Network-Wide Measurements of TCP RTT in 3G

Peter Romirer-Maierhofer[1], Fabio Ricciato[1,3], Alessandro D'Alconzo[1], Robert Franzan[1], and Wolfgang Karner[2]

[1] Forschungszentrum Telekommunikation Wien, Austria
[2] mobilkom austria AG
[3] Università del Salento, Italy
lastname@ftw.at

Abstract. In this study we present network-wide measurements of Round-Trip-Time (RTT) from an operational 3G network, separately for GPRS/EDGE and UMTS/HSxPA sections. The RTTs values are estimated from passive monitoring based on the timestamps of TCP handshaking packets. Compared to a previous study in 2004, the measured RTT values have decreased considerably. We show that the network-wide RTT percentiles in UMTS/HSxPA are very stable in time and largely independent from the network load. Additionally, we present separate RTT statistics for handsets and laptops, finding that they are very similar in UMTS/HSxPA. During the study we identified a problem with the RTT measurement methodology — mostly affecting GPRS/EDGE data — due to early retransmission of SYNACK packets by some popular servers.

1 Motivations

Third-generation (3G) cellular networks provide wireless Internet access to a growing population of mobile and nomadic users. Since the early deployment of GPRS and UMTS at the beginning of this decade, operational 3G networks have been continuously evolving. The introduction of EDGE and HSDPA/HSUPA (or HSxPA) respectively in GPRS and UMTS has increased the available radio bandwidth, while further upgrades are promised by the next wave of radio technologies like HSPA+ and LTE — refer to [7] for more details on 3G technology evolution. The combination of higher bandwidth and cheaper tariffs has produced a substantial growth of 3G user population and traffic volumes (see e.g. [9], which in turn led to major upgrades also in the Core Network. The functional complexity and ever-evolving nature of the 3G infrastructure increase its exposure to problems and errors. Therefore, it is compelling for 3G operators to be able to readily detect network problems and anomalous incidents. To this purpose the operators deploy a number of monitoring and alerting systems, each covering a different section of the global infrastructure and relying on different types of input data and sensors — both passive and active.

M. Papadopouli, P. Owezarski, and A. Pras (Eds.): TMA 2009, LNCS 5537, pp. 17–25, 2009.

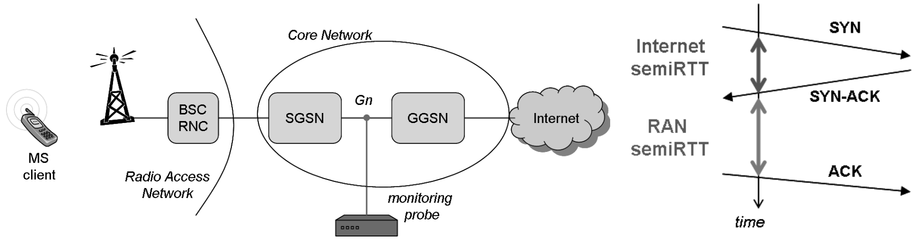

Fig. 1. Monitoring setting (left) and RTT computation scheme (right)

A class of anomaly sensors can be built upon the real-time analysis of packet-level traffic monitors. The basic idea is simple: extract a set of network performance indicators — we call them "network signals", e.g. delays percentiles or loss rate — out of the packet-level trace stream, and seek for deviations from the "normal" profile observed in the past. Such approach underlies two fundamental assumptions: (i) that network performances and the associated "signals" are stable under problem-free operation, and (ii) that network problems induce a recognizable deviation in at least a subset of the monitored network signals.

In this work we consider the possibility of using Round-Trip-Times (RTT) measurements, as obtained from passive analysis of TCP handshaking packets, as a possible "network signal" for detection of network anomalies. We present large-scale measurements from an operational 3G network and investigate the stability of the underlying distributions. Our results are based on very recent traces (January 2009) and include HSDPA/HSUPA and EDGE traffic.

The methodology of inferring RTT from passive TCP traces is not new. Benko *et al.* [1] reported large-scale measurements of TCP RTT from an operational GPRS network already in 2004. We adopt here the same methodology of [1] which considers exclusively SYN/ACK pairs, but provide results also for GPRS/EDGE and UMTS/HSxPA sections. Vacirca *et al.* [2] reported RTT measurements from an operational UMTS network, with data from 2004, considering also DATA/ACK packet pairs. Since then, the capacity of 3G network has increased considerably, due to the introduction of HSxPA and EDGE, and consequently the measured RTT values are now considerably lower. While some recent papers have investigated the delay process in HSDPA via active measurements (e.g. [3,4]), to the best of our knowledge this is the first study to report on large-scale passive measurement of RTT in a modern 3G network.

2 Measurement Setting

The measurement setting is depicted in Fig. 1. Packet-level traces are captured on the so-called "Gn interface" links between the SGSN and GGSN — for a detailed overview of the 3G Core Network structure refer to [5]. We use the METAWIN monitoring system developed in a previous research project and deployed in the network of a major mobile operator in EU — for more details

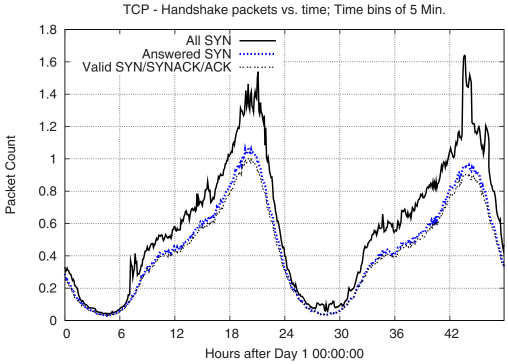

Fig. 2. Time-series of $N_{syn}(k)$, $N_{acked}(k)$ and $N_{valid}(k)$, 5 min bins (rescaled values)

refer to [6]. The monitoring system is able to extract IP packets from the lower 3GPP layers (GTP protol on Gn, see [5]) and discriminate the connections originated in the GPRS/EDGE and UMTS/HSxPA radio sections. In this study we provide network-wide measurements but in principle one can extract separate signals at finer spatial granularity, e.g. for individual SGSN areas or BSC/RNC areas.

The RTT measurement methodology works as follows (ref. Fig. 1). We consider only the TCP connection openings in uplink, i.e. initiated by the client Mobile Stations (MS), for all destination ports. We ignore the downlink connections opened by the Internet hosts — these are present due e.g. to peer-to-peer applications. The elapsed time between the SYN in uplink and the associated SYNACK in downlink is taken as an estimation of the (semi-)RTT in the wired part of the network, between the Gn link and the remote server in the Internet. Hereafter we refer to such quantity as "wired RTT". Similarly, the elapsed time between the SYNACK in downlink and the associated ACK in uplink is taken as an estimation of the (semi-)RTT in the Radio Access Network (RAN), between the Gn link and the Mobile Station. We shall refer to such quantity as "wireless RTT". We extract valid RTT samples only from unambiguous and correctly established 3-way handshakes, and discard all those cases where the association between packet pairs is ambiguous — e.g. due to retransmission, duplication, mismatching sequence number. All valid RTT samples within a measurement interval (e.g. 5 minutes or 1 hour) are aggregated into a logaritmically-binned empirical histogram. Additionally, for each measurement interval k we maintain three global counters: $N_{syn}(k)$ counts the total number of SYN observed in uplink, $N_{acked}(k)$ counts the number of SYN which received a SYNACK reply, finally $N_{valid}(k)$ counts the number of valid RTT samples after filtering out all ambiguous and incomplete sequences. The ratio $r_{inv} \triangleq 1 - \frac{N_{valid}}{N_{acked}}$ represents the fraction of invalid samples over the acknowledged SYN, i.e. the fraction of SYNACK packets that generate a valid RTT sample.

3 Measurement Results

In the following we present measurements taken in January 2009 from a subset of Gn links (exact number undisclosed) of a nation-wide operational network in Austria. Fig. 2 depicts the global counters $N_{syn}(k)$, $N_{acked}(k)$ and $N_{valid}(k)$ computed in 5 min intervals across two days. The values are rescaled in order to hide the absolute volume of connections, as required by the non-disclosure policy with the network operator. The time-of-day profile of network load achieves its peak between 7-9pm, while at night it drops below 5% of the peak. The spikes in the number of total SYN $N_{syn}(k)$ are due to some mobile stations occasionally performing high-rate scanning.

3.1 Wireless Client-Side RTT

Fig. 3 plots the empirical Cumulative Distribution Function (CDF) of the wireless RTT separately for GPRS/EDGE and UMTS/HSxPA. Each graph includes six curves for different measurement intervals of 1 hour each, at different time-of-day. Both empirical distributions are considerably stable in time, with only minor fluctuations between different measurement intervals. The upper tail of the RTT distribution (ref. Fig 3) achieves values as high as a few seconds. Recall from Fig. 1 that the wireless RTT values estimated by SYNACK/ACK pairs include the delay components internal to the client terminal, e.g. packet processing time and I/O buffer delay. In some cases such internal components can be very large. For example, the terminal I/O buffers can become congested due to many paralallel downloads (self-congestion), for example in case of greedy peer-to-peer file-sharing applications. Consider also that some mobile terminals might have limited processing power and/or suboptimal implementation of the TCP/IP stack. Besides terminal-internal causes, large delays can be due to *user mobility*: if the client is moving to another radio cell the incoming downlink packets are buffered in the network — at the SGSN for GPRS/EDGE and at the RNC

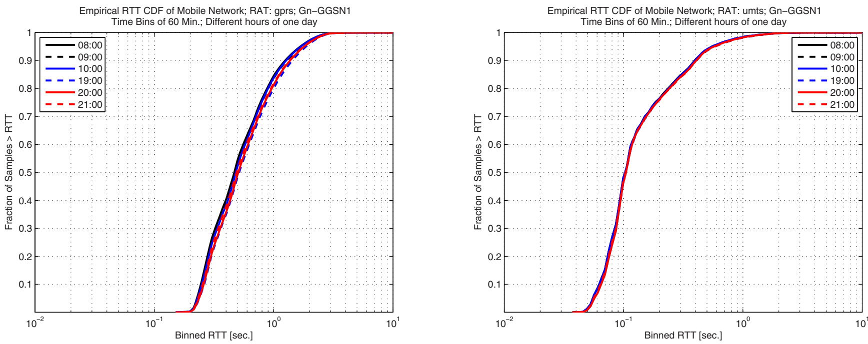

(a) GPRS/EDGE (b) UMTS/HSxPA

Fig. 3. Empirical CDF of wireless client-side RTT, six intervals of 1 hour

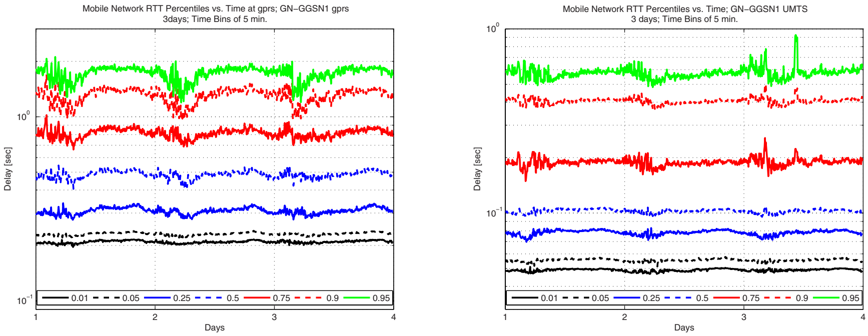

(a) GPRS/EDGE (b) UMTS/HSxPA

Fig. 4. Percentiles of wireless client-side RTT, 5 min bins

for UMTS/HSxPA — until the handover is complete. Another possible source of delay for downlink packets are the *flow control* mechanisms implemented by the 3GPP stack to accommodate temporary dips in the radio channel bandwidth. In a previous study [8] on the same network we have observed that handovers and flow control jointly cause at least 5% of the downlink packets in GPRS/EDGE to remain buffered above 1 second in the SGSN.

Fig. 4 reports various RTT percentiles computed at 5 min granularity over three days, starting from 00:00 of Day 1. The lower 1%-percentile is around 50 ms in UMTS/HSxPA and 200 ms in GPRS/EDGE, while the median values are respectively at 100 ms and 500 ms. Recall that the median of GPRS RTT in 2004 was around 900 ms [1, Fig. 4]. The intervals of higher fluctuation in Fig. 4 correspond to night hours when the number of active MS and network load are very low, and so is the number of RTT samples per timebin.

To complete the overall picture we need to look at the ratio of invalid samples r_{inv}, which is plotted in Fig. 5(a) separately for the two radio technologies. The actual values are surprisingly high: for UMTS/HSxPA it is constantly around 4%, while for GPRS/EDGE it varies from 5% at night to 15% at peak hour. Such values were largely unexpected since we were assuming that the dominant cause for invalid client-side RTT SAMPLES is the loss of SYNACK packets in the RAN. Instead, after a deep exploration of the traces we discovered that the dominant cause is the early retransmission of SYNACK by some popular servers. More specifically, we identified over a hundred servers — all of them within the `google` domain — that were retransmitting the SYNACK packets after only 300-500 ms instead of the recommended timeout value of 3 seconds [10]. This causes an ambiguity when the RTT is larger than the retransmission timeout (see Fig. 5(b)) since the ACK replying to the first SYNACK will be seen after the second SYNACK. In this case it is not possible to associate univocally the ACK to one of the two SYNACK packets, leading to a case of ambiguity that is discarded as “invalid sample” by the current measurement methodology — the same as in [1]. Recalling from Fig. 3(a) that the value of 300 ms falls within

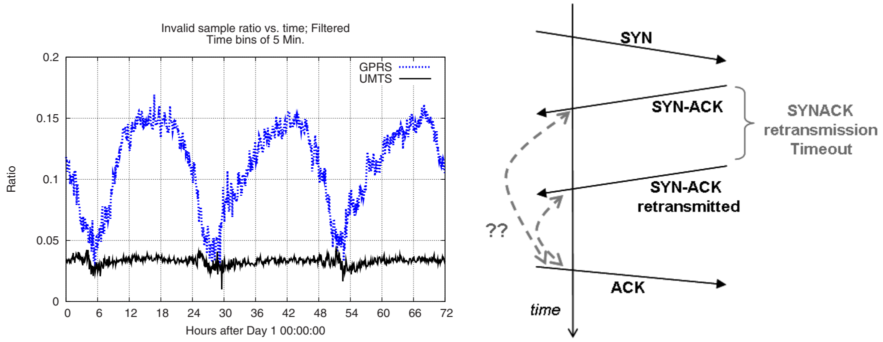

(a) Invalid ratio r_{inv} (b) Early SYNACK retransmission.

Fig. 5. Invalid ratio r_{inv} (a) and example of invalid RTT sample due to early SYNACK retransmission (b)

the range of client-side RTT values for GPRS/EDGE, and considering that a non-negligible fraction of connections are directed to `google` servers, it appears very likely that such high values of r_{inv} for GPRS/EDGE — and to a smaller extent also for UMTS/HSxPA — are due to this phenomenon. Clearly, this is introducing a bias into the RTT statistics, since for a certain share of the SYNACK (i.e. those sent by `google` servers) only the client-side RTT values that are smaller than 300 ms are recorded as valid samples, while the larger ones are discarded as invalid due to the duplicated SYNACK ambiguity.

In order to remove the bias on the RTT measurements we can follow two approaches. The simplest workaround is to ignore all SYNACK/ACK pairs coming from `google` servers, including the unambiguous ones, for example by filtering on the server-side IP address. This has the disadvantage of eliminating a non negligible part of the samples. More importantly, implementing such filtering would require to establish and maintain dynamically a list of filtered hosts, which is hurdle in practice. An alternative strategy would be to develop a method to resolve the duplicated SYNACK case, by inferring probabilistically which SYNACK to pick based for example on the RTT of other samples distribution. We leave the resolution of this problem as a point for further study.

Based on the presented results we can draw some conclusions for UMTS/HSxPA. We have seen that in the monitored network the performances of UMTS/HSxPA do not vary with the time-of-day (ref. to Fig. 4(b) and Fig. 5(a)), which means that they are poorly correlated with the network load. This indicates that the global network capacity is well provisioned. Instead for GPRS/EDGE we cannot draw any conclusion. Although it appears that some level of correlation with time-of-day is present for r_{inv}, this is not sufficient to quantify the degree of correlation between RTT and network load: in principle the daily profile of r_{inv} could be due differences in the traffic mix, and specifically in the relative share of traffic directed to `google`. More work is needed to resolve this issue.

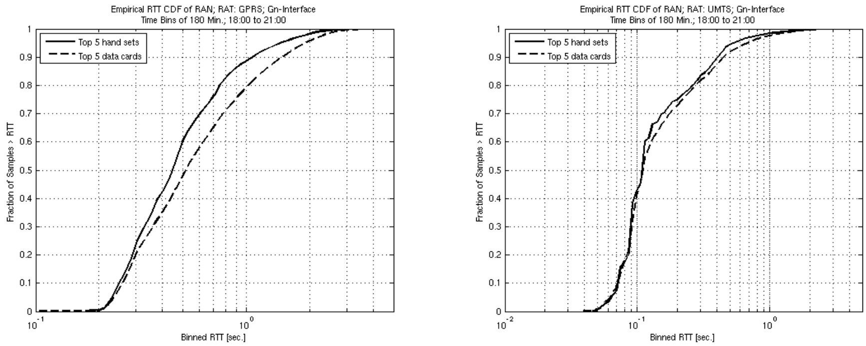

(a) GPRS/EDGE (b) UMTS/HSxPA

Fig. 6. Distinct RTT statistics for handsets and laptops (top-5 TAC in each group)

One interesting feature of our monitoring system [6] is the ability to correlate information extracted at different 3GPP layers and on different interfaces. Among other capabilities, the system is able to extract the Type Allocation Code (TAC) contained in the International Mobile Equipment Identity (IMEI) for each connection. Recall that the TAC identifies the terminal type, therefore we can use such information to extract separate RTT measurements for each class of terminal. In order to investigate whether there are differences between the wireless client-side RTT profile for handsets and laptops, we have extracted the top-5 TAC codes (ranked by the total number of valid RTT samples) for each of these two classes of terminals, and we have computed the RTT statistics separately for each group. The resulting CDFs are given in Fig. 6. In UMTS/HSxPA (Fig. 6(b)) the two distributions are very similar. Instead in GPRS/EDGE it appears that handsets have *lower* client-side RTT than laptops. A first possible explanation is that such difference is just an artifact due to the bias caused by retransmitted SYNACKs: if the share of connections to `google` servers is different for laptop and handsets, also the impact of RTT bias will be different for the two groups. A second alternative explanation is that handset users tend to generate less "aggressive" traffic than laptop users: for example, they tend to browse one page at time instead of opening several parallel pages, avoid visiting heavy websites, do not use peer-to-peer applications. Furthemore, many GPRS/EDGE handsets use WAP. In other words, the traffic produced by individual handsets tends to be smoother than laptops — shaped by applications and user behaviour — therefore producing lower queuing delay on limited bandwidth links. At the time of writing we are taking into consideration both hypotheses, and more exploration of the data is needed to confirm or reject them.

3.2 Wired Server-Side RTT

For the sake of completeness we report in Fig. 7(a) the percentiles of the wired RTT on the server side, for the whole traffic. The lower values in the range

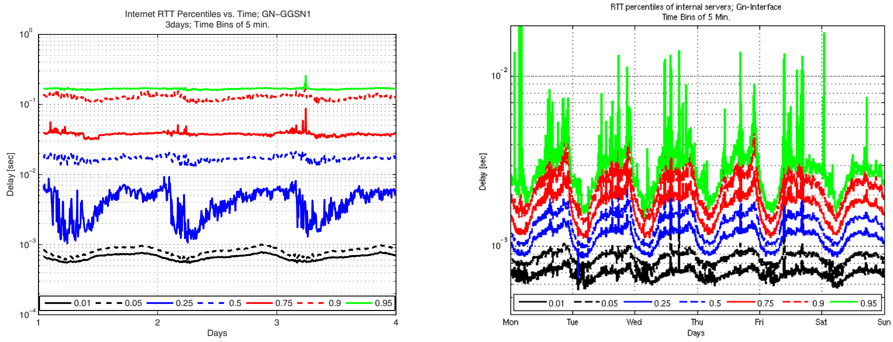

(a) All traffic.

(b) Internal traffic only.

Fig. 7. Percentiles of wired server-side RTT, 5 min bins

of a few milliseconds are partially due to connections terminated at operator's internal servers and proxies, located inside the Core Network, and partially to well-connected external servers, likely placed in the neighborhood of the peering links. It can be seen from Fig. 7(a) that the temporal profile of the 25%-percentile varies with the time-of-day, from 1-2 ms at night up to 6-8 ms in the peak hour. After further explorations we have found that this is due to variations of the external traffic mix, and specifically of the relative share of traffic directed to well-connected servers. In Fig. 7(b) we report the same percentiles only for the internal traffic, i.e. with SYNACK originated from the IP addressess internal to the CN domain. The RTT variations with time-of-day — hence with network load — are modest for the internal traffic, contained within 1-2 milliseconds, which indicates a relatively good provisioning of the internal servers.

4 Conclusions and Future Work

The results presented in this study confirm that modern 3G networks yield considerably lower RTT values than the initial GPRS deployment. We found that the network-wide performances — RTT distribution and invalid samples — are highly stable in time for UMTS/HSxPA, which indicates a negligible correlation with time-of-day and therefore a relative independence on network load. This is a first indication that the monitored UMTS/HSxPA network is currently very well provisioned. We have also shown that the global RTT distribution is essentially the same for handsets and laptops in UMTS/HSxPA.

During the study we have identified a limitation of the adopted RTT estimation methodology, namely the early retransmission of SYNACK packets after only 300 ms by some popular servers in the `google` domain. With the current methodology this leads to ambiguity in the RTT estimation and therefore to sample invalidation. The problem is present particularly on GPRS/EDGE, for which the typical RTT values are in the order of a few hundreds of milliseconds, where a non negligible fraction of samples are discarded. This leads to a bias in the RTT estimation which can not be quantified with our current data. In the progress of our work we intend to develop an effective method to solve the retransmitted SYNACK problem, either by probabilistic SYNACK resolution or by simple host-based filtering.

It is worth remarking that all presented time-series — RTT percentiles and invalid sample ratio — have pretty regular temporal profiles: flat or with regular daily cycles. This simplifies the task of detecting deviations in such signals that might reveal a network problem. On the other hand, it remains to be seen whether such signals can capture network anomalies, and of which kind. The present study could not address this aspect due to the absence of any network incident during the observation period. We are currently deploying on-line passive monitors in the operational network in order to collect long-term RTT measurements (weeks, months), so as to verify whether future network incidents are reflected in deviations of the network signals presented in this preliminary study.

References

1. Benko, P., Malicsko, G., Veres, A.: A Large-scale, Passive Analysis of End-to-End TCP Performance over GPRS. In: IEEE INFOCOM 2004 (2004)
2. Vacirca, F., Ricciato, F., Pilz, R.: Large-Scale RTT Measurements from an Operational UMTS/GPRS Network. In: Proc. of WICON 2005, Budapest (July 2005)
3. Jurvansuu, M., Prokkola, J., Hanski, M., Perälä, P.: HSDPA Performance in Live Networks. In: Proc. of IEEE ICC 2007, Glasgow (June 2007)
4. Barbuzzi, A., Ricciato, F., Boggia, G.: Discovering parameter setting in 3G networks via active measurements. IEEE Comm. Letters 12(10) (October 2008)
5. Bannister, J., Mather, P., Coope, S.: Convergence Technologies for 3G Networks: IP, UMTS, EGPRS and ATM. Wiley, Chichester (2004)
6. METAWIN and DARWIN projects: `http://userver.ftw.at/~ricciato/darwin`
7. Dahlman, E., Parkvall, S., Skold, J., Beming, P.: 3G Evolution: HSPA and LTE for Mobile Broadband, 2nd edn. Academic Press, Elsevier (2008)
8. Romirer-Maierhofer, P., Ricciato, F., Coluccia, A.: Explorative Analysis of One-way Delays in a Mobile 3G Network. In: IEEE LANMAN 2008, Cluj-Napoca, Romania (September 2008)
9. Heikkinen, M.V.J., Kivi, A., Verkasalo, H.: Measuring Mobile Peer-to-Peer Usage: Case Finland 2007. In: PAM 2009, Seoul (April 2009)
10. RFC 1122: Requirements for Internet Hosts — Communication Layers (October 1989)