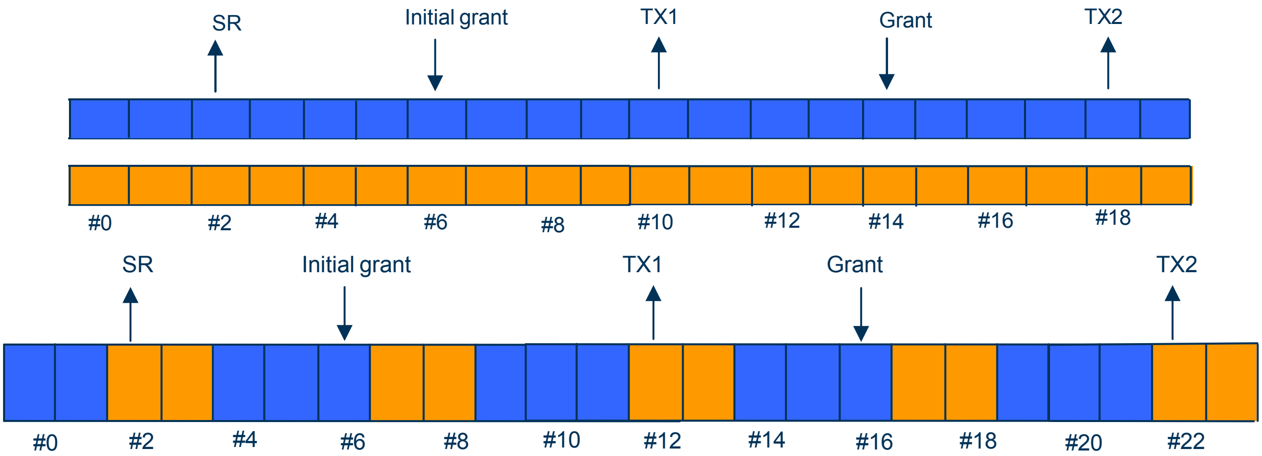

Figure 2. The user plane latencies in FDD (upper figure) and TDD 3:2 (lower figure). Blue subrames are for DL and orange for UL.

In ACK-NACK bundling, a logical AND is done over the HARQ feedback messages. If all DL HARQ processes in a so-called bundling window are correctly decoded, an ACK is sent by the UE. If one or more of the processes are not correctly decoded, the UE transmits NACK and all processes in the window have to be retransmitted. In this context, the bundling window refers to those DL subframes whose feedback messages are reported in a single UL subframe. If a multicodeword transmission is used for any of the subframes in a bundling window, then a bundled feedback is sent per codeword. ACK-NACK bundling with the TDD 3:2 configuration is illustrated in Figure 3. The feedback timing relations for all TDD allocations are defined in [5].

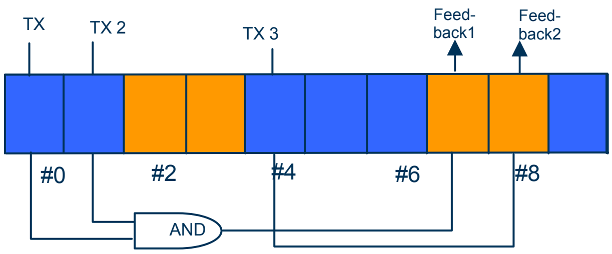

Figure 3. ACK/NACK bundling with TDD 3:2. The first bundling window is over subframes #0 and #1 and the second over subframe #4.

In the case when multiple DL transmissions are scheduled by the eNB in a bundling window, the UE might miss one of the assignments. The design of ACK-NACK bundling has considered this: The DL assignment carries a Downlink Assignment Index (DAI) indicating how many assignments the UE should have received so far within the current bundling window. If the UE detects that the DAI differs from the number of correctly received DL assignments, it does not send any HARQ feedback and the eNB can detect this. However, the eNB cannot know which of the transmissions was missed, and thus the whole bundle has to be retransmitted.

V. Simulation Model

The simulations are done with a detailed system-level simulator implementing the radio link with 3GPP L1/L2 and transport protocols. The path loss model is a typical distance-dependent model with wrap around.

The downlink-heavy TDD configurations, 3:2, 7:3, 4:1 and 9:1, are studied. The 8:2 configuration is excluded because it offers very similar DL:UL ratio to TDD 4:1. We choose configuration 7 for the special subframe including 12 DL OFDM symbols. The main parameters are listed in Table II.

TABLE II. Simulation Assumptions.

Parameter	Value
System Bandwidth	TDD: 10 MHz, FDD: 5 + 5 MHz
Network layout	Hexagonal grid, 21 cells
Site-to-site distance	500 m
Multipath model	Typical Urban, 3 km/h
Antenna model	SIMO, 1 transmit and 2 receive antennas in both uplink and downlink
Max. DL TX power	FDD: 20 W TDD: 40 W
Max. UL TX power	FDD and TDD: 0.25 W
Scheduling	Round robin
SR on uplink control channel	10 ms periodicity
HARQ transmissions	DL: max 7, UL: max 8
RLC mode	Acknowledged (AM)

VI. Simulation Results

In this section, results from multiple simulations are evaluated to compare the download performance of TDD and FDD. As a theoretical background, in Table III the number of DL and UL resource elements (RE) for data in a radio frame is presented for FDD and TDD. The numbers exclude overhead for reference signals and uplink and downlink control channels.

TABLE III. THE NUMBER OF RESOURCES ELEMENTS IN A RADIO FRAME.

	DL REs	UL REs
FDD	31500	27720
TDD 3:2	36600 (+16 %)	22176 (-20 %)
TDD 7:3	43500 (+38 %)	16632 (-40 %)
TDD 4:1	49200 (+56 %)	11088 (-60 %)
TDD 9:1	56100 (+ 78 %)	5544 (-80 %)

A. Single user download peak rates

First, download peak rates are studied. A single user at a time is placed in the system for requesting and downloading a single file with TCP. Once done, the user leaves the system and a new user arrives. The file size varies from 100 kB to 10 MB. The arrival position is selected randomly in the system according to a uniform distribution.

Figure 4 shows the CDF for the download performance in terms of the object bit rate, that is, the file size divided by the download time, when the file size is small, 100 kB. Even though the downlink capacity is larger in TDD, most of TDD configurations do not improve the performance when compared to FDD. This can be explained by longer uplink access delays of TDD as studied in Section III. The TCP congestion window increases exponentially for each received TCP ACK. If these feedback messages are delayed due to longer UL delays, the congestion window increases more slowly and the object bit rate remains lower. In addition, from Figure 4 it can be seen that TDD 9:1 has the worst performance at the lower 10^{th} percentile. This can be explained by too small instantaneous uplink capacity and power, especially when the UE is in bad radio conditions. TCP ACKs can arrive to the uplink buffer in bursts and it takes a relatively long time to empty the queue.

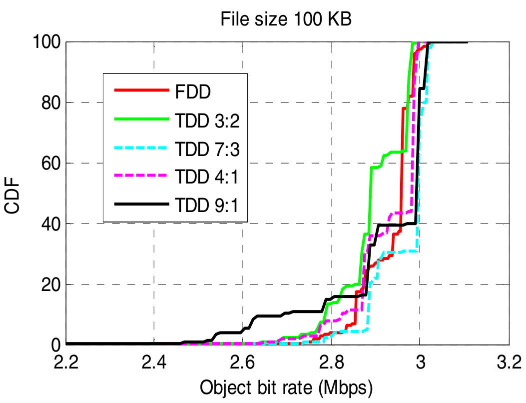

Figure 4. The download object bit rate, file size 100 KByte.

With a larger file size, the TCP initial phase does not have as much impact on the overall transfer time. Thus, uplink limitations do not have such a big effect. When the file size is medium, 1MB, in Figure 5, the object bit rate of TDD 4:1 is highest but TDD 9:1 suffers still from uplink limitations. Finally, with a large file size, 10 MB, in Figure 6, the extreme configuration TDD 9:1 provides the highest bit rate. However, the difference is not as large as the numbers in Table III may indicate. The peak rate increases 62% in simulation results whereas the number of DL resource elements increased 78% when comparing TDD 9:1 and FDD. This is due to many factors such as ACK/NACK bundling resulting in unnecessary retransmissions, and UL limitations and delays.

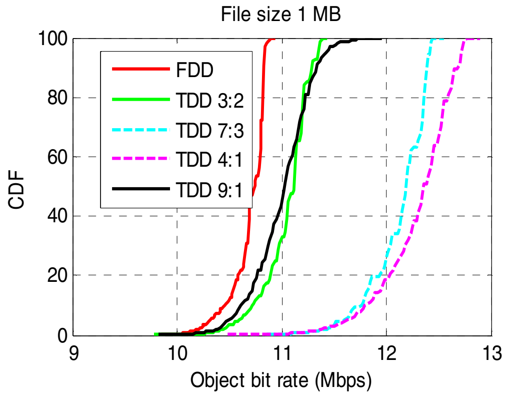

Figure 5. The download object bit rate, file size 1 MB.

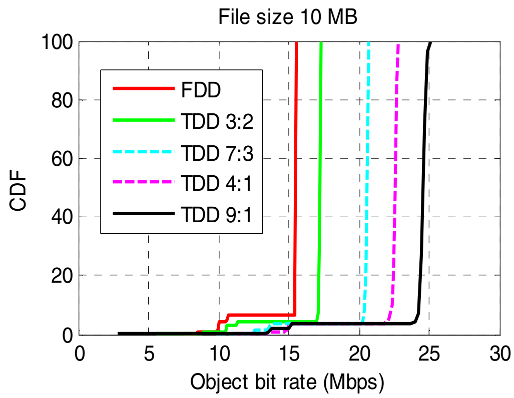

Figure 6. The download object bit rate, file size 10 MB.

B. Download mean object bit rates

In the following, we investigate the download performance in a multi-user scenario. The users arrive to the system according to a Poisson process with varying arrival intensity between scenarios. The file size is 1 MB. The user object bit rate as a function of the DL throughput is depicted in Figure 7. As the load and the number of users in the system increases (when increasing the user arrival rate), the overall throughput, i.e., the number of served bits, increases as well. On the other hand, when there are more users in the system, each individual user gets fewer resources and the object bit rate decreases. From Figure 7, we can see that the higher the load of the system, the bigger the difference between TDD and FDD. The increased DL capacity can be better utilized when there are more users.

C. Impact of control channel signaling errors on download

In this subsection, the impact of L1/L2 control channel errors on TDD download performance is studied. We focus on DL HARQ related control channels. The errors are independent and occur with probability p in each subframe. The values for p are selected according to the 3GPP targets:

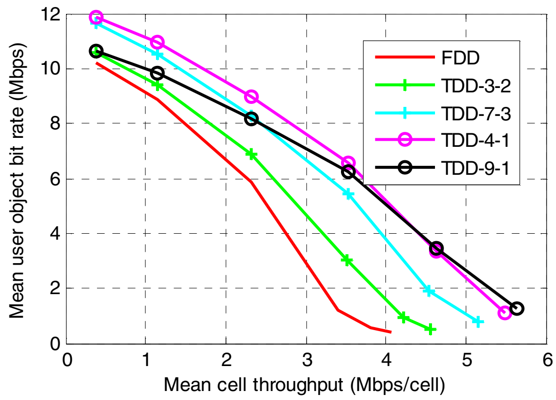

Figure 7. The mean object bit rate as a function of cell throughput in multi-user scenario for file download. File size 1 MByte.

1) DL assignment loss on Physical Downlink Control Channel (PDCCH) (p=10^{-2}). If this error occurs, the UE does not transmit any HARQ feedback. The eNB does not receive any feedback, i.e., it detects DTX, and retransmits all active DL HARQ processes of the corresponding ACK/NACK bundling window.

2) DTX-to-ACK error on Physical Uplink Control Channel (PUCCH) (p=10^{-2}). This error leads to data loss of the bundling window at the HARQ level and an RLC retransmission if operated in RLC acknowledged mode.

3) NACK-to-ACK error on PUCCH (p=10^{-4}). This error also leads to data loss at HARQ level and possible RLC retransmission.

In Table IV, the download object bit rates (median) in a single-user scenario are depicted with and without the above control channel errors. It can be seen that the impact of the errors is higher with TDD as compared to FDD. This difference is due to ACK-NACK bundling. One missed DL assignment leads to the retransmission of all HARQ processes in the bundling window. By operating the system such that the missed assignment probability is lower, the losses due to missed assignments can be made lower. Alternatively, the loss seen in TDD can be controlled by configuration selection.

TABLE IV. 50% PERCENTILE DOWNLOAD OBJECT BIT RATE. SINGLE USER SCENARIO, FILE SIZE 10 MB

Mode	No errors	With errors	Difference
FDD	15.5 Mbps	15.4 Mbps	-0.65 %
TDD 3:2	17.3 Mbps	16.8 Mbps	-2.9 %
TDD 7:3	20.5 Mbps	20.1 Mbps	-1.9 %
TDD 4:1	22.5 Mbps	21.5 Mbps	-4.4 %
TDD 9:1	24.5 Mbps	22.6 Mbps	-7.8 %

D. Mixed traffic scenario

Finally, a mixed traffic scenario including both download and upload users is studied. The file size is 1 MB and the control channels are error free. The overall user arrival intensity is fixed to 5 users per second but the fraction of download and upload users is varying. In Figure 8, the user object bit rate averaged over all users (download and upload) is depicted as a function of the fraction of upload users. We can see that when the fraction of upload users is over 40%, FDD outperforms TDD. This is due to shortage of UL resources in the chosen downlink-heavy TDD configurations. With an uplink-heavy configuration, TDD will naturally handle higher UL load.

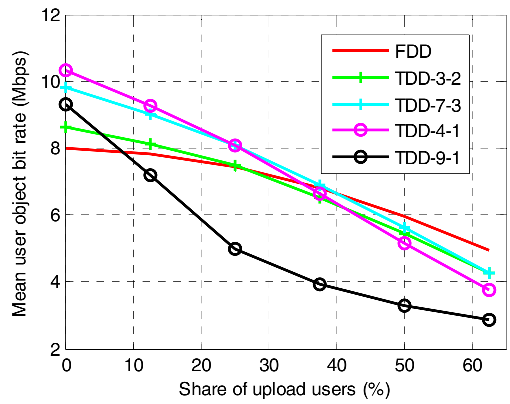

Figure 8. The mean user object bit rate as a function of the share of uploading users. File size 1 MB. The total user arrival rate is fixed.

VII. CONCLUSIONS

In this paper, we studied the performance of Internet access with TDD downlink-heavy asymmetries. By simulations, we show that the performance depends on many factors such as the file size and errors on the control channels. If the file size is small, around 100 kB, longer user plane delays as well as shortage of UL capacity and power does not allow getting the full advantage of the asymmetry of TDD. With a larger file size, the download bit rates can be improved as compared to FDD, even though the improvement is not as significant as the increase in resource allocation would suggest. We have also studied a mixed traffic model with both downloading and uploading users. In the studied scenario, FDD performs better than TDD in terms of the mean user bit rate if the share of uploading users is over 40%.

As discussed in the introduction, due to possible interference between uplink and downlink, the TDD configuration is preferably same over the geographical area. Thus the adaptation to instantaneous traffic changes in the cell is difficult and the large range of TDD configuration cannot always be fully utilized. Instead, it is preferable to select the configuration based on the long term traffic averages.

REFERENCES

[1] D. Astély, E. Dahlman, A. Furuskär, Y. Jading, M. Lindström, and S. Parkvall, LTE: The Evolution of Mobile Broadband, IEEE Communication Magazine, April 2009.

[2] M. Rahman, D. Astély, Link level investigation of ACK/NACK bundling for LTE TDD, IEEE VTC Spring, 2009.

[3] A. Eshete, A. Arcia, D. Ros, and Y. Jiang, Impact of WiMAX Network Assymetry on TCP, . IEEE WCNC, 2009.

[4] C-H Chiang, w. Liao, T. Liu, I. Chan, H-L Chao, Adaptive downlink and uplink channel split ratio determination for TCP based best effort traffic in TDD-based WiMAX networks, IEEE Journal on Selected Areas in Communications, February 2009.

[5] 3GPP TS 36.213, Evolved Universal Terrestrial Radio Access (E-UTRA); Physical Layer Procedures.

[6] 3GPP TS 36.211, Evolved Universal Terrestrial Radio Access (E-UTRA); Physical channels and modulation.

[7] W. Stewens, *TCP/IP illustrated*, Volume 1, Addison-Wesley, 1994.