

# Micro- and macroscopic analysis of RTT variability in GPRS and UMTS networks

\*Jorma Kilpi<sup>1</sup> and Pasi Lassila<sup>2</sup>

<sup>1</sup> VTT Information Technology, P.O.Box 12022, FIN 02044 VTT, Finland.

Jorma.Kilpi@vtt.fi

<sup>2</sup> Helsinki University of Technology, P.O.Box 3000, FIN 02015 HUT, Finland.

Pasi.Lassila@hut.fi.

**Abstract.** We study the data from a passive TCP/IP traffic measurement from a Finnish operator's GPRS/UMTS network. Of specific interest is the variability of Round Trip Times (RTTs) of TCP flows. The RTTs are analysed at micro- and macroscopic level. The microscopic level involves detailed analysis of the RTTs of individual flows, and we are able to detect, e.g., periodic behavior (via Lomb periodogram) and rate changes in the radio channel. At the macroscopic level we focus on the impact of so called self-congestion caused by bandwidth sharing at the mobile device itself, and it is shown how this seriously affects the RTTs observed by a given flow, both in GPRS and in UMTS.

**Keywords:** traffic measurements, RTT variability, GPRS, UMTS


## 1 Introduction

We study the data from a passive traffic capture measurement representing a 30 hour TCP/IP trace measured from one GGSN node of a GPRS/UMTS network in a major Finnish operator's network. The objective of the measurement is to analyze the variability of the Round Trip Times (RTTs) of TCP flows, where one end point of the flow is in the GPRS/UMTS network and the other in the public Internet. Large variability of the RTTs may cause problems to TCP's retransmission methods causing spurious timeouts (see, e.g., [1]). We perform a detailed analysis of the RTTs of selected individual flows (microscopic analysis) and the aggregate traffic (macroscopic analysis). Furthermore, as the measurement data contains TCP flows both from GPRS users and UMTS users, we are able to compare the properties of the RTT process for both technologies.

Some measurements of RTT variation have been recently made in mobile networks see, e.g., [1], [2], and [3]. Notably, in [1] an algorithm is provided for detecting spurious events in TCP, and [3] analyses RTTs, loss and throughput characteristics, among other things. However, our focus is different than in these studies. We study how individual flows observe the RTTs and the aggregate RTTs of all flows. We also experiment with other potentially useful statistical tools not used in earlier studies, namely Lomb periodograms (LP) as motivated by [4] and wavelets. However, the results on wavelets are not in this paper due to space limitations, but can be found in [5]. Finally, we discuss in

---

\* The authors thank Vesa Antervo from Elisa for his efforts in obtaining the trace and Marco Mellia for his help on `tstat`.



**Fig. 1.** Impact of flat rate population on traffic profile (left) and division into subsamples according to TCP connection arrival time (right).

the macroscopic level analysis the clear impact of simultaneous TCP connections (self congestion) on RTTs, which has not been addressed in previous studies.

## 2 Measurement Setup and Classification Methodology


The measurement point was the monitoring port of the Gi interface of one of the GGSN elements. The measurement was planned such that a statistically representative sample of the traffic was obtained. We then verified the accuracy of the time stamps in the data - the accuracy proved to be  $\pm 50\mu s$  (i.e., more than sufficient). The flow-level statistics were obtained by using `tstat`<sup>3</sup>. Additionally, `tstat` was modified to record all valid RTT samples of the TCP flows.

**Traffic profile:** The trace was obtained from a 30 hour measurement on Nov 10, 2004 and consists of traffic from GPRS and piloting UMTS users. Some of the subscribers have only volume based charging but some portion of the subscribers had also flat rate between 18.00-06.00. The effect of this *flat rate population* is significant, see Figure 1 (left). The data has been divided into 5 groups, as shown in Figure 1 (right), according to the TCP connection *arrival time*, which is determined by the time stamp of the SYN packet sent by the client. To study the impact of tariff change, we compare the data from Sample 1 and 2.

**Semi-RTT and RTT Count:** We focus on the properties of RTTs as experienced by TCP flows. As we are not measuring directly at the sender/receiver, the RTT process can not be fully observed. Instead, the notion of semi-RTT is used, similarly as in [2]. Semi-RTT refers to the difference between the time stamps of a (downstream) TCP/IP packet carrying data payload and of the corresponding ACK packet. The size of a flow is measured in terms of the field *RTT Count*, and it is one of the parameters `tstat` provides for each completely observed TCP flow. RTT Count represents the number of times a data segment and the corresponding ACK has been observed, and it is provided both for the upstream and downstream connections.

**GPRS vs. UMTS flows:** During the measurement there were only piloting UMTS users, but the absolute number of observed TCP connections that could be associated to UMTS was sufficient to make comparative analysis against GPRS. Due to the measurement setup, exact identification of whether a flow originated from GPRS or UMTS

<sup>3</sup> <http://tstat.tlc.polito.it>



**Fig. 2.** Time series of Flow 1 (left) and its LPs (middle and right).


was not possible. However, RTTs in UMTS are about one magnitude smaller than in GPRS [2]. Thus, flows with a minimum RTT less than a given threshold were identified as UMTS flows. The threshold value we used was 0.4 ms (in [2] a minimum of 0.476 ms was observed for GPRS).

### 3 Microscopic analysis of long TCP flows


Some specific flows are analyzed first to give the reader a flavor of what is behind the macroscopic level analysis. Example flows were chosen because the time series  $\{(t_i, RTT(t_i)) \mid i = 1, \dots, \text{RTT Count}\}$  had some distinctive features, where  $t_i$  is the time stamp of an ACK packet and  $RTT(t_i)$  is the RTT value calculated from the ACK packet. Additionally, we focused only on the very longest flows to be able to detect clear changes over a comparably long time interval.

**Flow 1 (GPRS):** The first flow was chosen since we wanted to understand what caused the improvement in RTT values after about half an hour as shown in Figure 2 (left), and we wanted to see if the rise in traffic due to tariff change at 18:00 affected this flow. Improvement in RTTs after about half an hour was easily seen to be due to change from large (1380B) to a smaller (536B) segment size. Analysis of the receiving rate of packets revealed that the source is receiving data at a constant rate of 18.2 kbit/s. More detailed inspection showed that a data burst of fixed size was sent every 3rd second on the average. It also takes about 3 seconds before every segment of one burst has been acknowledged by the mobile host. After the change in the segment size, the data burst could be sent slightly better within the 3 second interval. Thus, Flow 1 is probably produced by a streaming audio application, though it was behind TCP port 80. The application had a 3 second play-out buffer, and the application forced the change in the segment size. Finally, another thing to notice was that there were no simultaneous TCP connections from the same mobile.

To further analyse the periodic behavior, the LP has been computed corresponding to the beginning, middle and end of the flow, see Figure 2 (middle and right). From the middle plot it is clear that in the beginning the burst period is slightly above 3 seconds. After the change in the segment size, the period is somewhat below 3 seconds. Examining Figure 2 (left) one can observe that after the tariff change at 18:00 the level of the RTTs rises slightly again (possibly due to increased traffic in the network). Computing the LP from the end shows that the burst period increases to slightly above 3 seconds again, see Figure 2 (right).



**Fig. 3.** Time series of Flow 2 (left) and its LP (right).



**Fig. 4.** Time series of Flow 3 (left) and Flows 4 and 5 (right).

**Flow 2 (GPRS):** As seen in Figure 3 (left) RTTs at a level of 5.0-7.5 seconds seems to be normal for this connection. Unlike for Flow 1, the use of LP revealed no significant periodic structure for Flow 2, see Figure 3 (middle). There it can be seen that Flow 2 has a small spike in the LP at slightly below 1 second, but it is negligible compared to the pronounced spike of Flow 1 at 3 seconds. Moreover, as Figure 3 (right) indicates, there was a change in the downlink capacity from 40.2 kb/s to 26.8 kb/s, which corresponds to a loss of one downlink PDCH when using CS-2 channel coding scheme. Analysis revealed that this was due to self-congestion caused by other simultaneous TCP connections from the same mobile host. The high overall level of RTTs is probably due to a low terminal capacity.

**Flow 3 (GPRS):** Again, use of LP showed no regular periodicity. However, Flow 3 had some peculiar regular intervals of bad RTTs. In this case the self-congestion is also an explanation since the user was simultaneously running an application (MSNP) which initiated a new flow in port 80 regularly with approximately 8 minute intervals downloading a file of size 265.3 kB. These downloads occur exactly at the bad intervals visible in Figure 4 (left).

**Flows 4 and 5 (UMTS):** Both flows originated from the same mobile, and they are shown in Figure 4 (right). The level of RTTs of Flow 4 increase when Flow 5 starts. There were also several other long flows simultaneously from the same mobile. Almost all RTT variations are explained by these other flows.

## 4 Properties of RTT at the macroscopic level

**Changes in quantiles:** We computed the empirical CDFs from the aggregate RTT data. Figure 5 (left) compares Samples 1 and 2 and we can see a clear shift in the quantiles.




Fig. 5. Comparison of CCDFs (left) and increased web activity after 18:00 (right).





Fig. 6. Max RTTs with no upstream traffic (left) and with upstream traffic (right).

Based on packet level data analysis (results not shown here due to lack of space), we argue that the network is not congested but the shifts in the quantiles are due to self-congestion caused by an increased amount of simultaneous TCP connections from a given mobile (bandwidth sharing). Indeed, in Figure 5 (right) we show as a measure of web activity the maximum number of simultaneous TCP flows from a given mobile (IP address). Web activity clearly rises after 18:00 (corresponding to the data in Sample 2).

**Self-congestion:** To further study the effect of self-congestion on the RTTs, we divided such GPRS sessions, that had at least two simultaneous TCP connections, into two groups. The Group 0 did not send much data (Unique Bytes) into upstream direction whereas Group 1 contained those that had some significant data transfers in the upstream direction. More precisely, Group 1 was defined as the set of mobiles for which the total amount of upstream Unique Bytes was larger than 1 kB times the total number of TCP connections from that mobile. Figure 6 shows the maximum RTT of all TCP connections of a mobile against the maximum number of simultaneous connections observed from the same mobile for Group 0 (left figure) and 1 (right figure). The scales of axes are chosen to be the same for both plots in order to show that extremely large RTT values occur almost solely for Group 1 (significant upstream traffic), whereas very large number of simultaneous TCP connections occur for Group 0 (only downloading traffic). Because very large RTT values occur in Group 1, the context switching between transmitting and receiving packets sometimes causes problems.

It can be expected that observed maximum RTTs increase as the amount of simultaneous TCP connections increases. This is verified in Figure 7 (left), which shows an



**Fig. 7.** Robust estimates of conditional expectations of maximum RTTs for GPRS (left) and UMTS (right). Right plot also shows individual samples of maximum RTTs.

estimate of the expectation of maximum RTT over all simultaneous flows, conditioned on the maximum number of simultaneous connections.

Our UMTS sample is not large enough to make the same division into groups as with GPRS. Figure 7 (right) shows the individual values of maximum RTTs as a function of the number of simultaneous TCP flows during a session (red dots), and the conditional expectation of maximum RTT over all simultaneous flows, conditioned on the maximum number of simultaneous connections (solid line). Self-congestion is also a problem for UMTS mobiles, although less serious.

## 5 Conclusions

Results on the variability of Round Trip Times (RTTs) of TCP flows in a Finnish operator's GPRS/UMTS network were given. Microscopic analysis of the RTTs of individual flows was performed on some selected flows, and we were able to detect, e.g., periodic behavior and rate changes in the radio channel. At the macroscopic level it was shown how bandwidth sharing at the mobile device itself seriously affects the RTTs of flows, both in GPRS and in UMTS.

## References

1. Vacirca, F., Ziegler, T., Hasenleithner, E.: Large Scale Estimation of TCP Spurious Timeout Events in Operational GPRS Networks. In: COST 279. (2005)
2. Vacirca, F., Ricciato, F., Pilz, R.: Large-Scale RTT Measurements from an Operational UMTS/GPRS Network. In: First International Conference on Wireless Internet (IEEE WICON 05). (2005)
3. Benko, P., Malicsko, G., Veres, A.: A Large-scale, Passive Analysis of End-to-End TCP performance over GPRS. In: INFOCOM 2004. (2004)
4. Partridge, C., Cousins, D., Jackson, A., Krishnan, R., Saxena, T., Strayer, W.: Using Signal Processing to Analyze Wireless Data Traffic. In: ACM Workshop on Wireless Security, Atlanta, Georgia, USA, ACM (2002)
5. Kilpi, J., Lassila, P.: Statistical analysis of RTT variability in GPRS and UMTS networks. Technical report, VTT and TKK, <http://www.netlab.hut.fi/tutkimus/pannet/publ/rtt-report.pdf> (2005)