Network-wide measurement of TCP RTT in 2G networks

Bernd Pfitzinger[1], Tommy Baumann[2], Andreas Emde[2], Daragan Macos[3], and Thomas Jestädt[1]

[1]*Toll Collect GmbH, Linkstraße 4, 10785 Berlin, Germany*
[2]*Andato GmbH & Co. KG, Ehrenbergstraße 11,98693 Ilmenau, Germany*
[3]*Beuth Hochschule für Technik, Luxemburger Straße 10, 13353 Berlin, Germany*

Abstract

We analyze existing server-side log data of a large scale automatic toll system to measure the TCP round-trip-time (RTT) as experienced by the communication between the central system and the on-board units (OBUs) deployed for tolling heavy-goods vehicles. The RTT is estimated from passive monitoring by parsing server-side log files and aggregating fleet-wide statistics over time. Using this data we compare the characteristics of the four different types of OBU and the three GPRS (2G) networks used. We find the RTT data to be consistent with existing, smaller samples and extend the observed RTT range by an order of magnitude. The OBU types exhibit a markedly different behavior, most notably for long RTTs, and we find one of the 2G networks to 'hum' at 50 Hz and harmonics.

1. Introduction

Cellular networks are ubiquitous and by now part of many technical systems. Improving the bandwidth and lowering the network latency has been a focus of many network generations. In the context of machine-to-machine systems (M2M), cellular network access can be a substantial expense: The investment in hardware (e.g. provisioning mobile equipment with cellular modems and SIM cards) and the ongoing service fees may be moderate when compared to the price set for consumer smartphones and contracts but at the same time the operating costs of M2M systems often needs to be very low, e.g. for regularly collecting data from many millions of mobile units as it is the case in smart metering and smart grids [1] or electronic tolling applications [2].

Smart metering or tolling applications are typical examples of M2M systems where deferred (i.e. non-real-time) and low-bandwidth communication suffices between the mobile unit and a central system. Consequently, 2G networks are still an enabling technology,

chosen over more modern network types [3].

In this article we draw on experience designing and operating the German automatic toll system for heavy goods vehicles. With more than 1 000 000 on-board units (OBU) deployed it is a typical, moderately large M2M system. Since the toll fee is calculated by the OBU, infrequent and small data transmissions characterize the automatic tolling process. Updates to the geo and tariff data occur regularly but rarely. The system has been fully operational for more than a decade and over the past years we have introduced a detailed simulation model to aid the system (re-)design and operations (for details see [4] and references therein).

Our model of the German automatic toll system includes the technical features concerning the most important processes – transferring toll data from the OBUs to the central system and downloading updates to the geo and tariff data from the central system. The mobile data networks are treated as 'black box' systems, parametrized with different bandwidths and latencies for the three German network operators and the various types of OBUs deployed. No attempt was made to model the underlying TCP/IP protocol or even the GSM network stack.

While the simulation model adequately predicts the system behavior, the validation of the network abstraction remains an open task. It is of course no surprise that the most frequent process of the automatic toll system is the transfer of toll data from the OBU to the central system. In the current system design the toll calculation is a task executed by the OBU and only the results are transmitted, typically packaged according to various parameters (e.g. regarding the total amount of tolls due, the timing or the optimal technical data packet size). As a consequence, the average toll data transmission is very brief with less than 1 kB transmitted taking a total of several seconds.

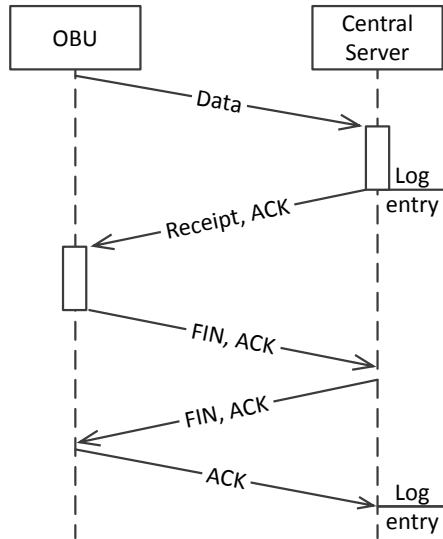The data transmission is a typical TCP/IP connec-

HℐCSS

Figure 1: Sequence diagram of the TCP/IP communication termination between an on-board unit (OBU) and the central system.

tion where the toll data most often fits into a single TCP packet (see figure 1) and is acknowledged (both on the TCP layer and the application layer) by a single TCP packet, followed by the TCP connection termination.

Upon validating a simulation model we noticed a lack of statistical data describing the dynamic behavior of cellular networks: Research on 2G networks apparently ceased years ago (an overview of the GPRS performance is given in [5]) and very few publications give statistically representative data e.g. for network latency and bandwidth. Yet the underlying statistics is very relevant in modeling real-world applications.

To aid our model validation we started gathering statistical data concerning the cellular network behavior, at present limited to measuring the network latency. This article summarizes our findings within the geographical coverage of the German mobile network operators.

In the next section we summarize the existing literature on 2G network latency and introduce our technical setup which is characterized by analyzing existing server-side log data to infer the TCP/IP round trip time. Accumulating data over months we analyze a total of several hundred million data points gathered from the whole OBU fleet. Section 3 discusses the results in particular regarding the heavy-tail of the latency probability density and the differences between various types of OBUs and mobile network operators.

Overlaying the data presented in section 3 is a pronounced 'hum', easily visible in the un-binned data: Section 4 revisits the data in the frequency domain to look at the modulation of the RTT as given by the time difference of two subsequent log entries. Section 5 uses a small sample to look at the intra-day changes of the latency.

2. Literature review and technical setup

The literature on the performance characteristics of GPRS networks beyond a laboratory setting is scarce, most notably [6], [7]. More recent articles cover GPRS/EDGE and 3G networks (e.g. [8], [9] and [10] giving a detailed breakdown of the latency in 3G networks). Typically these experiments attach passive monitoring to the GPRS core network to collect statistical data on the TCP connections either for specific TCP packets or by analyzing all data transmitted. The largest one, [6] gathers statistics in seven (unnamed) countries for a total of some 12 million TCP connections regarding the connection setup, i.e. the TCP SYN and ACK packets – where the initial retransmission timeout limits the maximum observable delay to 3 seconds. They report on observing a successful connection setup in 91.26 % of the cases, taking on average 1364 ms (most often to transmit a single SYN packet followed by a single ACK packet), degrading slightly with increasing network utilization and when the mobile unit is traveling [11]. The latency depends strongly on the type of mobile unit [8] and increases with the packet size [12].

Our approach to measure the RTT differs from the articles mentioned above: With emphasis on the end-to-end performance we do not access the network (either on the side of the mobile network operator or the central system of the automatic toll system) rather we analyze existing server-side application logs. To facilitate operational monitoring, the receiving server creates several log entries for each connection, e. g. comprising a timestamp with one milli-second time resolution and sufficient information to allow reconstructing individual connections as reported by the server application. Parsing the log data we observe more than 100 distinct patterns for "connections" – almost all corresponding to very rare error modes.

As expected, the dominant process of receiving toll data is the most frequent type of connection: 99.68% of all connections observed. Logging starts when the TCP connection is established which is immediately followed by the transfer of toll data. This time period is not uniquely dominated by network latency – the TCP connection has already been set up and the data transfer depends on the bandwidth and latency. Therefore to isolate the network latency as closely as possible we focus on the time difference between two successive log entries: The first entry acknowledges the successful toll data reception directly followed by a log entry noting the
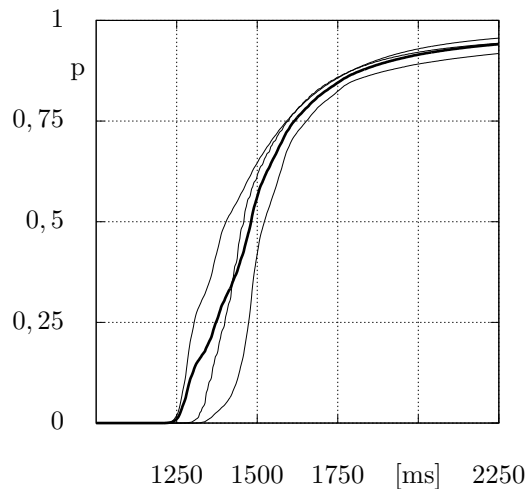
Figure 2: Empirical cumulative distribution function of the connection termination latency per cellular network (thin lines, normalized) and overall (bold line).

regular TCP connection termination (see figure 1, lower part). Note that lacking access to network traces, we are unable to verify the typical number of TCP/IP packets exchanged in this step.

The data collection also notes the mobile network operator and the type of OBU for each event. Four different OBU hardware platforms share the same tolling application but differ in their hardware and operating system. All OBUs are equipped with GSM modems (the most recent OBU type supporting up to GPRS multi-slot class 10) and a SIM card of one of the three national network operators. National roaming is not used and the OBUs are configured to discourage data transmissions abroad.

The receiving server uses ntp for time synchronization. To assess the potential timing uncertainty we compared the session duration as given by the time difference of the first and last log entry and the session duration computed and logged by the application. In most cases (99.9%) both values agree within $\pm$ 2ms.

3. Connection termination latency

To facilitate this analysis we collected a total of more than 300 million data points that can be separated by OBU type and network operator (but not e.g. by date). A quick and typical first overview for the data set is given in figure 2, where the empirical cumulative distribution function (CDF) of the connection termination latency is displayed (bold line in figure 2). The CDF remains close to zero up to approximately 1200 ms and quickly rises thereafter to reach the median at 1545 ms. As the latency increases the CDF flattens out leading to an average of

1806 ms.

Looking at the literature, [6] presents a large statistical data set dating back to 2004 for the connection setup RTT in several GPRS networks (see figure 4 in [6]) and observes the RTT involving two TCP packets to rapidly rise after 600 ms with a median close to 900 ms. Comparing the average RTT (of 1822ms in [6]) is misleading since the connection setup ends by definition after the initial retransmission timeout of three seconds whereas in our case timeouts occur much later.

With this information we deduce that in our case the 'connection termination' involves the exchange of two pairs of TCP packets (as shown in figure 1), i.e. including some processing time on the OBU.

Figure 2 also shows the CDF separately for each of the three national cellular networks. To improve readability of the figures we choose to normalize the probability distributions of each data subset independently to "1". While the shape and onset of the CDFs is similar, the latency differs perceptibly, e.g. at the median between 1500 ms and 1615 ms. However, in comparison the effect of different OBU types on the latency is even larger (ranging from 1481 ms to 1792 ms where the most recent OBU type gives the lowest latency).

After three seconds more than 95% of all connections are successfully terminated. The remainder extends over a considerable timespan which we can measure since the setup is not limited e.g. by short, initial retransmission timeouts and the data set is large enough to provide reasonable counts even for low probabilities (especially when we change to larger binning).

3.1 Behavior of the most recent OBU type

To give an overview we change the data display from the CDF to a log-log-plot of the empirical probability density function (PDF), emphasizing the extent of low-probability events. Figure 3 displays the histogram for the most recent OBU type for latencies up to 100 seconds (bold line). The histogram is converted to the probability density (per 1 ms bin) and short-term network effects – most notably the 20 ms granularity of the GPRS upstream traffic [13] – are smoothed out by binning 20 or more consecutive data points. In addition rather than drawing the data points we chose to display the connecting lines.

As seen in the CDF, the PDF exhibits a sharp rise followed by a brief plateau that is effectively an artifact of superimposing the three national networks (thin lines in figure 3 , each subset is normalized to 1) with one 'peak' each at neighboring positions on the time axis. Very few data points exist for the time period before the sharp rise, their cause or origin remaining unknown. Considering
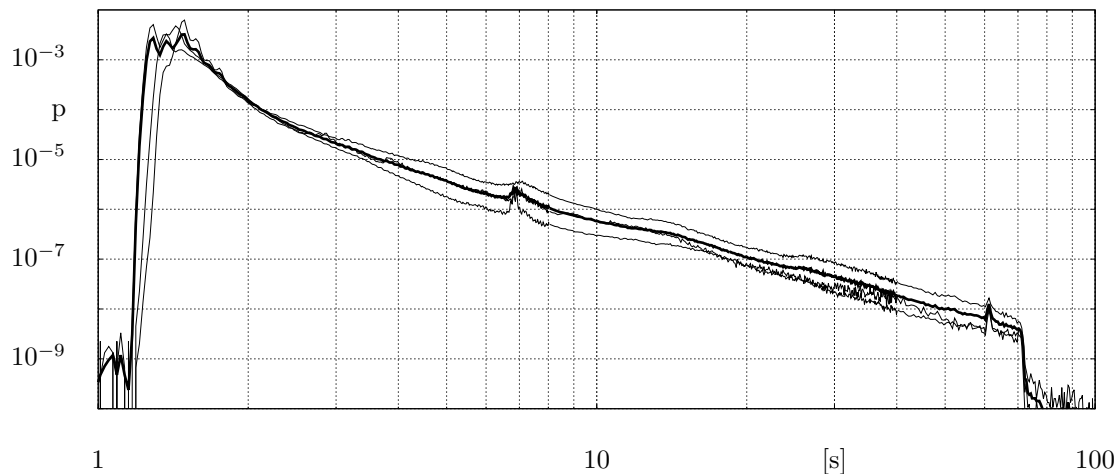
Figure 3: Empirical probability density function of the TCP session termination latency observed for the most recent OBU generation across all three national cellular 2G networks (bold line) and per network (thin lines, normalized). The data is aggregated using 20 ms, 100 ms and 500 ms binning for the regions below 8s, up to and beyond 40s. The distribution has a pronounced heavy tail extending up to and beyond the connection timeout of 70s.

that for this type of OBU 97% of all data points have latencies of less than three seconds, there remain sufficient events to suggest a power-law heavy-tail extending up to 70 seconds.

To extend figure 3 to very low probability densities we sacrifice the time resolution (x-axis) with a larger bin width of 100 ms (between 8 s and 40 s) and 500 ms beyond. The benefit is a reduced uncertainty of the empirical probability since more events fall into the same bin. As an example, the aggregated data (bold line in figure 3) corresponds to approx. 4500 events in the last 20 ms bin (at 8 s), switching to 100 ms binning enlarges the signal to some 22000 events allowing to extend the diagram further. Selecting arbitrarily the next transition to a lower time resolution at 40s we lift the signal strength from almost 400 events per 100ms bin to approx. 2000 events in the 500 ms bin. Even with this large bin width the low probability – 6 orders of magnitude below the peak – becomes visible as considerable 'noise'. As the signal strength attenuates further it falls below 100 events per bin only when the latency exceeds the preset application-layer connection timeout of 70 s.

This connection timeout lowers the signal by at least one order of magnitude and the total number of events beyond 70 s is too small to draw further conclusions. Approximately 1 in 300 connections ends with a timeout of the connection termination and consequently an average connection duration exceeding 100 s.

About 10 s prior to the connection timeout a pronounced peak is visible, again independent from the cellular network or OBU type. The peak in figure 3 is at 61470 ms corresponding probably to a process with a delay time of 60 s. Without further access to the trans-

mitted packets we can only speculate as to its origin. One possibility is the TCP persist timer (bounded to be at maximum 60 s) triggering a transmission when an TCP acknowledgement packet is lost to avoid the otherwise inevitable deadlock situation [14].

Returning to the peak probability, the data set for each cellular network exhibits some 'waves' before the heavy-tail of the PDF begins. Again, treating the data transmission as 'black box' we cannot investigate the cause. However, the literature also observes this behavior (see figure 18 in [7]) and argues that it originates from the buffering within the network, e.g. no buffering, buffers per SGSN or per mobile unit. Close to a latency of 7 s a pronounced peak is visible in all data sets differentiated per cellular network (thin lines in figure 3) and even when further differentiated by OBU type (not shown). The probable cause is the loss of one TCP packet and the successful retransmission after the retransmission timeout – depending on the average RTT and its standard deviation as experienced during each connection [14].

3.2 Observing the TCP retransmission timeout

So far the PDF of the connection termination latency followed a simple shape: The majority of the observed data points lies between 1200 ms and 3000 ms followed by a heavy-tail extending to long durations. In the log-log graph the heavy-tail is almost a straight line, i.e. a power-law fall off. This simple picture changes when the data coverage is extended to include all OBU types.

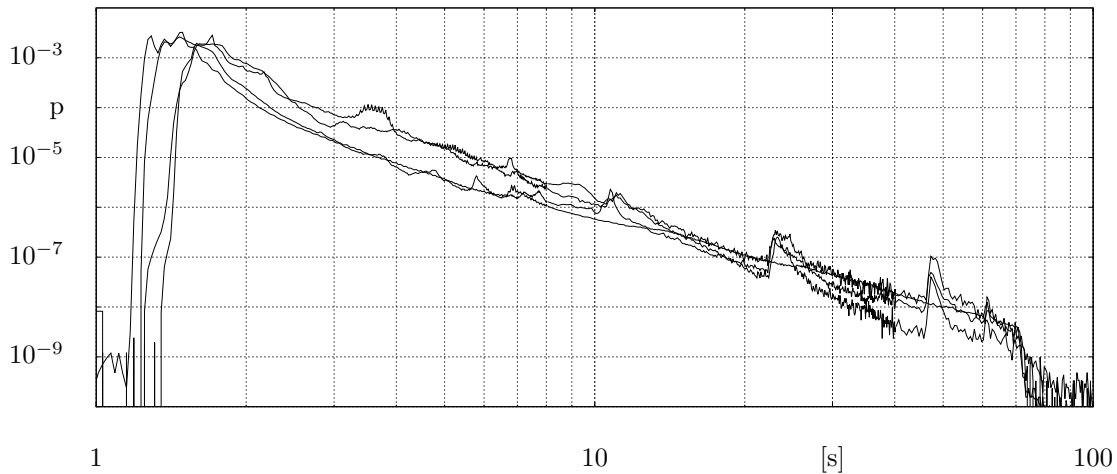The first look at figure 4 makes clear that the differ-

Figure 4: Normalized empirical probability density functions of the TCP session termination latency per OBU type aggregated across all cellular networks. Data binning as in figure 3 .

ences between the OBU types (i.e. the hardware and operating system) affect the latencies much more than the cellular network –regarding not only the onset of the rise in the PDF, but also the position of peak probability and the shape of the heavy tail. In addition the somewhat smaller sample size results in an increased sampling error of the probability measured for a given bin, i.e. the histogram becomes 'noisy'.

Several features of the PDF remain valid: The connection timeout after 70s concerns all connections, the peak at 61s is also visible for all OBU types as well as the (only) peak seen close to 7 s for the most recent OBU type.

At the far end of the heavy tail the most prominent change is the occurrence of pronounced broad peaks approx. at 47 s, 23 s and 11s, i.e. 24 s and 12 s apart – almost equidistant on the logarithmic x-axis. It is not obvious, whether this sequence extends to shorter time periods. We hypothesize that the behavior is caused by the exponential back-off algorithm used in the calculation of the retransmission timeout: To avoid causing congestion the next retransmission is delayed twice as long as previously. In principle the sequence of retransmission timeouts should continue with delays increasing up to 64s [14]. If this were the case the next 'peak' would occur 48 s later considerably after the 70s timeout. Even with the large data set available and increasing the binning to 2000 ms the number of events is so small that no statistically significant conclusion can be drawn.

We note that the PDFs exhibit many features for latencies up to 7 s, strongly depending on the OBU type (and sometimes the cellular network, not shown). Their cause is unknown and we are not aware of similar findings in the literature.

4. Frequency domain

2G networks use TDD (time division duplex) to separate radio transmissions into distinct time slots, with inbound and outbound transmissions on the same frequency but in different time slots. For GPRS this naturally leads to a granularity of 20 ms (a combination of the actual duration of 18.46 ms of a TDMA frame and a markup from the signaling frame every 12 frames, see [13]). So far we have intentionally disregarded the short term dynamics introduced by the cellular network by using at least 20ms binning of the data.

What is the short-term behavior of the latency? The data extraction is limited by the clock granularity of 1 ms, well below the 20 ms time slot duration. Figure 5 revisits the data gathered for the most recent OBU type (as shown in figure 3) in the vicinity of the peak of the probability density albeit with a bin width of 1 ms and linear scaling of both axis. To improve readability we have increased the binning to 2 ms for two of the data sets shown. At this scale it is possible to plot the actual data points and the lines are added to aid the eye.

The 20 ms rhythm induced by the cellular network time slot granularity is obvious in figure 5 for all three networks. Surprisingly, harmonics at 100 Hz and 150 Hz are present in one of the three networks deserving a closer look at the data set. To that end we take the event histogram with a bin width of 1 ms and switch to the frequency domain by applying a discrete fast fourier transform to the complete, unmodified data set. Since the data is always (considerably) above zero, the low frequencies are of no concern in this discussion – they are only needed to represent the non-zero average and the slope of the heavy-tail.
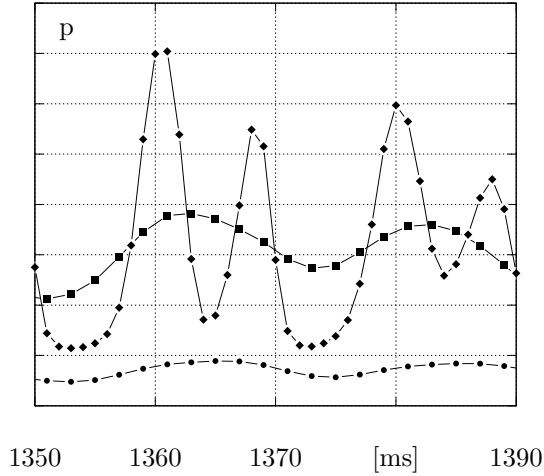
Figure 5: Normalized empirical PDF for the most recent OBU generation per cellular network showing the 50 Hz granularity of the 2G radio network. One network gives an additional 100 Hz signal (and further harmonics) visible in the 1 ms binning. The harmonics are not present in the other networks (the time resoultion in the graph is reduced to 2 ms binning).
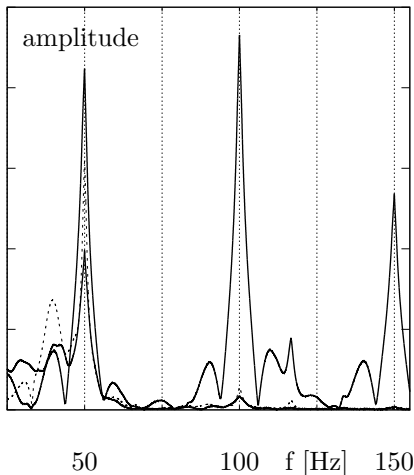

Figure 6: Fourier transform of the empirical PDF from figure 5: One out of three 2G networks shows strong harmonics at 100 Hz and 150 Hz.

The resulting spectrum is shown in figure 6: Clearly, all three cellular networks exhibit the 50 Hz 'hum' (i.e. 20 ms time slot granularity). One of the cellular networks has additionally an equally strong signal centered at 100 Hz and similarly at 150 Hz and negligible beyond (not shown). Considering that the same OBUs operate in all cellular networks the origin of these signals is probably neither the OBU (mobile unit) nor the central server. These higher frequency signals could be harmonics of the 50 Hz signal – yet they are not present in two out of three networks.

5.   Variability of the latency

In order to analyze the heavy-tail of the PDF we have gathered and analyzed statistical data over a long period of time and differentiated only according to the OBU type and the cellular network. As cited above the literature suggests that the network latency varies over time: Increasing network utilization during the day leads to increasing latencies.

To verify these observations we extended the data analysis by gathering statistics grouped by the time of day within a week, i. e. modulo 24*7 hours. Due to the late addition in gathering the data the number of events is considerably smaller (encompassing only slightly above 10 million events in total).

Using this small subset of the data we compute the median and average connection termination latency for each hour of the week. The typical behavior is visible in figure 7: During each day the latency increases during the day time impacting the average considerably more than the median, i. e. the network utilization affects the heavy-tail more than error-free transmission. Somewhat surprisingly the week-end does not noticeably improve the peak-hour latency.

The strong peak in the median and average latency around midnight is probably due to server-side maintenance jobs slightly degrading the responsiveness of the server during the off-peak hours.

6.   Summary and outlook

Trying to validate an existing simulation model we noted the lack of statistically representative data on the network latency. With a data set gathered from server-side passive monitoring of the connection termination, containing almost 30 times more events than previously reported in the literature we were able to extend the empirical PDF accordingly up to and beyond the communication timeout of 70 s. The data shows not only that the PDF is sensitive to the network but even more so to
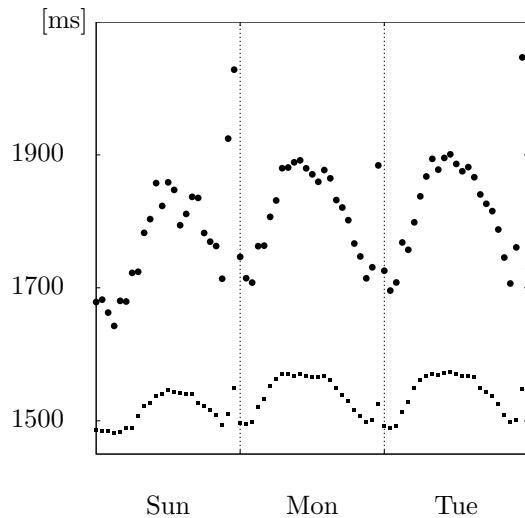
Figure 7: Variability of the median (lower part) and average latency (upper part).

the OBU type where some exhibit a behavior similar to the exponential back-off algorithm. In the frequency domain we confirm the well-known 20 ms time slot granularity of the GPRS networks. Surprisingly, the observed latency oscillates in one of the networks at additional frequencies of 100 Hz and 150 Hz with unknown cause.

To complete the validation of our 'black box' network model additional work is needed regarding the data transmission bandwidth. In practice, applying the results presented here would be simplified by modeling the empirical PDF (rather than using tabulated data).

References

[1] V. C. Gungor, D. Sahin, T. Kocak, S. Ergut, C. Buccella, C. Cecati, and G. P. Hancke, "Smart grid technologies: Communication technologies and standards," *IEEE Transactions on Industrial Informatics*, vol. 7, no. 4, pp. 529–539, Nov. 2011, ISSN: 1551-3203. DOI: 10.1109/TII.2011.2166794.

[2] A. T. W. Pickford and P. T. Blythe, *Road user charging and electronic toll collection*. London: Artech House, 2006, ISBN: 978-1-58053-858-9.

[3] D. Niyato, L. Xiao, and P. Wang, "Machine-to-machine communications for home energy management system in smart grid," *IEEE Communications Magazine*, vol. 49, no. 4, pp. 53–59, Apr. 2011, ISSN: 0163-6804. DOI: 10.1109/MCOM.2011.5741146.

[4] B. Pfitzinger, T. Baumann, D. Macos, and T. Jestädt, "On the necessity for high-availability data center backends in a distributed wireless system," in *2017 50th Hawaii International Conference on System Sciences (HICSS)*, accepted, 2017.

[5] T. Halonen, J. Romero, and J. Melero, *GSM, GPRS and EDGE performance*. Chichester, West Sussex, England: John Wiley & Sons Ltd, 2004, ISBN: 978-0-470-86696-2. DOI: 10.1002/0470866969.

[6] P. Benko, G. Malicsko, and A. Veres, "A large-scale, passive analysis of end-to-end TCP performance over GPRS," in *IEEE INFOCOM 2004*, vol. 3, 2004, pp. 1882–1892. DOI: 10.1109/INFCOM.2004.1354598.

[7] J. Kilpi and P. Lassila, "Statistical analysis of RTT variability in GPRS and UMTS networks," Tech. Rep., 2005. [Online]. Available: http://www.netlab.hut.fi/tutkimus/pannet/publ/rtt-report.pdf (visited on 05/08/2017).

[8] P. Romirer-Maierhofer, F. Ricciato, A. D'Alconzo, R. Franzan, and W. Karner, "Network-wide measurements of TCP RTT in 3G," in *Traffic Monitoring and Analysis: First International Workshop, TMA 2009, Aachen, Germany, May 11, 2009. Proceedings*, M. Papadopouli, P. Owezarski, and A. Pras, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 17–25, ISBN: 978-3-642-01645-5. DOI: 10.1007/978-3-642-01645-5_3.

[9] F. Ricciato, "Traffic monitoring and analysis for the optimization of a 3G network," *IEEE Wireless Communications*, vol. 13, no. 6, pp. 42–49, Dec. 2006. DOI: 10.2514/1.11950.

[10] M. Laner, P. Svoboda, and M. Rupp, "Latency analysis of 3G network components," in *European Wireless 2012; 18th European Wireless Conference 2012*, Apr. 2012, pp. 1–8.

[11] M. Fornasa, N. Zingirian, and M. Maresca, "Extensive GPRS latency characterization in uplink packet transmission from moving vehicles," in *IEEE Vehicular Technology Conference*, Mar. 2008, pp. 2562–2566. DOI: 10.1109/VETECS.2008.563.

[12] L. Fabini, L. Wallentin, and P. Reichl, *The importance of being really random: Methodological aspects of IP-layer 2G and 3G network delay assessment*, Proceedings of the IEEE International Conference on Communications, Aug. 2009. DOI: 10.1109/ICC.2009.5199514.

[13] G. Sanders, L. Thorens, M. Reisky, O. Rulik, and S. Deylitz, *GPRS Networks*. Chichester, West Sussex, England: John Wiley & Sons Ltd, 2004, ISBN: 978-0-470-86955-0. DOI: 10.1002/0470869550.

[14] W. R. Stevens, *TCP/IP Illustrated, Volume 1*. Boston, MA, USA: Addison-Wesley, 1995, ISBN: 978-0201633467.