

Explorative Analysis of One-way Delays in a Mobile 3G Network

Internal Technical Report
Telecommunications Research Center Vienna (ftw.)

Date

May 2008

Authors

Peter Romirer-Maierhofer and Fabio Ricciato

Abstract

In this paper we investigate the dynamics of one-way delays in an operational mobile core network. Our final goal is to develop anomaly detection schemes for the packet delay process in order to reveal network and equipment problems. This requires a preliminary exploration of the delay process in the core network, which we undertake in this study.

We present one-way delay measurements extracted from passive traces captured at an operational GPRS/UMTS network. We find that queuing is the only source of delay at GGSNs, while for SGSNs the mobility and flow-control further complicate the characterization of the “normal” delay behavior. Moreover, the presence of unwanted traffic has an impact on the delay statistics and should be taken into account. We find that information about actual bandwidth conditions in the Radio Network may be inferred by investigating packet delays in the core network. Our explorative results are promising about the possibility of leveraging one-way delay measurement for troubleshooting in such networks.

I. INTRODUCTION

Third-generation (3G) mobile networks play an increasingly important role in the telecommunication market. Operators of a 3G infrastructure must cope with a usage environment that is rapidly evolving, as new applications emerge and terminal capabilities increase continuously. In such a scenario, it is highly desirable to automatically detect network problems such as equipment malfunctioning, point of congestion, misconfiguration and alike. One natural approach to the problem is to measure some key performance parameters on the data plane, e.g. one-way delays, and to raise an alarm whenever a significant deviation occurs from the “typical” behavior observed in the past. This approach underlies two fundamental assumptions: *i*) that the behavior of the process under normal conditions is highly predictable, and *ii*) that anomalous phenomena (e.g., malfunctioning) generate appreciable deviations on the observed process. Therefore, the practical implementation of this approach requires a preliminary understanding of the “normal” patterns, to be achieved by an explorative analysis of a sample dataset.

In this work we carry out the explorative analysis of the single-hop delay process at the main elements of a 3G Core Network. We find that the SGSN delays are influenced by user mobility and flow control, in addition to queuing. Searching for simple indicator signals, we find that the fraction of packets exceeding some fixed delay threshold should be preferred to delay percentiles. As a valuable side effect, we identify new possible research directions for inferring the global condition of the radio network from the delay process at the SGSN. To the best of our knowledge this is the first work reporting on detailed measurements about the SGSN delay process.

II. RELATED WORKS

The initial study on single-hop one-way delays in an operational IP network was performed by Papagiannaki et al. in [1]. There the methodology based on IP header hashing and matching was presented for the first time. The analysis was then extended to the end-to-end delays across a geographical network in [2]. Those studies were performed on the Sprint backbone, while our study focuses on the core section of a cellular network, where the delay process is not only shaped by queuing, but also by user mobility and flow control procedures. Some previous works have studied the delay process in 3G mobile networks based on passive traces, e.g., [3]–[5]. They all focused on the analysis of Round-Trip-Times (RTT), estimated from TCP DATA/ACK or SYN/SYNACK pairs captured at a single monitoring point on Gi. Other works have resorted to active measurements to evaluate one-way end-to-end delays, e.g., [6]–[8]. In general, it is not possible to exploit the RTTs or the end-to-end delays to infer the internal dynamics of the Core Network. This is because the delay on the radio link is the dominant component in the total delay budget, and it exhibits a high degree of variability due to the changing radio conditions. Therefore, the fine-grain dynamics taking place within the wired section can not be observed with those approaches.

III. MEASUREMENT SETTING

A. Network Overview

Here we give a short overview of the structure of a 3G cellular network, for more details refer to [10]. In the network under study four different access schemes are available depending on the geographical location of the Mobile Station (MS) and its terminal capabilities: GPRS, EDGE, UMTS and HSDPA [10]. The MSs are connected via radio link to the antennas. A set of geographically neighboring antennas is connected to the Base Station Controller (BSC) in GPRS/EDGE and the Radio Network Controller (RNC) in UMTS/HSDPA. These are then connected to a set of so-called Serving GPRS Support Nodes (SGSN) via the Gb (for GPRS BSC) and IuPS (for UMTS RNC) interface. The overall set of antennas, BSC/RNC and the links to the SGSNs constitute the Radio Access Network (RAN). The primary role of the SGSN is to perform mobility management, which involves frequent signaling exchanges with the MSs. The data plane traffic collected by the SGSN is concentrated within a small set of so-called Gateway GPRS Support Nodes (GGSN), which act as IP-layer gateways for the user traffic. The set of SGSNs and GGSNs is interconnected by a wide-area IP network that will be referred to as the “Gn network” following the terminology of 3GPP specifications (“Gn interface”). The GGSNs are connected to the public Internet via the Gi interface.

In this work we consider traffic of MSs connected via the GPRS/EDGE RAN. The protocol stack of the Gb mode is sketched in Fig. 1. The SGSN is directly connected to the MSs via the LLC protocol which provides link layer functions such as flow control and error correction. The Base Station System GPRS Protocol (BSSGP) is used for adapting the transmission rate between the SGSN and the Base Station System (BSS) dynamically according to the receive buffer utilization at the BSS. Radio bandwidth at the air interface (“Um interface”) is multiplexed among connected MSs via the Radio Link Control (RLC) protocol by transmitting short RLC blocks in Time Division Multiple Access (TDMA) fashion.

B. Monitoring System

The present work is based on packet-level traces captured in the operational network of a mobile provider in Austria. The monitoring system was developed in a past project [11]. The capture cards are Endace DAG [12] with GPS synchronization. For privacy reasons, we store only the packet headers up to the transport layer, i.e., payload is stripped away. All packets are captured and no sampling is implemented. Similarly to timestamps, a unique MS identifier is stored as an additional information

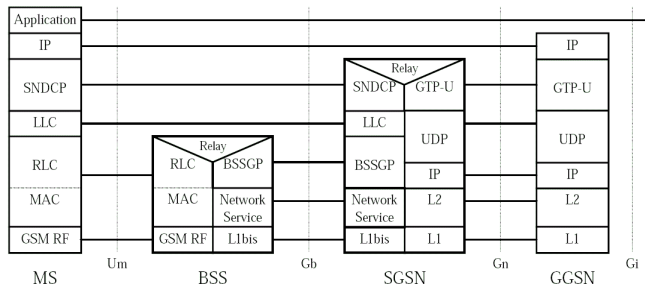


Fig. 1: User plane for Gb mode

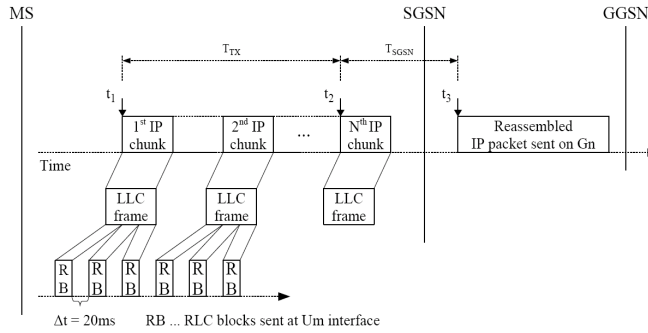


Fig. 2: Framing of GPRS traffic

label for each frame. To protect the user privacy, the MS identifiers are chosen as arbitrary strings, decoupled from the real MS identity.

C. Delay Computation

The one-way delay was extracted in post-processing with a similar methodology as in [1]. In our work, each IP packet, excluding some selected header fields (e.g., Time-To-Live), is hashed into a string of length $N = 128$ via the MD5 function [13] (instead CRC-32 was used in [1]). This guarantees a negligible collision probability. For each processed IP packet, the hashed string, the arrival timestamp, and both IP addresses are written to a binary file. The same procedure is repeated for the traces captured at every interface. In order to extract the delay samples from A to B, we seek for string matches in the two corresponding binary files, and take the difference of the respective timestamps (we will elaborate on the timestamp definition in § III-D). While this scheme is conceptually simple, we reported a number of practical issues that must be taken into account [9].

We can combine delay computation and packet sampling - denote by $0 < r_s \leq 1$ the sampling rate - or alternatively compute the delay for *all* packets. Our monitoring system is capable to run online sampling and delay computation with tunable value of r_s . However for offline analysis we often resort to complete delay computation ($r_s = 1$). The delay samples are then aggregated into $N = 300$ exponentially spaced bins between 0 and 60 seconds in order to summarize their distribution. The bin boundary is given by the following formula:

$$b(n) = e^{(n-\delta)\cdot\alpha} \quad n = 0..N - 1, \quad (1)$$

with the parameter setting $\alpha = 0.07$ and $\delta = 240$.

D. Which Timestamp to Use?

We now elaborate on the arrival time considered for the delay computation. We focus on packets at the IP layer. When fragmentation occurs at lower layers (e.g., at LLC layer on the Gb links), we can choose the arrival timestamp of the first or last fragment. Fig. 2 shows the framing of an IP packet transmitted from a GPRS-connected MS to the GGSN (uplink). At the Gb interface IP packets are encapsulated into LLC frames. Let L denote the IP packet length and S the maximum number of IP bytes carried per LLC frame. If $L > S$, the IP packet has to be fragmented into two or more LLC frames. Hence, we can choose either the arrival time of the first LLC frame (t_1 in Fig. 2) or the last one (t_2) for calculating the delay between the Gb and Gn interface. We initially chose a delay definition of $d' = t_3 - t_1$, which in fact includes the time T_{TX} between the arrival of the first and the last LLC frame of an IP packet plus the time T_{SGSN} needed for reassembling and forwarding the IP packet inside the SGSN. This definition conforms to the strategy used by the DAG cards (see e.g. [14, p.102-112]), where the timestamp marks the *beginning* of the frame reception. In the second part we turned to the definition $d'' = t_3 - t_2$ because

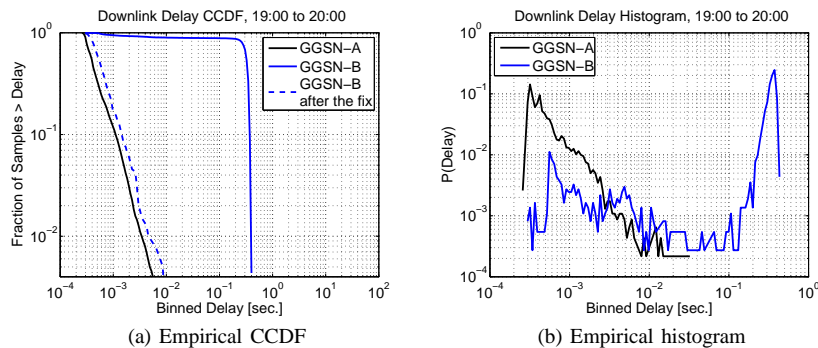


Fig. 3: Delay distribution at two GGSNs

of the reasons explained in § V. The noted delay definitions are also applicable to downlink traffic, which may be fragmented at the Gb interface as well.

E. Input Datasets

The presented work is based on three different packet traces. The first is a 1 hour trace extracted from the Gi and the Gn interface of two GGSNs in November 2007, at peak hour. The second trace was collected at the Gn and Gb interfaces of a GPRS-capable SGSN (“SGSN-A”) during a full day in November 2007. Similarly, the third trace is from another GPRS-capable SGSN (“SGSN-B”) in December 2007. Each trace includes both data and signaling traffic, which enables exploring the correlation between the delay process on the user plane and the signaling traffic.

IV. GGSN DELAYS

We start presenting the sampled one-way delay measured between the Gi and the Gn interface of two GGSNs denoted by “GGSN-A” and “GGSN-B”. In Fig. 3 we plot the Empirical Cumulative Distribution Function (ECCDF) along with an empirical histogram for the interval 7-8 pm, representing a period of high load. At GGSN-A 90% of the delay samples take very low values, below 1 ms (ref. Fig. 3a), while less than 1% of the delay values are higher than 4 ms. The histogram at GGSN-A (ref. Fig. 3b) shows an absolute peak at only around 0.3 ms. However, the delay distribution of GGSN-B tells a different story. A significant fraction, around 70% of all samples at GGSN-B suffer from delays greater than 30 ms, while only 10% of the delays are lower than 6 ms (ref. Fig. 3a). Most of the values fall into a histogram bin at around 400 ms (ref. Fig. 3b). This is clearly an anomalous behavior indicating congestion. The overload was caused by a hardware problem within the reassembly of IP packets, resulting in substantial packet delays under high traffic load¹. The problem was readily fixed by reconfiguration. In Fig. 3a we also report the ECCDF of GGSN-B one day after the fix. Clearly, the delay distribution at GGSN-B looks “healthy” again. Notably, both ECCDF in Fig. 3a show a very regular shape (straight in log-log scale).

Having learned that a hardware problem may manifest itself in a change of the delay distribution, one could think of implementing an automatic online agent, continuously measuring the delay and reporting alarms in case of abnormal deviations. Alarming based on a simple threshold on some delay percentile would be sufficient to detect similar congestion events in the GGSN.

V. SGSN DELAYS IN UPLINK

In principle, detecting anomalies in the GGSN delay process is a relatively easy task. This is because the delays internal to the GGSN are caused only by queuing. Hence, the typical delay values are very low in a “healthy” and well-dimensioned GGSN. This is not the case with the SGSN, where packets can be buffered due to *user mobility* and *flow control*. As we will show later, long-delays emerge in the SGSN delay process, in addition to the short-delays due to queuing. Before inspecting the delay distribution, we have to select a suitable observation period. In Fig. 4 we report the time series of the all-day traffic load arriving at the Gn interfaces of SGSN-A and SGSN-B in uplink and downlink. The rate is measured as byte counts in time bins of 1 minute. The total traffic load of SGSN-A is substantially lower (each graph is rescaled by an arbitrary factor²). The time window ΔT from 8 am to 4 pm is our observation period for the delay analysis. That means we are not focusing on

¹Such kind of anomalous events are extremely infrequent in the operational network under study. When they occur, the packet-level traces around the anomaly provide highly valuable (and rare) information about the actual traffic patterns during a real anomaly. Therefore, we are eager to store such data and explore them carefully, in order to develop and deploy practical methods for detecting future occurrences of similar events. As a side value, the analysis of the anomalous dataset often leads to the removal of the very root cause for the anomaly (e.g., an improved equipment configuration), so as future occurrences become even less likely. In this sense, the research plays the role of an additional “learning loop” that contributes to increasing the robustness and performance of the network infrastructure.

²Several quantitative values are considered business-sensitive by the operator and subject to non-disclosure, e.g. absolute traffic volumes, number and capacity of monitoring links. For this reason the following graphs reporting absolute traffic values have been rescaled by an arbitrary undisclosed factor.




Fig. 4: Traffic rates at two SGSNs




Fig. 5: Uplink delay at two SGSNs

the peak hour (8-9 pm in this network) but on a period of medium traffic load. Since our aim is to investigate the “normal” delay process in a well-dimensioned SGSN, choosing an off-peak observation period eliminates the risk of bias due to overload conditions.

To investigate the uplink delay induced by a GPRS SGSN, we report the ECCDF of the uplink packet delay at SGSN-A and SGSN-B in Fig. 5a. The distributions have comparable shapes for both SGSNs. However, we observe that 80% of the packets of SGSN-A, and only 60% for SGSN-B, are forwarded after less than 1 ms. Recall that the traffic load is higher at SGSN-B, but this alone cannot explain the difference between the distributions. In fact, 2% of the samples at SGSN-A, but only 0.9% for SGSN-B experience a delay of more than 10 ms. For both SGSNs the mass of the distribution stays below 5 ms, while a noticeable bump is present at higher delays. To further investigate these bumps we resort to the scatter plot of packet length L vs. packet delay D at SGSN-A in the period 8-10 am (ref. Fig. 5b). We observe clear delay clusters for packet lengths $L > 500$ bytes. This is due to packet fragmentation, as explained below.

As noted in § III-D, at the Gb interface an uplink IP packet of length $L > S$ is segmented into at least two LLC frames. At the SGSN the packet is reassembled and forwarded at once via the Gn interface. The value of S is negotiated between the MS and the SGSN, the default value in the specifications is 500 octets [15] and is typically used by mobile handsets, while data cards are sending up to 1500 payload octets in a single LLC frame. As shown in Fig. 2, at the radio interface Um between the MS and the BSS, one LLC frame is packed into several radio blocks of the RLC protocol [16]. These radio blocks arrive at the BSS every 20 ms [16, p.24-33] indicated by Δt in Fig. 2. Hence, radio blocks carrying the last chunk of an LLC frame have an arrival time T_B of an integer multiple of 20 ms. T_B depends on the LLC frame length, the number of LLC bytes carried per radio block and the actual transmission bandwidth (including possible error correction). According to the definition of $d' = t_3 - t_1$ (ref. Fig. 2), the time T_{TX} between the arrivals of the LLC frame carrying the first and the last chunk of a segmented IP packet is included in the delay measurement. Since LLC frames arrive at a periodicity of T_B , the interval T_{TX} includes the T_B 's of those LLC frames, which are received after the one carrying the first chunk of an IP packet. Hence, plotting $T_{TX} + T_{SGSN}$ vs. packet length L reveals the delay clusters observed in Fig. 5b. Packets shorter than 500 bytes (i.e., default value of S) are never fragmented, hence, we observe no clusters for those packets. The blue dashed lines manually fit into these cluster, have an offset of exactly $\Delta t = 20$ ms to each other. By applying $D = L/B$, we can determine the slope $h = 1/B$, where B denotes the actual bandwidth corresponding to these blue lines. In Fig. 5b, we find $B = 192$ kbit/s. This bandwidth stems from the fact that several BSSs are connected to the SGSN via E1 links. These links provide a total capacity of 2 Mbit/s, which is multiplexed into 32 time slots, resulting in 32 channels with a bandwidth of 64 kbit/s [17]. In our case the E1 connections are obviously configured to aggregate three E1 time slots to one channel, resulting




Fig. 6: Downlink delay distribution at two SGSNs

in the observed bandwidth of 192 kbit/s.

Since the LLC frame arrival T_B depends on the actual radio conditions, we can infer the uplink bandwidth available to MS from the scatter plot in Fig. 5b. We manually fit three different slopes for three different bandwidths 8, 20 and 40 kbit/s - represented by the dashed black, solid black and red line - into the observed clusters. The $B = 8$ kbit/s can be considered as the minimum observed bandwidth. A bandwidth of 20 kbit/s corresponds to the specified data rate of one time slot of GPRS Coding Scheme 4 (CS-4), while a rate of 40 kbit/s is available to mobile handsets supporting either two CS-4 time slots or E-GPRS also known as EDGE [16].

The above discussion shows that the analysis of the IP packet delay d' on the user plane of the core network, can provide insight into the current radio conditions. In principle, also the relative usage of different access technologies like GPRS and E-GPRS might be distinguished. We leave this aspect as an interesting direction for further research.

In this work we aim at analyzing the delay in order to reveal equipment problems at the SGSN. Therefore, we need to separate the delay component that depends on the RAN conditions from the delay component inside the SGSN. This can be achieved by changing the delay definition from $d' = T_{TX} + T_{SGSN}$ to $d'' = T_{SGSN}$ (ref. Fig. 2), such that t_2 is regarded as the arrival time of the complete IP packet at the Gb interface. In Fig. 5c we report the ECCDF of the refined uplink packet delay at SGSN-A and SGSN-B. The bumps in the tails of the distributions reported in Fig. 5a have been cleared out. In Fig. 5c queuing is the only source of packet delay. Note also that the tail of the higher utilized SGSN-B shows constantly larger delay than of SGSN-A. We also verified, that the fragmentation clusters reported in Fig. 5b are eliminated with the new delay definition d'' .

VI. SGSN DELAYS IN DOWNLINK

Our investigations of the downlink delay at the monitored SGSNs, are based on the delay definition d'' (see § V). In Fig. 6 we report the ECCDF and the empirical histogram of the downlink delay at SGSN-A and SGSN-B from 8 am to 4 pm. The mass of the ECCDFs in Fig. 6a is comprised by low delays for both SGSNs. At SGSN-A $\approx 93\%$ and at SGSN-B $\approx 80\%$ of the samples are lower than 3 ms. The larger fraction of low delays at SGSN-A may again result from the lower traffic load compared to SGSN-B. Both ECCDFs show a significant step at around 2 sec. From the empirical histogram in Fig. 6b we see that most delay values are lower than 10 ms, while there is a clear bump for higher delays. These are due to *user mobility* and *flow control*.

As noted in [18], the wireless downlink rate in GPRS networks is time varying. To avoid overflow and/or underflow of the downlink buffers at the BSS, the downlink transmission rate from the SGSN to the BSS is regulated for each data connection by the flow control mechanism of the BSSGP protocol. In other words, downlink packets may be buffered inside the SGSN if the downlink buffer at the BSS is exhausted. Besides BSSGP flow control, a further cause for downlink packets buffering at the SGSN are mobility management procedures like *handover* and *paging*. The handover procedure occurs when an active user is changing to another cell and/or routing area. In this case, the SGSN may temporarily store the corresponding packets until the location change is completed and the packets can be forwarded to the new location. The paging procedure will be described in detail below.

We believe that the delay induced by flow control and mobility is the main cause of the bumps observed in the empirical histogram. It is not trivial to discriminate and quantify the relative contribution of flow control and mobility to the bump of large downlink delays. The combination of passive and active traffic measurements would be needed in order to gain a detailed description of the delay process in this region. Again we leave the research on this aspect for further study.

Fig. 6b shows a significant delay spike at around 2 sec. As reported in [9] the SGSN may directly forward downlink packets to MSs in READY state, since at the SGSN the location of those MSs is known at cell granularity. After some time of traffic inactivity MSs fall back into STANDBY state, where location changes are announced to the SGSN only at routing area level.




Fig. 7: Stability of percentiles at SGSN-A

If a downlink packet destined to a MS in STANDBY state arrives at the SGSN, it has to perform GPRS paging to reveal the MS's concrete location prior to forwarding the packet to the MS. By performing several active measurements reported in [8], we validated that packets forwarded after a successful GPRS paging experience a delay of around 2 sec. This explains the spike observed in Fig. 6b, which is caused by packets being delayed due to GPRS paging. The ECCDF of two shorter periods at SGSN-A in Fig. 6c will be discussed in the next section.

From a previous work [9], we know that scanning traffic is a possible source of paging. Compared to this previous study, the presence of scanning traffic in the new datasets considered in this work is dramatically lower and most of the paging is caused by legitimate applications. By manual analysis we found that two third of the total paging is due to the blackberry service [21], while the remaining third is likely caused by low-interaction services as, e.g., messaging and notification applications.

VII. STABILITY OF DELAY PERCENTILES

After having dissected the delay process at the SGSN, we now look at the stability in time of the delay distribution. In fact, a high level of stability (small random fluctuations) is an important prerequisite for using the delay percentile as input for anomaly detection.

In Fig. 7, we plot six different percentiles vs. time for SGSN-A. Every percentile is computed from an observation period of 5 min. The uplink percentiles in Fig. 7a (linear vertical scale) are very stable over time. The 0.1-0.5 percentiles are permanently below 1 ms, and the 0.99 percentile is only fluctuating between 2.1 and 2.7 ms. However, the downlink percentiles in Fig. 7b (logarithmic vertical scale) show a different behavior. While the 0.1-0.9 percentiles are rather constant and all below 3.5 ms, the 0.95 percentile is heavily fluctuating between 3 ms and 2 sec! That means a small *vertical* shift in the ECCDF (ref. Fig. 6c) in this flat region results in the 0.95 percentile (i.e., *horizontal* distance) fluctuating widely between two orders of magnitude. This is due to the bimodal shape of the delay distribution (ref. Fig. 6b). Since shifts of the flat regions in the ECCDF are caused by regular variations within the global mobility and flow control processes, the observed fluctuation in the 0.95 percentile should be considered physiological of the normal delay condition. The lesson to be learned is, that higher percentiles in downlink, are inadequate for detecting abnormal conditions in the SGSN. An alternative choice might be to use low percentiles and/or vertical distance as, e.g., the relative fraction of samples that exceeds some delay threshold. In Fig. 7c we show this delay crossing share for a number of different threshold values computed in time bins of 5 min. The time series are much more stable compared to the high delay percentiles. This suggests that these indicators should be preferred over the canonical percentiles for the task of anomaly detection on the SGSN delay process.

VIII. CONCLUSIONS AND OUTLOOK

In this work we have explored the single-hop delay distribution at the core elements of a 3G Core Network. At the GGSN, packet buffering is due to queuing, and the resulting delays are very low and generally stable. In a few real cases of congestion and equipment malfunctioning (one case being reported in this work) we have observed that the mass of the delay process jumps to abnormally large values, making the task of anomaly detection relatively easy. Basically, simple thresholding on selected key percentiles would suffice to detect the problem.

The task is more complex for the SGSN, where a subset of packets experience large delays also under normal operation conditions, due to flow control and user mobility. Here, fixed-threshold crossing frequencies (i.e., the relative fraction of packets exceeding a certain delay value) appear to be better indicators than canonical percentiles. In our project [11] we are implementing online alarming agents based on these criteria, to be used experimentally on a production basis. This would provide a direct and reliable means to promptly detect a broad class of network problems.

During the study we noted that *packet fragmentation* introduces a correlation between the reception time of long uplink packets at the SGSN and the actual conditions of the radio link (bandwidth, loss). Similar interdependencies exist in downlink,

due to flow control. These observations indicate the possibility of monitoring *indirectly* the global status of the radio section from the analysis of one-way delays captured in the Core Network, specifically at the SGSN. We believe this is an intriguing direction for further research.

ACKNOWLEDGMENTS

The Telecommunications Research Center Vienna (ftw.) is supported by the Austrian Government and the City of Vienna within the competence center program COMET. This work is part of the DARWIN project [11] run at ftw.

REFERENCES

- [1] K. Papagiannaki et al., "Measurement and Analysis of Single-Hop Delay on an IP Backbone Network," *IEEE JSAC*, vol. 21, no. 6, August 2003.
- [2] B. Choi et al., "Analysis of Point-To-Point Packet Delay in an Operational Network," *IEEE INFOCOM'04, Hong Kong*, March 2004.
- [3] P. Benko, G. Malicsko, A. Veres, "A Large-scale Passive Analysis of End-to-End TCP Performance over GPRS," *IEEE INFOCOM'04, Hong Kong*, March 2004.
- [4] F. Vacirca, F. Ricciato, R. Pilz, "Large-Scale RTT Measurements from an Operational UMTS/GPRS Network," *1st Int'l Conference on Wireless Internet (WICON'05), Budapest*, July 2005.
- [5] J. Kilpi, P. Lassila, "Micro- and macroscopic analysis of RTT variability in GPRS and UMTS network," *Proc. of Networking 2006, LNCS 3976, Coimbra, Portugal*, May 2006.
- [6] Y. Lee, "Measured TCP Performance in CDMA 1x EV-DO Network," *Proc. of 7th Passive and Active Measurement conference (PAM 2006), Adelaide, Australia*, March 2006.
- [7] J. M. Cano-Garcia, E. Gonzalez-Parada, E. Casilari, "Experimental Analysis and Characterization of Packet Delay in UMTS Networks," *Proc. of 6th Int'l Conference on Next Generation Teletraffic and Wired/Wireless Advanced Networking, St. Petersburg*, May 2006.
- [8] A. Barbuzzi, F. Ricciato, G. Boggia, "Discovering parameter setting in 3G networks via active measurements," *in preparation*.
- [9] F. Ricciato, E. Hasenleithner, P. Romirer-Maierhofer, "Traffic analysis at short time-scales: an empirical case study from a 3G cellular network," *IEEE Trans. Network and Service Management*, to appear, 2008.
- [10] J. Bannister, P. Mather, S. Coope, *Convergence Technologies for 3G Networks*. Wiley, 2004.
- [11] METAWIN and DARWIN projects. [Online]. Available: <http://userver.ftw.at/~ricciato/darwin>
- [12] Endace Measurement Systems. [Online]. Available: <http://www.endace.com>
- [13] "The MD5 Message-Digest Algorithm," *RFC 1321*, April 1992.
- [14] S. F. Donnelly, "High Precision Timing in Passive Measurements of Data Networks," Ph.D. Dissertation, CS Dept., Waikato Univ., Hamilton, New Zealand, June 2002.
- [15] "Mobile Station - Serving GPRS Support Node (MS-SGSN); Logical Link Control (LLC) Layer Specification," *3GPP TS 44.064, Version 7.2.0, Release 7*, October 2007.
- [16] "Overall description of the GPRS radio interface; Stage 2," *3GPP TS 03.64, Version 8.12.0, Release 1999*, April 2004.
- [17] "Physical/electrical characteristics of hierarchical digital interfaces," *ITU-T Recommendation G.703*, November 2001.
- [18] A.S. Bedekar, R. Agrawal, R. Ranjan, "A lossless algorithm for BSSGP flow control in GPRS and EDGE," *IEEE GLOBECOM'03, San Francisco*, December 2003.
- [19] M. Allman, V. Paxson, J. Terrell, "A Brief History of Scanning," *ACM Internet Measurement Conference (IMC'07), San Diego*, October 2007.
- [20] R. Pang et al., "Characteristics of Internet Background Radiation," *ACM Internet Measurement Conference (IMC'04), Taormina*, October 2004.
- [21] Blackberry Wireless Solutions. [Online]. Available: <http://www.blackberry.com>