Measurements in a Laboratory UMTS Network with time-varying Loads and different Admission Control Strategies

Andreas Mäder[1], Bartosz Wagner[1], Tobias Hossfeld[1], Dirk Staehle[1] and
Hans Barth[2]

[1] Department of Distributed Systems, Institute of Computer Science
University of Würzburg, Am Hubland, D-97074 Würzburg
`[maeder,bwagner,hossfeld,staehle]@informatik.uni-wuerzburg.de`
[2] T-Mobile Deutschland GmbH
Landgrabenweg 151, D-53227 Bonn, Germany
`hans.barth@t-mobile.de`

Abstract. Most measurement studies in the literature on UMTS focus on scenarios which are set up as "real" as possible. The focus of this study is the performance and the system behaviour in a nearly ideal laboratory system, such that external influences like other-cell interference or fading are eliminated. We performed experiments on load, throughput and timeliness behaviour with UMTS NodeB hardware of two different brands. The results show different admission control and hardware resource management strategies in high load situations which may impact the performance of applications like voice over IP and the modeling of admission control for dedicated channel radio bearers.

1 Introduction

Mobile wireless networks based on the universal mobile communication system (UMTS) are now in the early operational phase in many countries. The most common variant of UMTS, which is in operational use in west European countries, uses the FDD-WCDMA air interface. The few measurement studies we have found mostly focus on the TCP and IP performance in terms of throughput or round trip times. In [1], the TCP goodput for a FTP-like data transfer of different volumes in a public UMTS network is measured. The authors found out that the first TCP connection is affected by bearer setup signaling which lowers the goodput. Subsequent connections show a good performance. In [2], the performance of P2P-traffic over TCP connections was found to be stable and close to the maximum. The authors of [3] performed TCP measurements in near-ideal network conditions. The study shows stable, low variance round trip times and only small packet loss rates, leading also to a good TCP performance. In [4], the TCP performance for several scenarios is investigated. In the static scenario, the application-level goodput for a MTU size of 1500 bytes is 370 kbps, which is also close to the optimum.

Our work concentrates on the effects of time-varying loads on the system behaviour. We designed several load scenarios where a number of user equipments (UEs) are subsequently connected and/or disconnected to the network to capture the system behaviour under time-dynamic load situations. The measurements were taken in three different setups on equipments of two different hardware vendors. The capacity of the system was limited by the hardware configurations on the one hand and different resource management strategies on the other hand.

The paper is organized as follows: In the next section, the measurement setups are explained. In Sec. 3, we show the measurement results for different scenarios. In Sec. 4, several implications for the modelling of dedicated channel (DCH) bearers are identified. We conclude our paper in Sec. 5.

2 Measurement Setup

The measurements were performed with six test UEs, five Nokia 6650 and one LG from the U8000 series with a Qualcomm chipset. The UEs were connected via USB to the measuring workstation, see Figure 1. At the air interface the UEs are connected via cable to the NodeB. The UEs were separated from each other with a 20dB upstream resistor at the splitter. This splitter was further connected to a duplexer which then was attached to the isolated NodeB. With this configuration the signal strength at the UE receiver was around -70dBm, which can be seen as an ideal case.

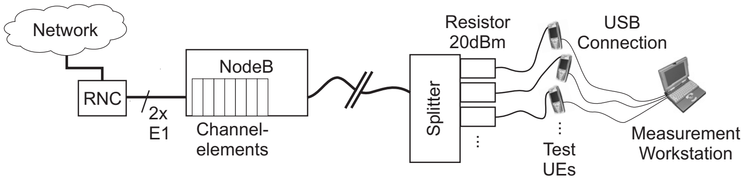

Fig. 1. Measurement Setup

On the UTRAN side, three NodeBs, further NodeB A, B1 and B2, of two different vendors can be selected. (NodeB B1 and B2 came from the same Vendor, but had different software versions). The NodeBs were connected to RNCs of the same vendor with two E1 links with together approximately 2.2Mbps available for user data traffic. The hardware configuration of the channel elements (CE) which are required for the processing of the radio signals are different for the three NodeBs. Together with the capacity of the Iub interface and the air interface the amount of these channel elements is one of the factors which define the capacity of a cell, see e.g. [5]. The number of CEs required for a radio access bearer (RAB) depends on the channel bit rate and on the vendor, but as a general rule, one 16kbps channel needs one CE, which does not necessarily mean that a linear dependency between bit rate and channel elements exist. However, in our case, the hardware capacity of NodeB A was relatively small such that at most two 384kbps and additionaly two 128kbps bearer were supported. The hardware capacities of NodeBs B1 and B2 were sufficient high such that in this case theoretically the OVSF-codes (up to seven with spreading factor 8) and the Iub-interface (2392Kbps for user data traffic) limited the capacity.

The measurements has been performed with the 3GMA-software from Focus Infocom on a Windows 2000 workstation. We used the integrated FTP client (configured for passive mode) for file transfers.

3 Experimental Results

The experiments were performed for three scenarios: *Single UE* (SU), *delayed start increasing load* (DSIL) and *simultaneous start decreasing load* (SSDL). For the first experiment, we measured the throughput of a single UE and compared it with the theoretical value under consideration of the layer 2 overhead. The second and third experiments were designed for the evaluation of time-dynamic effects with different numbers of connected UEs. For this reason, FTP file transfers of different data volumes are subsequently started. We distinguish two cases: Delayed start, i.e. between the transfer begins is a pause of several seconds, and simultaneous start, where all UEs try to connect at the same time. All measurements have been performed several times and we present one representative for each experiment to identify the observed effects.

3.1 Single UE Experiments

The single UE experiments consisted of large file transfer downloads for the 384 kbps and 64 kbps RABs. The file size was 14.4 Mbyte for the first and 2.4 Mbyte for the latter. Although the UMTS protocol stack is fairly complex, the influence of the layer 2 overhead in virtual absence of block errors on the throughput is quite low. The protocol overhead depends on the network configuration, where the most influential layer below TCP is the radio link control (RLC). The protocol header volume is approximated by

$$O_{\mathrm{UMTS}} = H_{\mathrm{PDCP}} + LI + E[N_r] \cdot SEG_{\mathrm{RLC}}. \tag{1}$$

The maximum achievable throughput on application level is then with the inclusion of the TCP/IP header overhead given by

$$B = B^* \cdot \left(1 - \frac{H_{\mathrm{TCPIP}}}{SEG_{\mathrm{MTU}}} - \frac{O_{\mathrm{UMTS}}}{SEG_{\mathrm{MTU}} + O_{\mathrm{UMTS}}}\right), \tag{2}$$

where B^* is the bearer bandwidth calculated from the RLC payload size and the TTI, see e.g. [6]. Note that we neglect the (very low) influence of FTP overhead. The meaning and values of the variables in this formula are shown in Tab. 1; N_r is the number of RLC retransmissions, which we approximate from the reported block error rate. The systems were configured for RLC acknowledged mode.

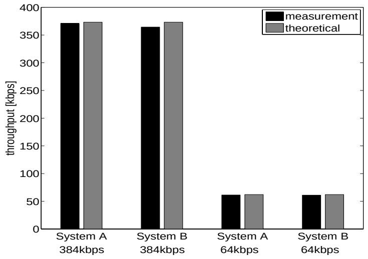

Fig. 2. *SU* scenario: Theoretical bandwidth versus measured throughput

variable	meaning	value [bytes]
$\mathrm{H}_{\mathrm{PDCP}}$	PDCP header	1
LI	length indicator	2
$\mathrm{SEG}_{\mathrm{RLC}}$	RLC segment	40
$\mathrm{H}_{\mathrm{TCPIP}}$	TCP/IP header	40
$\mathrm{SEG}_{\mathrm{MTU}}$	TCP MTU	1500

Table 1. Network parameters and settings

With the values in Tab. 1, the total overhead is only around 2.7% of the brutto channel bandwidth. The measurement results in Fig. 2 show that as expected the measured FTP throughput (black) is very close to the maximum achievable throughput including the protocol overhead, which is shown as gray bars. The only visible difference is for the 384 kbps, where system A performs slightly better than system B, which has an BLER of 0.02% in contrast a BLER of 0 for system A. However, since the BLER is such low it is unlikely that it is the cause of the lower performance due to effects like TCP restransmissions. For 64kbps, neither for system A nor for system B block errors have been observed.

3.2 Experiments with Increasing and Decreasing Load

The first load experiment, the "delayed start increasing load" investigates the system behaviour while the number of active UEs is successively increased. Table 2 shows the starting times and file sizes for the 384 kbps and 64 kbps Bearers. The files were transfered from a standard FTP server (Microsoft IIS 5.0) in binary mode. The pause between the start of a new file transfer was 120 s. The file sizes were chosen such that in the case all UEs get the same bandwidth, all transfers would finish at nearly the same time.

Table 2. Starting times and file sizes for the *delayed start increasing load* experiment

UE	starting time	file size 384 kbps	file size 64 kbps
1	0 s	43.2 Mbyte	7.2 Mbyte
2	120 s	37.4 Mbyte	6.2 Mbyte
3	240 s	31.7 Mbyte	5.3 Mbyte
4	360 s	25.9 Mbyte	4.3 Mbyte
5	480 s	20.2 Mbyte	3.4 Mbyte
6	600 s	14.4 Mbyte	2.4 Mbyte

Figure 3 shows the throughput traces of the individual UEs and the cumulated throughput for system A. The small, upper figures are the individual throughput traces of the UEs, while the lower figure shows the aggregated

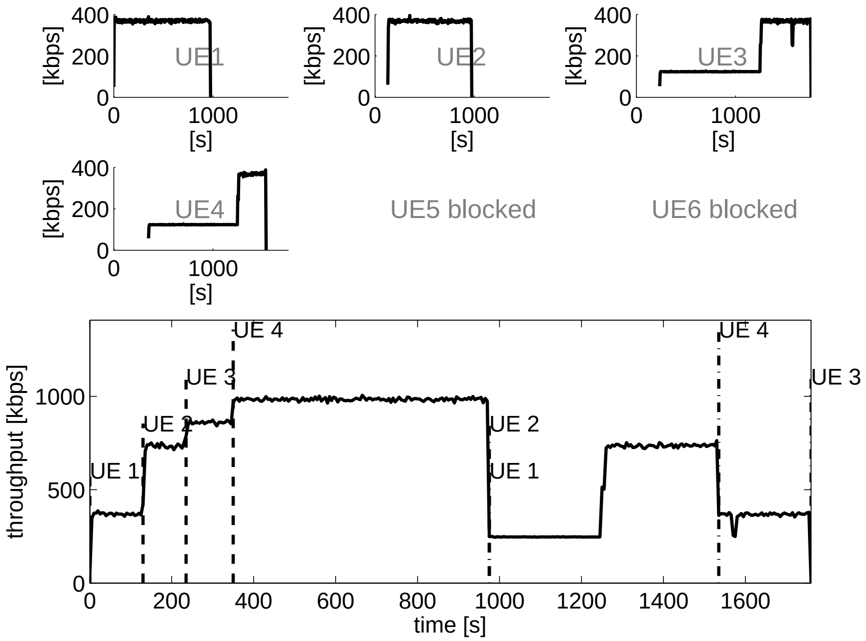

Fig. 3. *DSIL* scenario with 384 kbps bearers on system A: Hardware limits the system capacity

throughput. The vertical dashed and dotted lines indicate the file transfer starting and ending time marks for each UE respectively. 3GMA measurement software uses an event driven logging mechanisms, which means that the measurement points in the UE traces are not synchronized. For this reason, we first established a common time axis and added then the resampled individual throughputs of the UEs to get an approximated aggregated throughput.

It can be seen that in system A, only four UEs could connect to the system. While UE 1 and UE 2 get 384 kbps bearers, the later connecting UE 3 and UE 4 only get 128 kbps bearers. So with in total 4 UEs active, the total throughput is around 1Mbps, as shown in the aggregate trace. The reason for this behaviour are the hardware elements, since the number of required CEs is $2*16+2*4 = 40$, which is exactly the maximum number of available CEs. The link quality of the individual UEs seem to be independent of the number of active UEs, which can be seen on the very smooth development of the traces.

After UE 1 and UE 2 have finished the downloads and disconnected, system A graded the connections of the two remaining UEs up to 384 kbps after a significant amount of time (ca. 250 s).

Figure 4 shows the results of the same experiment for system B1. In contrast to system A, in this case all 6 UEs where able to connect, although UE 6 only gets an 128 kbps bearer, while all others get 384 kbps. Since the NodeB of this system has theoretically a sufficient number of hardware elements, the limitation in this case could be either due to additionally signalling traffic on the Iub

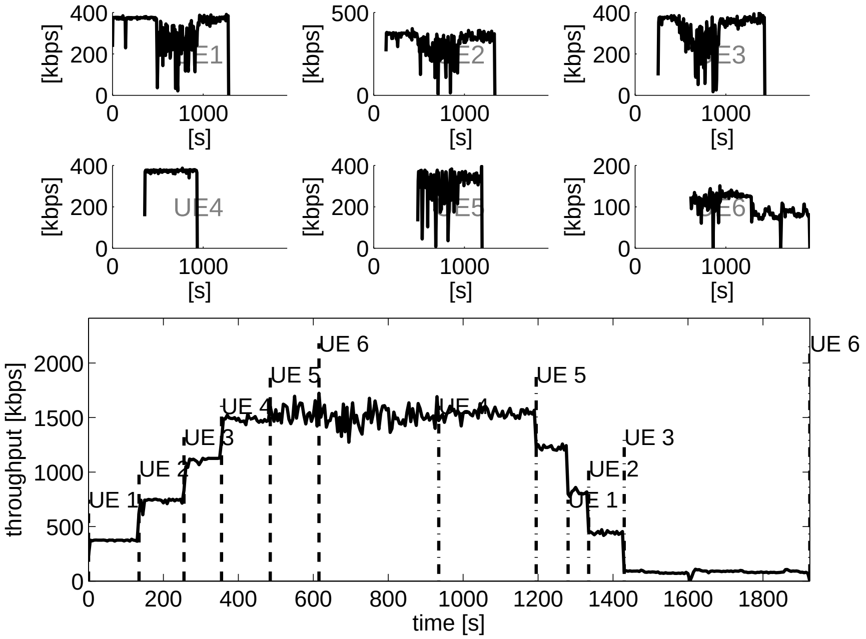

Fig. 4. *DSIL* scenario with 384 kbps bearers on system B1: Jitter grows with number of UEs

interface, or due to a more conservative admission control strategy regarding hardware elements. It can be noted, however, that the traces get increasingly bursty with an increasing number of active UEs which could affect traffic with stringent jitter requirements like voice over IP negatively. While the aggregated throughput follows roughly the formula $4 * 384$ kbps for the first 4 UEs, this is not the case for the last two UEs. With start of the file transfer of UE 5 at ca. 500 s, the traces of all UEs become bursty and a significant decline can be observed. This changes again at approx. 950 s when UE 4 finished the download, which leads to a more smoothly developing of the traces.

For system B2, which had a newer firmware version and more (96) CEs, the effect of performance degradation with a higher number of active user could not be observed. Figure 5 shows the corresponding throughput traces. All 6 UEs were able to connect and also got the requested bandwidth of 384 kbps. The UE traces are very smooth especially in comparison to system B1. The exception is UE 5, which shows periodically dips in the throughput trace. Whether the improvements are subject to the better hardware equipment, to the higher firmware version or to a different network configuration could not be evaluated. The traces show, however, that under ideal conditions the performance of bulk transfer via TCP is close to the optimum in UMTS.

The same experiments have also been performed with 64 kbps bearers. We do not show the traces here since no interactions, blocking or other effects are

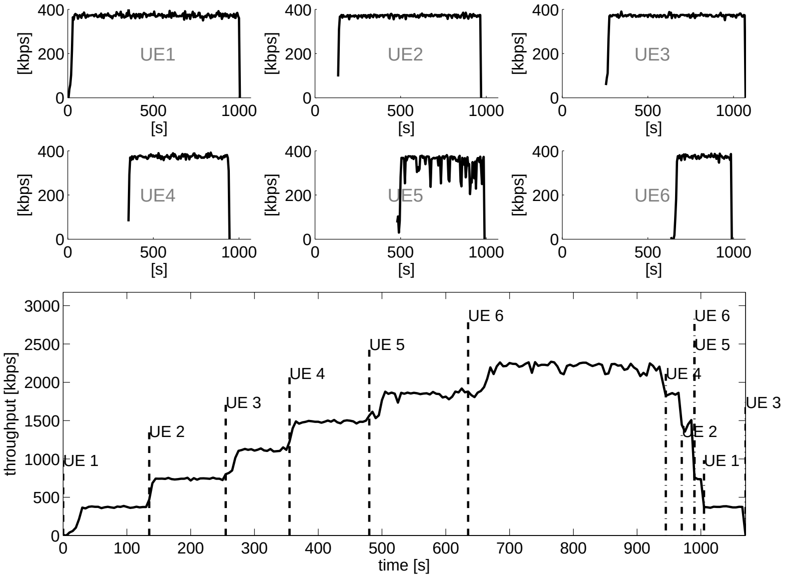

Fig. 5. *DSIL* scenario with 384 kbps bearers on system B2: No blocking, low jitter

observable and the total achieved cumulated throughput corresponds exactly to the number of active file transfers.

The *simultaneous start decreasing load* experiment was designed to observe the behaviour of the system if many UEs try to connect at the (almost) same time and then successively leave the system within a certain time interval. The download file sizes of this experiment correspond therefore to the previous experiment, but the starting time was set to 0 s for all UEs. The results for system A are shown in Fig. 6. What strikes first is the fact that this time, all UEs were able to establish a connection, so the admission control seems to behave different according to the time interval in between new connection attempts fall. It can be further observed that the bit rates of the first 4 UEs jump between 64 kbps and 384 kbps, while the bit rates of the last 2 UEs stay constant. The aggregated trace show that the cumulated bit rate does never exceed 1000 kbps, which is plausible if we remember that the capacity of system A is limited by the channel elements such that two 384 kbps bearer and two 128 kbps bearer can be served at the same time, but not more. We can see at the individual traces, that at most one 384 kbps and three 128 kbps bearer are active, which corresponds to a hardware resource requirement of $3 \cdot 4 + 16 = 28$ CEs.

The first UE has finished it's download at ca. 420s, and after another 50s the RRM tries to maximize the resource utilisation by upgrading the previously downgraded bearer of UE 3 to 384 kbps. The at this time instance remaining UEs, however, are not exceeding 128 kpbs in the remaining time, although two of them, UE 2 and UE 4, already had the full 384 kbps for a short time.

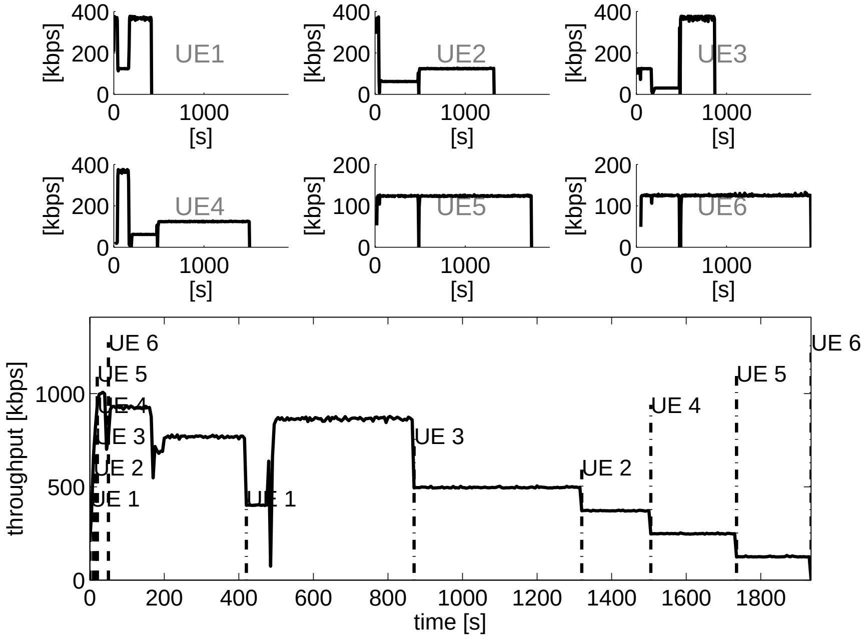

Fig. 6. *SSDL* with 384kbps on system A: Varying bit rates

The results for the systems B1 and B2 are shown in Fig. 7 and Fig. 8, resp. Like in the previous increasing load scenarios the individual throughput of system B1 is lower in the case that several UEs transmit at the same time. Further no irregularities could be observed in other respects, which also holds for system B2. One difference one can observe to the previous experiments where all six UEs could connect is that now in both systems only five UEs were able to do so. In system B2 effectively only 4 UEs were active, since UE 4 lost it's bearer before the file transfer could start, which was a common event throughout the measurement campaign. We do not show the results for the 64 kbps bearers, since in this case as in the previous scenario no particularly effects have been observed.

All in all, the behaviour of the different systems under load is quite intransparent without knowledge of the internal parameters for the network configuration, especially the admission control. We can however state that the two vendors implement different admission control and hardware resource management strategies, which we summarize in the next section.

4 Implications on the Modeling of DCH bearers

The results in the previous section reveal several implications on the modeling of packet switched DCH bearers. While with a low number of active RABs the throughput is very good as shown in the single UE experiment, recall that we observed the following effects for the load experiments:

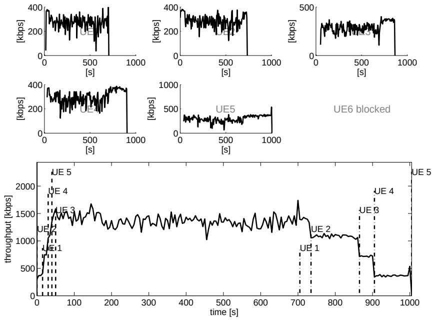

Fig. 7. *SSDL* with 384kbps on system B1: Blocking, high jitter

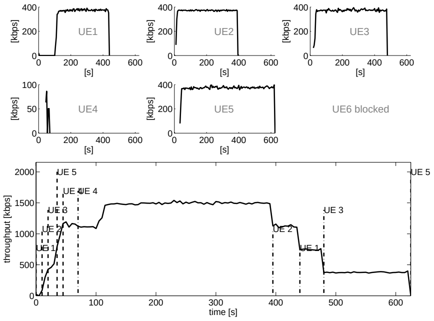

Fig. 8. *SSDL* with 384kbps on system B2: Blocking, low jitter

1. With many (here five 384 kbps DCH bearers) parallel active RABs, the individual DCH throughput may be decrease and the jitter may increase, as been observed in system B1.
2. If the system is close to the maximum capacity, incoming connections may not be blocked as commonly assumed in the most traffic models, but instead downgraded to a bearer with lower bit rate. This effect has been observed in both system A and system B1.
3. Depending on the vendor implementation, an downgraded bearer may be upgraded to it's requested bit rate if the load situation does it permit. This has been observed in the *DSIL* scenario in system A.
4. An incoming connection may force other already persisting connections to downgrade to avoid blocking. This has been observed in the *SSDL* scenario in system A.
5. If more high bit rate connections request are in the system then capacity is available, the system may switch the bit rates of all connections on a quite large time scale, possibly to share the available bandwidth more fairly between the UEs.

The link quality decrease in high load situations is important for non-best effort traffic like voice over ip, since this applications are sensible for jitter. The downgrading of DCH bearers to lower bit rates may also be a problem for real-time traffic, but affects especially the modeling of admission control and therefore the blocking probabilities. Although this effect has only been observed in the context of insufficient hardware elements, it may also used as strategies for other hardware or radio resources. The result is that the blocking probabilities assumed without downgrading are overestimated. While the initial downgrading and possible later upgrading – if sufficient resources are available – does not affect other, already existing connections, this is not the case for the forced downgrading. This strategy does further reduce the blocking probabilities, but for the sake of service degradation. Furthermore, they surely have an influence on the system behaviour which may be significant if the system operates near the maximum capacity and the hardware elements are a limiting factor. Table 3 summarizes the observations.

Table 3. Observed effects

Observation	affects:	system A	system B1	system B2
Blocking/connection losses	block. prob.	X	X	X
Jitter/lower throughput	RT/BE app.		X	
Initial downgrading	block. prob.	X	X	
Forced downgrading	block. prob., RT app.	X		
Upgrading if possible	RT/BE app.	X		
"Fair" bandwidth assignment	RT/BE app.	X		

5 Conclusion

We have presented the results of a measurement campaign in an ideal laboratory UMTS environment. The results show, that the TCP throughput performance on dedicated channels is very close to the optimum under this circumstances. However, with a higher number of connected UEs and parallel data transfers, several vendor specific effects could be observed like lower individual UE throughputs and higher jitter, obviously due to a degradation of the link quality. More significant are the different admission control and rate assignment/rate adaptation strategies. We have observed the downgrading of radio bearers to lower bit rates, also for already established and active connections, and the subsequently upgrading if the resources are available again. This behaviour affects both the blocking probabilities as well as the download times and should be considered in admission control models, simulation models or in the analysis of measurements of real-time traffic like voice over IP or video streaming.

Acknowledgments: The authors want to thank Gabriel Zgunea and Laslo Slawitschka of T-Mobile Deutschland for their very helpful support.

References

1. Pentikousis, K. Palola, M.J.M.P.P.: Active goodput measurements from a public 3G/UMTS network. IEEE Communications Letters **9** (2005) 802–804
2. Hoßfeld, T., Tutschku, K., Andersen, F.U.: Mapping of file-sharing onto mobile environments: Enhancement by UMTS. In: Proc. of MP2P, Kauai Island, Hawaii (2005)
3. Riihijärvi, J., Kohlwes, M., Mähönen, P.: Measurements of TCP Performance over UMTS Networks in Near-Ideal Conditions. In: Proc. of IEEE VTC, Stockholm, Sweden (2005)
4. Catalan, M., Gomez, C., Viamonte, D., Paradells, J., Calveras, A., Barcelo, F.: Performance evaluation of internet access over a precommercial UMTS network. Technical Report TD(05)002, COST 290 (2005)
5. Mäder, A., Staehle, D., Schröder, B.: Dimensioning of Hardware Components in UMTS Networks with Sectorized NodeBs. In: Proc. of ITC19, Bejing, China (2005)
6. 3GPP: 3GPP TS 34.108 V5.4.0 Common test environments for User Equipment (UE) conformance testing. Technical report, 3GPP (2005)