

The Illusion of Being Deterministic – Application-Level Considerations on Delay in 3G HSPA Networks (Work in Progress)

Joachim Fabini¹, Wolfgang Karner², Lukas Wallentin¹,
and Thomas Baumgartner²

¹ Institute of Broadband Communications,
Vienna University of Technology
Favoritenstr. 9/E388, A-1040 Wien
Joachim.Fabini@tuwien.ac.at

² Mobilkom Austria AG, Obere Donaustr. 29, A-1020 Wien

Abstract. The delay experienced by mobile applications in HSPA networks depends to a large extent on highly dynamical global context like, e.g., cell load or algorithms and thresholds governing radio resource scheduling, and on local context like, e.g., user-generated load or load history. These complex uncertainty factors are outside of an applications sphere of influence and result in applications perceiving HSPA link behavior as non-deterministic and non-reproducible.

This paper analyzes accurate round-trip and one-way delay measurement results for three public HSPA networks to demonstrate the high degree of network non-determinism which mobile applications are likely to encounter in practice, particularly significant payload-dependence and halved delay on slightly increased user-generated link load. We argue that current HSPA radio link schedulers, relying on instantaneous user load as decision criterion for channel capacity allocation, neglect real-time application requirements. Cross-layer optimization is one solution which enables deterministic scheduler decisions based on application requirements.

Keywords: 3G, One-way Delay, HSPA, Measurements, User Experience.

1 Introduction

Next Generation Network (NGN) architectures aim at merging circuit switched (CS) voice networks and packet switched (PS) data networks into one common infrastructure. Central to these IP-based NGN architectures is their access-agnostic nature, meaning that applications need not be aware of the underlying access network technology. However, due to fundamentally different behavior of NGN access network technologies, e.g., with respect to delay and transfer rate, applications face difficulties in hiding access technologies from users. Moreover, high transfer capacities supported by technologies like HSPA in uplink and in

downlink direction force mobile operators to increasingly optimize radio resource usage of their networks, meaning that radio resources, e.g., dedicated channels, are allocated to users for a short time interval based on momentary user-generated load. Additional factors which may influence on this resource scheduling optimization include, among others, the user's mobile device type, radio provisioning, user load history, and current cell load generated by all users.

From a single mobile application's perspective the combination of various coding schemes and network scheduling parameters along with the uncertainty factors mentioned earlier yields an unpredictable mobile network behavior with respect to delay and transfer rate. This holds true for static measurements and the more for scenarios involving terminal mobility. Mobile applications, specifically real-time applications relying on deterministic network conditions, must handle delay which varies significantly, depending on the application's own traffic pattern, as well as on other background traffic generated by the mobile user.

Main aim of this paper is to raise NGN application developers' awareness concerning the peculiarity of mobile access networks. Specifically, the measurement which we present in this paper demonstrate that it is grossly negligent to infer from application tests in core- or fixed access networks onto application behavior in mobile access networks. The measurement results presented in this paper question the relevance of many scientific publications which rely on simplifying assumptions like, e.g., "the round-trip delay of UMTS is 150 ms". Our measurement results demonstrate that delay in today's mobile access networks is payload dependent and can effectively halve on slight increase in user-generated network load. Mobile real-time applications must, therefore, be prepared and tested to appropriately handle this special and apparent paradox behavior.

1.1 Related Work

There are only few publications analyzing 2.5G and 3G packet-switched performance aspects based on real network measurements. The authors of [1] use ICMP echo messages to infer on conversational and delay aspects of GPRS and EDGE networks. [2] analyzes impact of network load on TCP RTT in live GPRS and UMTS networks, whereas [3] proposes the use of passive monitoring for optimizing mobile network performance. In recent work, [4] relies on synchronized packet-level captures in the core network to infer on one-way delay. In [5] we have presented early GPRS and UMTS network measurement results as well as the impact of application-layer protocols like HTTP, POP3, SMTP or FTP on user-perceived performance, while [6], [7] and [8] present previous results of payload-dependent access network measurement and access network emulation results.

This paper's main contribution is to raise mobile application developer's awareness concerning the huge dependency of delay and, generally, of the behavioral determinism in mobile networks on user-generated traffic load and traffic patterns. The measurement results presented in this paper quantify for the first time application-perceived uncertainty of HSPA networks based on empirical measurement results and comparison of several public HSPA networks.

1.2 Structure of This Paper

In this paper we present payload-dependent delay measurement results for round-trip-, uplink- and downlink delay in public HSPA networks to point out the impact of resource allocation in mobile networks on real-time applications expecting deterministic traffic conditions. The remainder of this paper is structured as follows: Section 2 proposes a measurement methodology for round-trip- and one-way delay assessment based on round-trip delay decomposition into request and reply, followed by Section 3 which presents measurement results for this methodology in three distinct public HSPA networks depending on measurement configuration. In Section 4 we discuss the impact of our measurement results on NGN applications and NGN technologies, present methodological implications and conclude with an outlook on future work.

2 Measurement Methodology

The proposed measurement methodology for assessing HSPA network delay relies on two randomness factors, specifically on Internet Control Message Protocol (ICMP) packets having *random payload size* which are sent at *random start times*. Whereas random payload sizes distribute the impact of temporal network overload over the whole measurement payload space, start time randomness as proposed by the IP Performance Metrics (IPPM) framework in RFC 2330 eliminates correlations between send time and periodical network behavior.

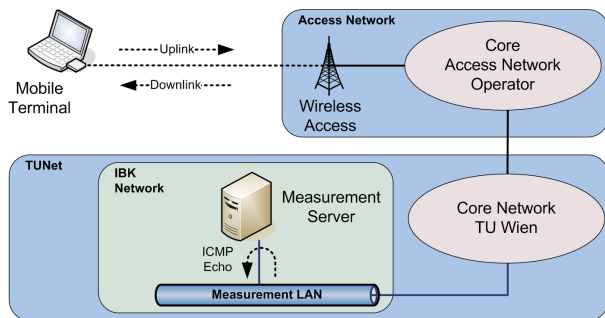


Fig. 1. Generic setup for automated round-trip delay measurement

The measurement setup in Figure 1 depicts the mobile client, a laptop computer connected to public HSPA networks using a Huawei E870 HSPA modem, and the measurement server which is connected to the Institute of Broadband Communication’s switched Ethernet network. Depending on the measurement methodology, ICMP messages are sent either by the mobile client and reflected by the measurement server or vice-versa. For one-way measurements the clocks of mobile client and measurement server are continuously and accurately synchronized against a common global time base, whereas tcpdump [9] processes

log incoming and outgoing packets along with their headers and timestamps on both hosts. Correlating the information stored in the two tcpdump trace files yields accurate uplink and downlink delay values, the precision being better than 0.5 ms. However, only ICMP *request* messages have been used for one-way delay computation.

The mobile terminal's position was fixed during all measurements, this static measurement regime minimizing the impact of variable measurement parameters like, e.g., of varying radio conditions and handovers because of terminal motion onto measurement results.

3 Measurement Results

Any single dot in the diagrams presented in this section represents the value of one ICMP round-trip- or one-way delay measurement, one diagram displaying approximatively 10,000 single measurements. Measurement ICMP payload size is selected uniformly distributed between 12 bytes and 1450 bytes, the start time of any ICMP-packet being chosen randomly between 50 ms and 500 ms following the last transmission.

3.1 Round-Trip Delay Measurement Results

Figure 2 depicts round-trip delay measurement results using the random measurement methodology presented in Section 2 for three public Austrian HSPA networks labeled as Operator A to Operator C.

Figures 2(a) to 2(c) illustrate that payload-dependent ICMP response patterns of the three measured public HSPA networks differ significantly. Due to the random measurement methodology, temporary changes in coding and scheduling influence on a wide range of payload sizes. All three diagrams show clustered groups of delay values which can be approximated by straight lines, some of them having a visible block structure. Figure 2(d) depicts the groups to which in the following we refer to as *delay lines*. Differentiating factors, which uniquely identify any delay line, include the payload interval covered by the delay line, its block size, delay value increase between subsequent blocks, and the delay line's initial delay offset.

For instance Operator A's diagram shows three main delay lines, as depicted by straight lines and corresponding labels in Figure 2(d). The first one, *Delay line 1*, positioned at the top left is composed of 40-byte wide payload blocks. It starts at an initial delay offset of approximatively 75 ms for 12 bytes ICMP payload size and ends at a delay value of 200 ms for 439 bytes payload size. The second delay line, *Delay line 2* starts at a delay value of 160 ms for 440 bytes payload but uses a block size of 160 bytes. The third and lowest one of these three main delay groups, labeled *Delay line 3* spans the entire payload range. It starts at a delay value of 75 ms for 12 bytes ICMP payload and terminates at 250 ms delay for 1450 bytes payload. Sparse measurement results below *Delay line 3* indicate the existence of a fourth, dotted, high-performance and low-delay *Delay line 4* which, however, is only rarely used for the tested load profile.

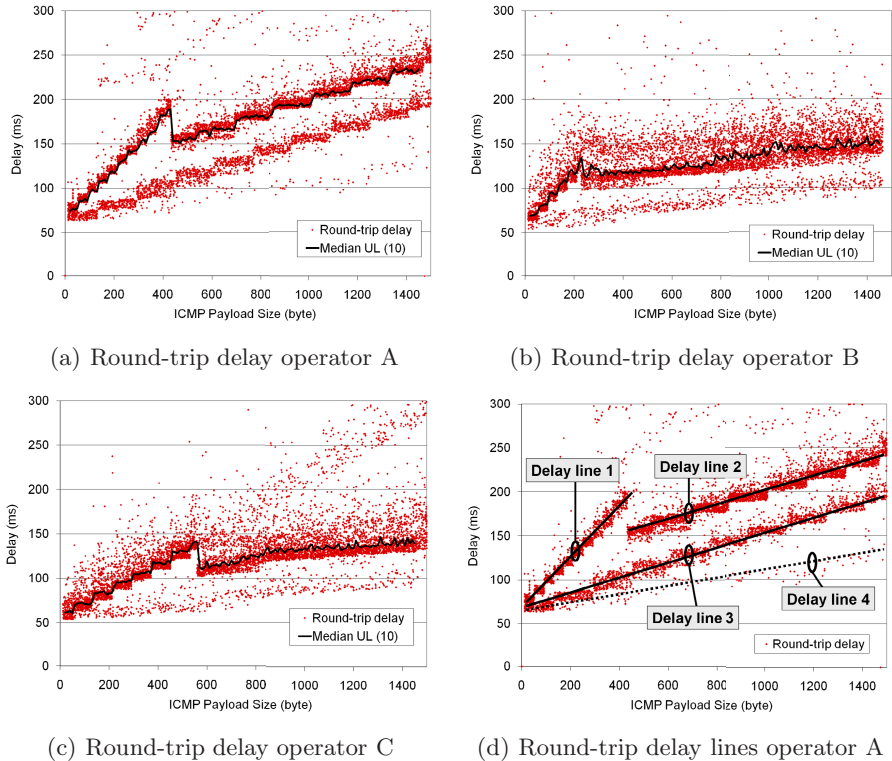


Fig. 2. HSPA Round-trip delay measurement results (measurement load only)

Due to the average function’s sensitivity to outliers we have selected the median as representative function for the “common” network behavior. To compute the diagram’s median curve, labeled in the diagram’s legend as *Median(10)*, we have segmented the entire measurement payload range into sets which cluster 10 adjacent payload sizes. The diagram displays the median delay values for these sets, positioned in the center of the respective 10 payload values, interconnecting these points by straight lines.


The round-trip delay diagrams for Operator B in Figure 2(b) and for Operator C in Figure 2(c) exhibit different patterns, although the basic diagram structure is similar to Operator A. All diagrams start with steeply increasing low-blocksize delay lines and continue with larger block sizes. However, the “typical” delay at large payload sizes is significantly lower for Operator B and Operator C (150 ms) than for Operator A (200 ms to 250 ms).

Fundamental to the diagram discussion is the statement that the slope value of these delay lines equals the inverse of the channel’s transfer capacity. Therefore, delay lines having a high slope value represent low-bandwidth channels whereas low slope values correspond to high-bandwidth channels. This finding raises an additional question, namely how can it happen that the diagrams in Figure 2

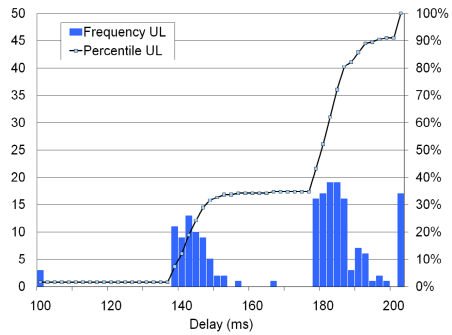
show two delay lines having equal transfer capacity but a delay offset of more than 50 ms. This analysis requires decomposition of any HSPA round-trip delay measurement value in its one-way delay components, which we present exemplary for Operator A.

3.2 One-Way Delay Measurement Results

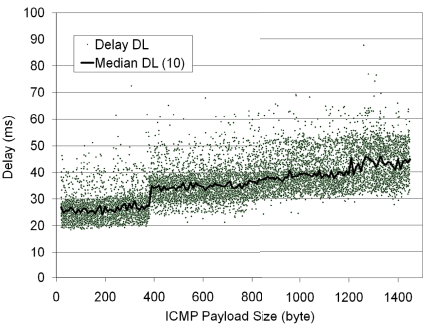
Figure 3 depicts uplink delay and downlink delay measurement results and histograms for Operator A. The diagrams evidence that the HSPA uplink originates the multi-layering which we have noticed in the round-trip delay diagram in Figure 2(a). Therefore we can infer that the delay lines can be mapped to specific HSUPA uplink grants. The first (top leftmost) delay line uses E-TFCI 4 (Enhanced Dedicated Transport Channel Transport Format Combination Indicator) carrying 372 bits of RLC payload, whereas the lower delay line can be mapped to E-TFCI 15 which supports 1362 bits of payload for 10 ms Transmit Time Interval (TTI) according to Annex B.4 - Table 1 of TS 25.321 [10]. E-TFCI



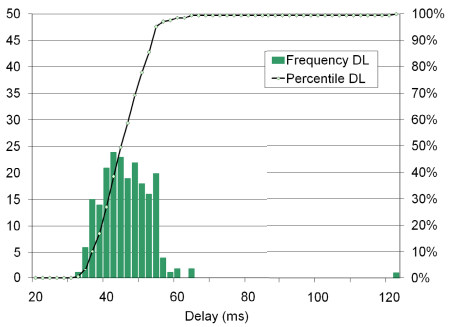
(a) Uplink delay



(b) Uplink histogram (1420-1469 bytes payload)



(c) Downlink delay



(d) Downlink histogram (1420-1469 bytes payload)

Fig. 3. HSPA one-way delay measurement result and histogram Operator A

payload includes higher layer headers and therefore is larger than the IP layer block sizes which have been presented in section 3.1.


Our interpretation of the diagram in Figure 3(a) is that the mobile network grants E-TFCI 15 whenever a packet with payload larger than 440 bytes is to be sent while the current grant is E-TFCI 4. In this case the terminal starts transmitting using E-TFCI 4 while switching to E-TFCI 15 and continues using E-TFCI 15 afterwards. This “slow-start” causes the delay- and payload offset of the upper right E-TFCI 15-delay-line relative to the lower E-TFCI 15 layer. We also note that the terminal never obtains a higher grant if it uses E-TFCI 4 and a packet with payload size of up to 440 bytes must be sent.

The histogram analysis of uplink delay measurement results reveals that the upper left and right delay lines account for the majority of all measurement values. Figure 3(b) depicts exemplary the uplink delay histogram for the payload interval between 1420 and 1469 bytes. The percentile curve points out that approximately 30% of all measurement values belong to the lower delay line and 55% to the upper line, this ratio being representative for the entire measured payload interval. Knowing that the delay between two subsequent measurement packets is chosen randomly between 50 ms and 500 ms one likely explanation for this distribution is that Operator A optimizes the network aggressively, downgrading terminal grants from E-TFCI 15 to E-TFCI 4 shortly after the packet has been transmitted. In a less aggressive optimization case, i.e., when the E-TFCI 15 grant persists until the next packet is ready to send, the majority of delay values would have been located in the lower (E-TFCI 15) delay layer.

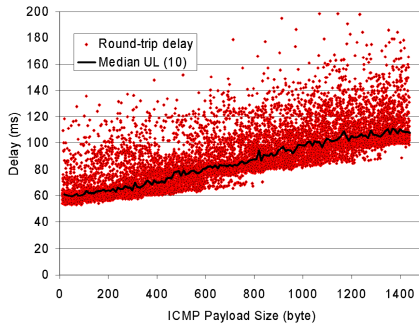
Explanations analogous to the ones for Operator A, though with other grants and payload limits, apply to the diagrams for Operator B and Operator C which we have presented in Figures 2(b) and 2(c). The low delay at large payload sizes indicates that both operators use higher grants than E-TFCI 15 for large packets. Operator B uses the same E-TFCI 4 grant at low packet sizes but a significantly smaller threshold than Operator A (approximately 230 byte) which triggers the switch to a higher grant. Finally, Operator C uses a higher grant than Operator A and Operator B (most likely E-TFCI 7 or E-TFCI 8) for packets up to 560 bytes in size.

3.3 Round-Trip Delay Measurement Results (Background Load)

Assuming that user-generated load – i.e., the load due to measurement packets in the case of measurements – biases on the transfer capacity which mobile networks allocate to users we have repeated the HSPA measurements with deterministic background load. The modified measurement methodology uses identical setup and methodology as presented in section 2 but adds two constant bit rate (CBR) flows between the mobile terminal and the measurement server, one flow loading the uplink and one the downlink. Main aim of this user load is to prevent the network from withdrawing higher grants suddenly after a measurement packet has been sent while minimizing collisions with measurement packets.



(a) Round-trip delay operator A (load) (b) Round-trip delay operator B (load)



(c) Round-trip delay operator C (load)

Fig. 4. HSPA Round-trip delay measurement results (deterministic background load)

Figure 4 presents round-trip delay measurement results for the deterministic background load scenario. Uplink and downlink have been loaded using two periodic 50 kbit/s data rate 100 byte payload UDP flows generated by IPerf [11], one in each direction. Although the number of outliers in the background-loaded measurement diagrams increases visibly when compared to the original diagrams in Figure 2, the delay pattern in the background load case in Figure 4 is more deterministic. Specifically, the multi-layering clearly visible in Figure 2(a) disappears when measuring using deterministic background load.

Main result of this new measurement methodology is that the median round-trip delay value *decreases* significantly for all three operators throughout the entire measured payload interval. Most prominent, for Operator A the delay difference is substantial when comparing Figure 2(a) against 4(a). The median round-trip delay value at large payload sizes effectively halves from 225 ms to 110 ms when adding background load to the network. For Operator B and Operator C the decrease in round-trip delay is smaller, though still significant. Round-trip delay drops from 150 ms to 120 ms for Operator B and from 140 ms to 110 ms for Operator C, taking into account the median values at large payload sizes.

3.4 One-Way Delay Measurement Results (Background Load)

Analogous to the non-loaded measurements we have decomposed round-trip delay measurement values into their uplink and downlink delay components, enabling accurate analysis of performance improvements. All one-way diagrams which we have presented for the non-loaded case in Figure 3 are shown in Figure 5 for the background-load measurement scenario.

A comparison of uplink delay diagrams and their median delay curves in Figures 3(a) and 5(a) illustrates that adding background load to measurements decreases the uplink delay significantly. The histograms showing delay frequency for large payload sizes (1420-1459 bytes) in Figures 3(b) and 5(b) confirm this finding. Whereas in Figure 3(b) the percentile curve starts to increase at 140 ms and reaches 90% close to a delay value of 200 ms, the percentile curve in Figure 5(b) starts at 50 ms and crosses 90% at an uplink delay value of 110 ms. However, the uplink jitter in the background load case is significant and affects more than 12% of all measurement values for large payload sizes.

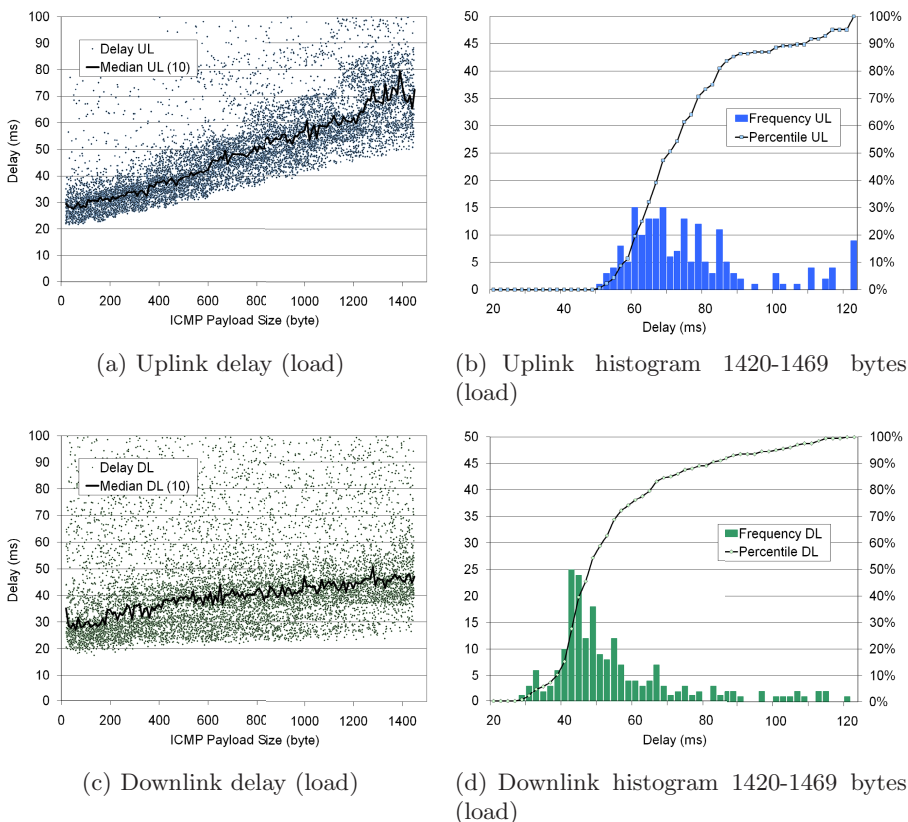


Fig. 5. HSPA one-way delay measurement results Operator A (background load)

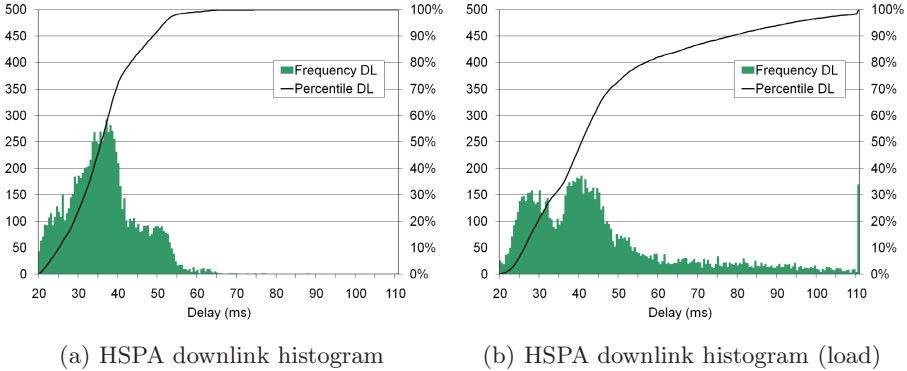


Fig. 6. HSPA downlink delay histograms for Operator A

Non-regarding this, the difference in uplink delay median values for the payload interval mentioned above is substantial: adding background load decreases the uplink delay median value from 182 ms in the non-loaded case to 72 ms in the loaded case. However, the measurement results indicate that background load does not improve HSPA downlink performance, on the contrary. A comparison of HSPA downlink delay diagrams shown in Figure 3(c) for the non-loaded case and in Figure 5(c) for the background load case points out that the median value curve is affected mainly for low payload values up to 400 bytes, for which median delay increases by 5 ms to 10 ms. Detrimental is the huge jitter which in the background load scenario affects more than 15% of all downlink packets and spreads over the entire measured payload interval. This finding is confirmed by the evolution of percentile curves in Figure 6(a) for the non-loaded case and Figure 6(b) for the background load scenario, the diagrams depicting downlink delay histograms for the entire measured payload interval.

4 Conclusions and Future Work

From an application perspective HSPA networks behave non deterministic. The measurement results presented in this paper demonstrate that the HSUPA uplink delay *decreases* significantly with *increasing* uplink load and associated higher grants in terms of radio resources. This radio link scheduling strategy is at least questionable from a fairness perspective, resulting in delay penalties for applications which manage their resources economically. However, additional load does not have a positive effect on downlink delay. Inherent reason for this asymmetric load influence on performance is the radio link scheduler’s location in the NodeB. Due to the NodeB’s low-latency fixed network connectivity towards RNC and core network the scheduler can react almost instantaneously by allocating appropriate radio resources for packets in the *downlink*, whereas for *uplink* transmission it is the mobile terminal’s task to request uplink radio resource allocation according to the momentary load. Uplink resource requests therefore

require an additional radio access round-trip delay until the mobile terminal receives the notification whether a specific request can be fulfilled or not fulfilled, depending on the radio scheduler policy and momentary context.

Underlying reason to applications perceiving radio link scheduler decisions as non-deterministic is eventually missing context information in both, the mobile terminal and in the radio link scheduler. Mobile applications which do possess a priori or momentary context information – e.g., having configured a specific real-time traffic profile, stating what amount of data is to be sent at which point in time – currently have no means to communicate this context information to lower layers or the radio link scheduler. In other words, data at IP-layer is considered by lower layers and by the radio link scheduler as being opaque, all packets having same priority and importance. It is therefore legitimate to question the traditional OSI layering model which does not foresee cross-layer information exchange. We expect context propagation in mobile access networks to be implemented using cross-layer optimization in the near future, marking IP packets by appropriate priorities in terms of delay, loss, and priority requirements, to enable and support appropriate radio link scheduler decisions.

Our measurements demonstrate that randomness in terms of payload and packet send time can unleash mobile operator specific configuration parameters and thresholds. Active measurements can therefore serve as a simple, straightforward tool to optimize own mobile networks or to reverse-engineer highly sensitive configuration parameters of competing network operators.

Seen from standardization perspective, the measurement results presented in this paper question the meaningfulness of compression techniques proposed for NGNs like, e.g., Signaling Compression [12]. The results demonstrate that these techniques, which are supposed to decrease delay in narrow-band mobile networks and are mandatory part of NGNs like, e.g., the IP Multimedia Subsystem (IMS), might show the contrary effect of increasing delay because of their reduction of user-generated network load.

Even more important, this paper's results point out a major drawback of existing standards like the ETSI Technical Specification TS 102 250 Part 1 to Part 6, which target the definition of fair metrics for mobile network measurements but completely disregard measurement result dependence on measurement methodology. Facing the huge influence of measurement methodology, specifically of traffic profiles, on delay results, we conclude that these standards must be substantially revised to fulfill basic fairness requirements. As future work we therefore plan to contact ETSI and contribute our findings in order to enhance TS 102 250 with respect to measurement methodology specification.

Acknowledgment

The authors would like to thank mobilkom austria AG for technical and financial support of this work. The views expressed in this paper are those of the authors and do not necessarily reflect the views of mobilkom austria AG.

References

1. Ball, C., Masseroni, C., Trivisonno, R.: Introducing 3G like conversational services in GERAN packet data networks. In: IEEE 61st Vehicular Technology Conference (VTC 2005), Stockholm, Sweden, May 2005, pp. 2186–2191 (Spring 2005)
2. Vacirca, F., Ricciato, F., Pilz, R.: Large-scale RTT measurements from an operational UMTS/GPRS network. In: Proceedings of the First International Conference on Wireless Internet (WICON 2005), Budapest, Hungary, July 2005, pp. 190–197 (2005)
3. Ricciato, F.: Traffic monitoring and analysis for the optimization of a 3G network. *IEEE Wireless Communications* 13(6), 42–49 (2006)
4. Ricciato, F., Hasenleithner, E., Romirer-Maierhofer, P.: Traffic analysis at short time-scales: an empirical case study from a 3G cellular network. *IEEE Transactions on Network and Service Management* 5(1), 11–21 (2008)
5. Reichl, P., Jordan, N., Fabini, J., et al.: Wireless Inter-System Quality-of-Service: A Practical Performance Analysis of 3G and Beyond. In: *Kommunikation in Verteilten Systemen (KiVS 2005)*, Kaiserslautern, Germany, March 2005, pp. 230–241 (2005)
6. Fabini, J., Reichl, P., Egger, C., et al.: Generic Access Network Emulation for NGN Testbeds. In: *4th International Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities (TridentCom 2008)*, Innsbruck, Austria, March 2008, p. 10 (2008)
7. Fabini, J., Reichl, P., Poropatich, A.: Measurement-Based Modeling of NGN Access Networks from an Application Perspective. In: *Proceedings of the 14th GI/ITG Conference on Measurement, Modeling, and Evaluation of Computer and Communication Systems (MMB 2008)*, Dortmund, Germany, April 2008, pp. 45–59 (2008)
8. Fabini, J., Reichl, P., Poropatich, A.: A Generic Approach to Access Network Modeling for Next Generation Network Applications. In: *Proceedings of the 4th International Conference on Networking and Services (ICNS 2008)*, Gosier, Guadeloupe, March 2008, pp. 254–260 (2008)
9. Lawrence Berkeley National Laboratory, University of California, Berkeley, CA: TCPDump Monitoring Tool (December 2007), <http://www.tcpdump.org/>
10. 3GPP: 3rd Generation Partnership Project: 3GPP TS 25.321, Medium Access Control (MAC) protocol spec., version 6.15.0 (March 2008), <http://www.3gpp.org/>
11. National Laboratory for Applied Network Research (NLANR): IPerf Performance Measurement Tool (December 2007), <http://dast.nlanr.net/Projects/Iperf/>
12. Price, R., Bormann, C., Christofferson, J., et al.: IETF standards track RFC 3320: Signaling Compression (SigComp) (January 2003)