Large-Scale RTT Measurements from an Operational UMTS/GPRS Network

Francesco Vacirca, Fabio Ricciato, Rene' Pilz
Telecommunications Research Center Vienna (ftw.)
Donau-City St. 1, 1220 Vienna
{vacirca,ricciato,pilz}@ftw.at

***Abstract*— In this paper we present some observations about TCP RTT as captured in the live traffic of an operational GPRS/UMTS network. RTT samples are extracted from traces collected by passive monitoring the Gn interfaces of one of the major providers in Austria, EU. We compare results for GPRS and UMTS, and expose some methodological issues involved in performing this type of analysis, for instance the potential bias introduced by heavy users. We also explore the correlation of RTT with time-of-day.**

The primary motivation for this research is to gain a better understanding of the dominant causes that shape the behaviour and the distribution of RTT across the TCP connections that populate a GPRS/UMTS network. The results presented here are useful for the scientific community, for instance to set more realistic simulation scenarios for other research activities. In addition, RTT measurements as obtained by large-scale passive monitoring can be used in combination with other indicators to build a summary indicator of the performances experienced by the TCP users at the scale of the whole network, i.e. a global RTT-based KPI (Key Performance Indicator).

I. Introduction

In this work[1] we report some observations from an ongoing research on TCP Round Trip Time (RTT) in GPRS and UMTS based on passive measurements from an operational network. The traces are collected by passive monitoring the Gn links of one of the major providers in Austria, EU. The primary motivation for this research is to move towards a better understanding of the dominant causes that shape the behaviour and the distribution of RTT across the TCP connections that populate a GPRS/UMTS network. Such knowledge would pave the ground towards the exploitation of passive RTT monitoring for practical purposes. Possible applications include:

- Detect performance bottleneck on the side of the network and/or of the terminals.
- Optimize parameter setting and resource allocation in the network (e.g., buffer).
- Build more realistic scenarios for simulations.
- Use some statistic parameters extracted from TCP monitoring (e.g., RTT samples) to build a summary indicator of the performances experienced by the users at the scale of the whole network.

The latter point is useful to detect medium- and long-term (from hours to weeks) drifts in the performances caused by macroscopic modifications in the volume and composition of the traffic at the scale of the whole network. Such modifications might be caused for example by changes in the relative composition of the terminal population (e.g., laptop, smartphones, handsets) or by changes in user behavior due to new preferred applications or new billing models. This is particularly important nowadays for 3G networks, which are serving a population of terminals and users highly heterogeneous, growing fast and still far from reaching a stable composition. For instance, in fall 2004 several operators have successfully pushed GPRS/UMTS datacards for laptop into the market. It is evident that usage patterns are inherently different for laptop and handset. Since the relative ratio between the two populations of terminals - therefore of users - has not stabilized yet, we can expect macroscopic shifts in the network-wide traffic composition in the near future, which makes the case for a synthetic monitoring of performances at the scale of the whole network. Among other parameters, RTT samples as derived by passive monitoring in the 3G Core Network can contribute to this goal.

In this paper we report several observations that have emerged so far during our investigations on the RTT in a real operational network, and on the criticalities underlying this type of research.

A. Related works

There are several papers that analyse RTT of TCP connections in wired networks. Most of them rely on the triple handshake of TCP to collect RTT samples of the TCP SYN, a methodology that holds intrinsic limitations as shown for instance in [7]. Only a few works consider RTT samples of TCP DATA packets: for instance in [8] the authors collect and analyse RTT from a trace of eight hours at the University of North Carolina. Other previous RTT measurement studies in wired networks are referenced therein. The results derived from the analysis of RTT in wired networks can not be directly conveyed to 3G due to the huge differences of the two environments, and more specifically to the peculiarities of the 3G network. Regarding RTT measurements in operational GPRS networks, passive measurements are reported in [10] (with data from seven different networks, captured on Gi), while active measurements are used in [9]. To our knowledge, this is the first paper presenting large-scale passive RTT measurements from an operational UMTS network and compare with GPRS. Also, to the best of our knowledge there were no previous papers presenting an analysis of the stability of the RTT statistic over a period of several weeks.

The rest of the paper is organized as follows. Section II presents the sources of RTT variability in a GPRS/UMTS network. Section III provides a brief description of the measurement setting and of the RTT estimation method.

[1]This work is part of the METAWIN project [3] supported by the Austrian research program Kplus and by the industrial partners of ftw.

IEEE COMPUTER SOCIETY

IV report several measurement results and explore the RTT behavior comparatively for GPRS and UMTS. In Section V we conclude, highlighting the open points and some hints for further research.

II. RTT VARIABILITY

In a GPRS/UMTS network there are several dimensions of RTT variability. A first dimension is *temporal variability*: the RTT of a single connection might fluctuate in time due to the following causes:

- Radio channel conditions: fluctuations in the radio channel and errors causes retransmissions at the radio link layer, therefore variability in the packet transfer time.
- Packet scheduling and channel assignement dynamics (e.g. DCH setup delay in UMTS, see [17, pp. 182-189]).
- Queuing and congestion produced by data traffic from/to other terminals in the same cell.
- Queuing and congestion produced by traffic from/to the same terminal: one or more TCP connections from the same terminal might have the congestion window large enough to fill the per-user transmission buffer located in the PCU (Packet Control Unit, a module of the BSS) and/or in the SGSN.
- Voice call preemption: resource units (e.g. time-slots in GPRS) dedicated to data traffic can be pre-empted by voice traffic, reducing the bandwidth available to data traffic.
- Mobility events: in some cases - particularly in GPRS - the handover procedure causes an additional delay in the packet delivery, since the packet is buffered in the PCU/SGS until the mobility procedure is completed.

A second dimension of RTT variability is *terminal variability*, that includes:

- Terminal capabilities (e.g., maximum number of time-slot in uplink and downlink for GPRS).
- Terminal type: handset, smart-phones, palmtop and laptops with 3G datacards can have different processing power, buffering capacity, etc.
- SW version: terminals have different implementations of the protocol stack supporting different options and versions of protocols.

A third dimension is *spatial variability*, that involves aspects related to the configuration of cells and network equipments:

- Radio Channel: radio planning (e.g. cell size, cell population, etc.) and physical layer settings (e.g. Power Control settings, etc.).
- MAC and RRC: resource units allocated to data traffic (e.g. time-slots in GPRS and codes in UMTS), etc.
- Link Layer: Maximum number of transmissions per PDU, configuration of Time-outs, etc.
- Buffer provisioning (e.g. per user buffer vs. per cell buffer, buffer size, etc.).
- SW stack of radio network equipment (e.g., BSC vendor).

In principle all such factors contribute to shape the RTT statistic. However, it is possible that in practice only a few of the above factors are dominant. One of the ultimate goals of our research is to gain a more clear view of whether some sources of variability are dominant over the others. But beforehand, the global RTT statistical behaviour has to be explored and this work offers a contribution towards this goal.

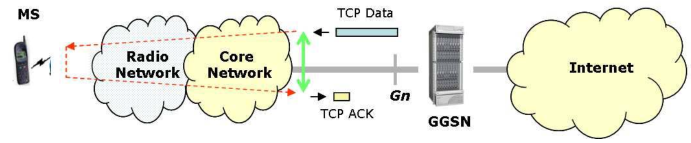

Fig. 1. Semi-RTT between Gn and MS: estimation from passive monitoring of Gn.

III. MEASUREMENT SETTING

The development of a large-scale passive monitoring system - including a parser for the whole protocol stack of the 3G Core Network (3GCN) - and its deployment in the operational network were accomplished within the METAWIN project [3]. While we passively monitor all core network interfaces (Gi, Gn, Gb, IuPS) the results presented in this work are based exclusively on Gn traces. For proprietary reasons we can not disclose several absolute quantitative values (e.g., traffic volumes, number of users, number of Gn links, etc.) nor any other information that might indirectly lead to them (e.g., absolute number of RTT samples). Therefore, for this quantities we will provide only relative values, i.e. fractions, rather than absolute values.

The dataset used in this work is derived by packet traces collected at the Gn interface. All links were monitored, covering 100% of GPRS and UMTS traffic from home subscribers, while traffic from roaming subscribers is not considered. Packets are captured with Endace DAG cards [1] and recorded with GPS synchronized timestamps. For privacy requirements the traces are anonymized by hashing any field related to user identity at the lower layers of the 3G stack (e.g., IMSI, MSISDN, etc.). The traces include TCP/IP headers enabling the analysis of several TCP statistics. For this work we considered more than 3 weeks of traces, from 6th to 30th December 2004.

Based on the collected traces, RTT samples were estimated with a modified version of `tcptrace` [2]. Instead of end-to-end RTTs, we considered the semi-RTTs between the monitored interface (Gn) and the Mobile Station (MS). Investigating semi-RTTs instead of end-to-end RTTs suffices to analyze realistic RTT variations that occur in the GPRS/UMTS network and removes the delay components in the wired Internet. For sake of simplicity, in the rest of this work we will refer the Gn-MS-Gn semi-RTT simply as "RTT". Note that the RTT estimated in this way includes three components: the downlink delay (Gn→MS), the uplink delay (MS→Gn) and the delay component internal to the MS (processing and I/O buffering). Only the first two delay components are accountable as network-dependent, while the latter depends exclusively on the terminal. However, the TCP dynamics, and ultimately the user-perceived performances, will be impacted by the cumulated total RTT.

For estimating the RTT values we exploit the TCP protocol behaviour. The process is depicted in Figure 1: a RTT

IEEE COMPUTER SOCIETY

is defined as the elapsed time $t_{data} - t_{ack}$, where t_{data} and t_{ack} are the timestamps respectively of a TCP DATA packet arriving from the Internet and of the associated ACK, as "seen" on Gn. Only non-ambiguous DATA-ACK pairs are considered to produce a valid RTT sample. The acknowledgment number of ACK must be at least one byte greater than the last sequence number of the DATA packet. Furthermore, it is required that the packet being acknowledged was not retransmitted, and that no packets that came before it in the sequence space were retransmitted after t_{data}. The former condition invalidates RTT samples due to the retransmission ambiguity problem. This is the same procedure that TCP utilizes to estimate the RTT and to set the value of the Retransmission TimeOut (Karn's algorithm, see [5]). The latter condition invalidates RTT samples since the ACK could acknowledge cumulatively the retransmitted packet rather than the original DATA packet.
Before the analysis we filtered out all packets on port tcp:4662 that is used by a popular peer-to-peer file sharing tool (eDonkey). Since it typically runs with many parallel TCP connections it is likely to induce self-congestion on the radio bandwidth and/or on the terminal internal resources (e.g., transmission buffer). The exploratory analysis of such connections (not reported here) displayed RTT values systematically larger than for other traffic, therefore we filtered them out to prevent biases on the whole statistics.
During the exploitative analysis it emerged the presence of a large number of packets directed to ports tcp:135 and tcp:445, mainly TCP SYN in the uplink direction. This is likely due to some self-propagating worms attached to some infected 3G terminals. The presence of such unwanted traffic should be expected in any 3G network since laptops with 3G datacards - often equipped with popular operating system for PC - coexist nowadays with handsets and smartphones in the 3G network, and it is well-known that unwanted traffic is a steady component of the traffic in the wired networks since years (see for instance [4]). The detailed analysis of such traffic and its impact on the 3G network will be covered in a following separate paper. What is important here is that most of such packets did not bring any valid RTT sample while consuming resources in the analysis code (i.e., memory state in `tcptrace`), therefore filtering them out speeds up the analysis process.

TABLE I
SUMMARY STATISTICS OF SOME TCP OPTIONS.

	UMTS		GPRS	
day	6/12	26/12	6/12	26/12
SACK	93.4%	94.3%	91.2%	94.4%
TS	3%	7.8%	49.7%	53.8%
MSS low	1%	1.5%	1.5%	1.4%
MSS high	89.5%	83.4%	88%	87.9%
Wnd low	1.1%	1.9%	10%	9.9%
Wnd med	14.2%	20.6%	8%	6.3%
Wnd high	84.7%	77.5%	81.9%	83.8%

In table I we reported the observed incidence of some TCP options as seen in both GPRS and UMTS. All the values indicate shares of TCP connections. We considered the same metrics of [10, Table I]:

- "SACK" and "TS" indicate the percentage of TCP connections with the SACK [13] and Timestamp [14] options enabled respectively.
- "MSS" refers to the percentage of TCP connections with Maximum Segment Size of 300-600 bytes ("MSS low") and 1300-1600 bytes ("MSS high").
- "Wnd" refers to the percentage of TCP connections with Maximum Advertised Window of $\leq 10KB$ ("Wnd low"), $10-20KB$ ("Wnd med"), and $> 20KB$ ("Wnd high").

It is worth to notice that the TS option is used in GPRS by approximately 50% of the connections, whereas in UMTS it appears less frequently (3-8%). This is a sign of difference between the terminal population in the two networks. We also notice that the values for GPRS are sensibly different from those reported in [10, Table I] which were based on older measurements (at least 1 year) and from different networks.

IV. RTT ANALYSIS

In the first two subsections we analyze and compare the detailed RTT statistics for GPRS and UMTS during a sample day (6th December 2004): the analysis of other days showed that this can be considered as a good representative of the "typical" RTT behavior for both GPRS and UMTS at the time of monitoring. Later in section IV-C we present summary statistics along the whole monitored period, i.e. more than 3 weeks in December 2004.

A. Global distributions

The empirical Complementary Cumulative Distribution Functions (CCDF) of the RTTs captured along one full day are shown in Figure 2. The samples are divided into four groups corresponding to four 6h periods. Note that during the 00:00-06:00 period the global network load is lower, while the busy period is found in the evening (18:00-24:00). From the graphs we observe that the minimum RTT is 476ms and 127ms for GPRS and UMTS respectively. A sensible fraction of RTTs are very large: 70% of the GPRS samples and 10% of UMTS are above 1 second.
It appears that the correlation with time-of-day is rather weak: the RTT distribution does not change much in different periods, despite large differences in the traffic loads. This is consistent with the finding reported in [10, fig. 4] that the difference in RTT between off-peak and peak hours is minimal. Note that the graph in [10, fig. 4] and fig. 2(a) can not be directly compared. First in that work the authors only considered SYN-ACK pair, while here we consider also DATA-ACK pairs, consequently the RTT reported therein were smaller since SYN packets are typically much shorter than full DATA packets, and the packet size has a major impact on the downlink delay component in GPRS. Second, the methodology used in [10] can not capture RTT larger than 3 seconds, with a clear impact on the shape of the CCDF also in the range below 3 seconds due to re-normalization.
However, further exploration of our data showed that there is a 3h period (03:00-06:00) where the global network load is sensibly lower and also the RTT samples are systematically lower than during the rest of the day. This does not app

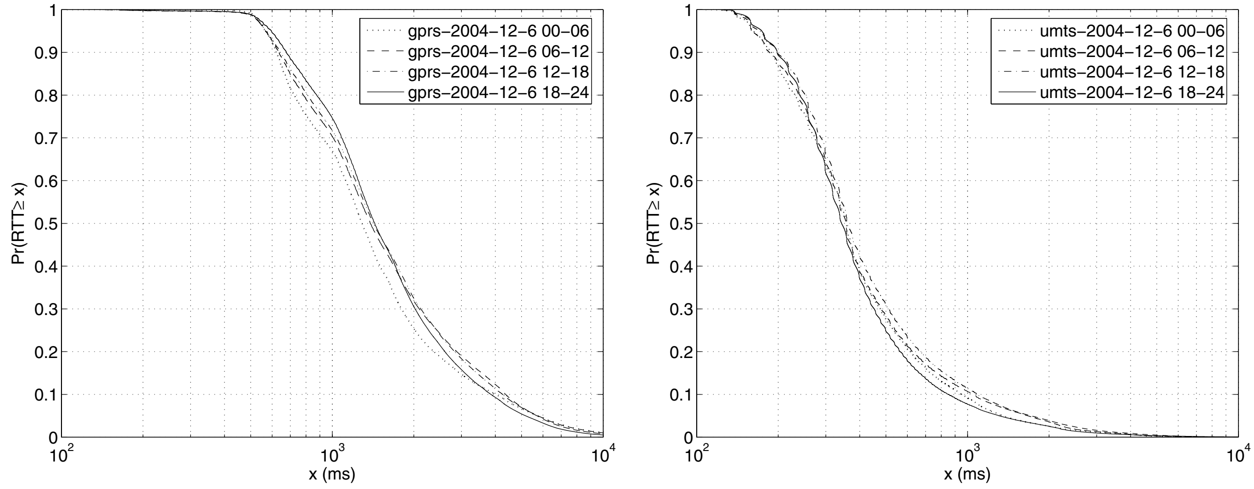

(a) GPRS

(b) UMTS

Fig. 2. Empirical CCDF of RTT for GPRS and UMTS in four different 6h periods (traces from 6th December 2004).

fig. 2 because the low-load period 03:00-06:00 contributes with only a small number of RTT values to the statistic of the 00:00-06:00 period due to lower generated traffic, therefore the samples from 00:00-03:00 dominate the statistic. In fig. 3 we compare the CCDF in the low-load period with the distribution in the rest of the day, and the difference in the RTT distribution emerges more clearly especially for higher values. The lesson to be learned here is that the choice of a 6h period between 00:00 and 06:00 was mistaken, since the phenomenon could not be assumed stationary along the whole 6h. A better choice is to consider shorter integration periods, between 1h and 3h. Also, as expected the global network load plays a role in shaping the RTT distribution in GPRS, while for UMTS the impact of the network load on the RTT was only marginal.
Remind that in GPRS the radio traffic channel is shared among the terminals in the same cell, while in UMTS dedicated channels are adopted (DCH, [17, p. 27]. Such difference certainly plays an important role in the RTT statistics and explains to some extent the different RTT behaviour with respect to the network load that is reflected into a different dependence with the time-of-day, as discussed below in Section IV-C.

B. Correlation with session size

When a 3G terminal becomes active it establishes a so called "PDP-context" with the network that is maintained during the whole activity period. Hereafter we will use the term "session" as a synonymous of PDP-context. During a session the terminal is assigned a dynamic IP address by the GGSN. In principle, after closing a session the same IP address might be immediately re-assigned to another terminal for a new session. However, we had the means to check off-line that in the network under monitoring the level of address re-usage by different terminals is very low within few hours, that means only a very small fraction of IP addresses where used by two or more different terminals within each 6h period.

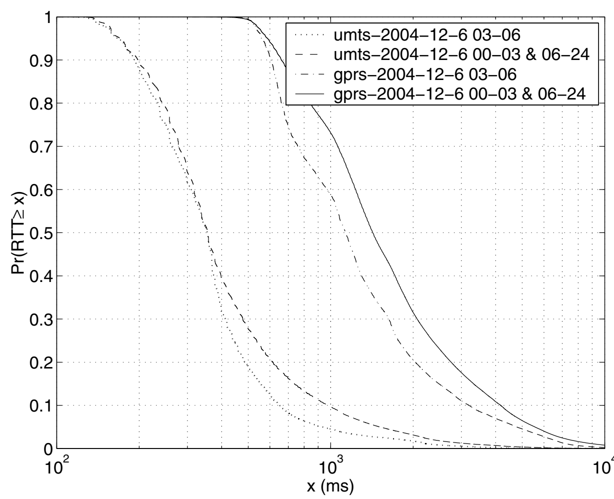

Fig. 3. Comparison of RTT for GPRS and UMTS in low-load period 03:00-06:00 and rest of day (traces from 6th December 2004).

Therefore we will use IP addresses to discriminate sessions, accepting the small error caused by sporadic address re-usage. On such basis, we assign each valid RTT sample to a session (i.e., to an IP address), and denote the total number of samples within a session as the "session size". Note that the session size defined in this way is not the same as the volume of traffic exchanged in the same session, despite the two quantities are highly correlated.
Since sessions are generated by terminals with different characteristics (terminal variability), operated by users with different behaviors and preferred applications, it is expected that the RTT process displays some heterogeneity across different sessions. Additionally, those sessions which generate very high volumes of traffic produce also more RTT samples and might bias the overall statistic. In this section we explore the correlation between session size and RTT.

We display in fig. 4 the qq-plot of the distribution of RTT samples across sessions. We ranked the sessions according to their size within each 6h period: on the X-axis sessions are sorted in ascending order (the top-session is placed at $x = 1$) and the rank is normalized to the total number of sessions seen in that period (whose value can not be disclosed). The value $y(x)$ on the Y-axis indicates the fraction of RTT samples generated by the lower quantile x of session.

For both GPRS and UMTS the distributions display a high level of disparity: a low fraction of sessions accounts for a very large fraction of RTT samples. The authors of [12] propose to use the *crossover* as a synthetic metric for disparity. With the above notation, the crossover is defined as the fraction c for which $y(1-c) = c$. We found for GPRS a crossover of 85/15 during the busy periods and 90/10 during night, while in UMTS it stays around 80/20. The stability in the level of disparity will be analyzed further in section IV-C.

In order to explore the correlation between RTT statistics and session size, we computed the mean and several percentiles of the RTT over all the samples from the lower quantile x of sessions, for progressively increasing x. The results, limited to the mean (top graphs) and 90-percentile (bottom graphs) for all time periods are plotted in Figure 5.

For GPRS the graphs expose a strong correlation between the session size and the RTT: as more heavy sessions are included in the statistic, the global values of RTT clearly drift towards higher value. We do not have yet a full and confirmed explication for this behaviour, and our investigation are still ongoing. However such correlation suggests that self-congestion on the radio bandwidth and/or terminal resources might have a dominant role in shaping the RTT values in GPRS. In order to shed light on this phenomenon our next investigations will address correlation of RTT with other traffic parameters like for example TCP goodput, number of parallel TCP connections, etc.

In case of UMTS the correlation with session size is null, and the RTT summary values (mean, 90-percentile) stay flat for the whole range of x (note that for very low values of x the statistics is highly volatile due to the small number of samples). This result implicitly suggests that in UMTS the RTT is fairly independent of other processes that are indirectly correlated to the session size like for example TCP connection goodput.

It is likely that the different approach to resource allocation between GPRS (shared channel) and UMTS (dedicated channel) is accountable - at least to some extent - of the different dependency between RTT and session size for the two technologies.

C. Long-term analysis

So far we have provided detailed measurement results for one sample day. In this section we present some summary statistics over a longer timeframe, specifically 25 days from 6th to 30th December 2004, for both GPRS and UMTS. Each parameter was computed in time bins of 1 hour. The gap of 12 hours in the last week is due to missing data consequent to a temporary unavailability of the monitoring system.

In fig. 6 we gather four types of graphs for both GPRS and UMTS. The top subgraph reports the behavior of the total traffic volume in bytes, cumulated for uplink and downlink. The volume has been normalized to the peak value in order to avoid disclosing the absolute value. As expected there is a strong 24h periodicity for both GPRS and UMTS.

The next subgraph displays the crossover in the distribution of RTT samples across sessions, a metric that was proposed in [12] to quantify the level of disparity. We found for GPRS a crossover of 85/15 during the busy periods and 90/10 during night, meaning that during the off-peak period (night) the relative incidence of top-talkers (also called “elephants”) is higher, and so is the potential for bias. It is possible that such behaviour is dictated by a relatively flat presence of elephants coupled with cyclical volume of mice, but we are still in the process of investigating the relative proportions of the two classes of users.

For UMTS the crossover was slightly lower than for GPRS, between 80/20 and 85/15.

While in GPRS the crossover seems to be highly stable in time, with a clear periodicity with the time-of-day, the fluctuations of the UMTS crossover appear quite randomly distributed. This is due to the fact that in the monitored network the number of users - therefore of sessions - in UMTS is still lower than GPRS, which explains the higher volatility in the statistics associated to session count - as is crossover. On the other hand in UMTS sessions generate individually higher traffic volumes than in GPRS - due to higher available radio bandwidth - which prevents the volatility to pass into the global RTT statistics.

The next subgraph (second from bottom) displays some summary statistics of RTT within each time bin: mean, median, 5% and 95% percentiles. The mean and the lower percentiles appear rather stable in GPRS, while the higher percentile displays a certain time-of-day dependency related to 24h volume cycles.

In UMTS some large fluctuations are present also for the mean. The largest “atypical” values are found on the 26th. A closer look at the full RTT distribution in UMTS during the 26th is shown in fig. 7. From fig. 7(a) we see that the 10% percentile almost doubled, from 1 to 2 seconds, and the total average increased dramatically. From the second graph (fig. 7(b)) it is immediately evident that this was due to one or two top-talkers (“elephants”) that biased the overall statistic on the 26th. Without them, the overall distribution on the 26th would have been quite similar to the 6th, which was selected as a representative of the “typical” (i.e., unbiased) case. Note that big elephants were also present on the 6th and in several other days, but then they behaved similarly to the other sessions without introducing any bias. In fact, by definition the bias is not necessarily introduced by all elephants, but exclusively by those “elephants that behave differently”. We do not know yet why the elephants on the 26th had such dramatically different RTT values from the rest of the traffic, and we are performing additional investigations in order to assess whether it was due to some special feature of a custom application or instead to particular network conditions.

Incidentally we remark that the type of plot shown i

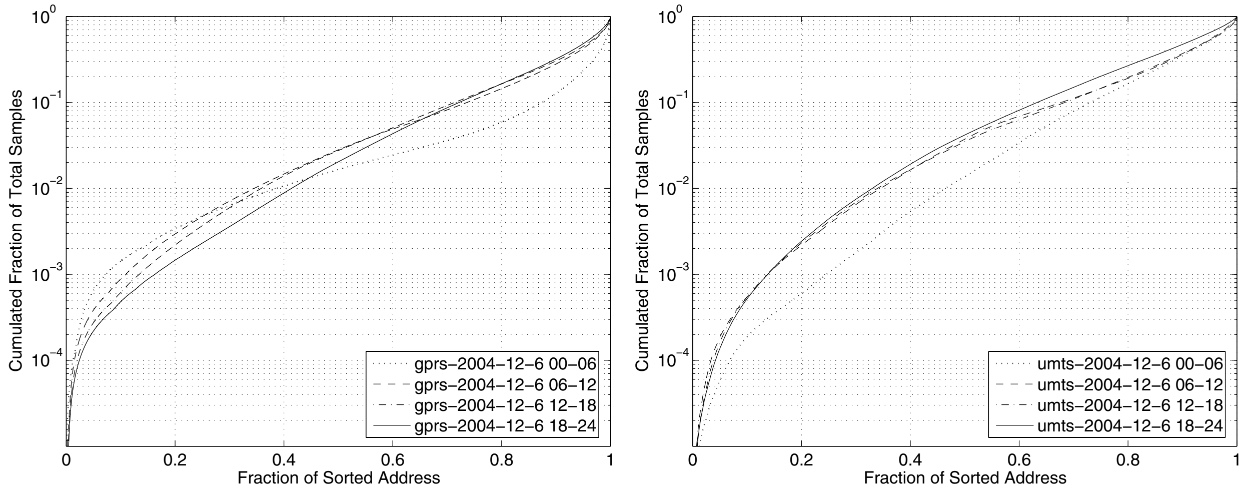

(a) GPRS.

(b) UMTS.

Fig. 4. qq-plot of RTT samples per session in a sample day (6th December 2004).

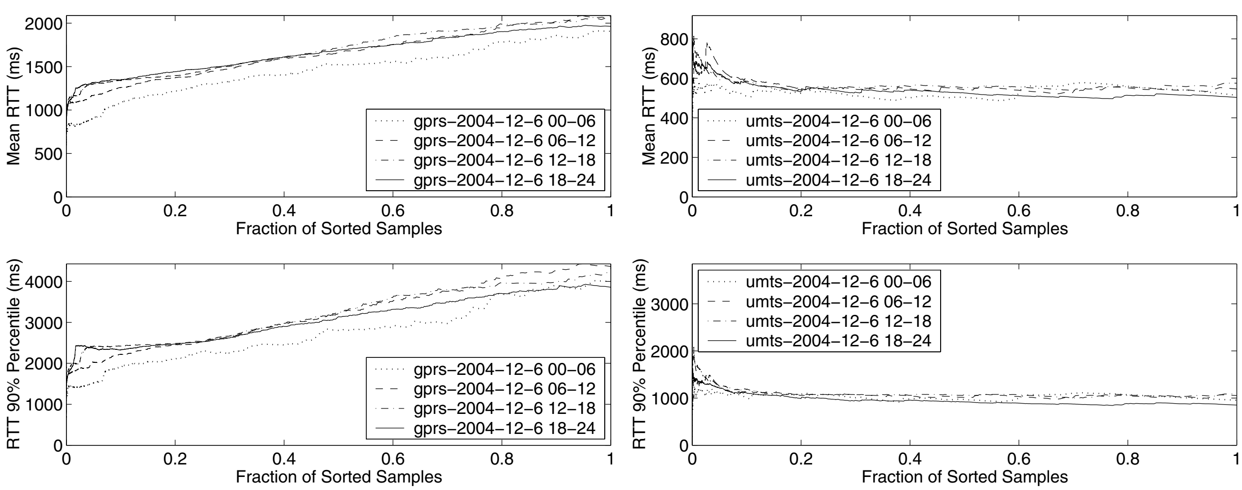

(a) GPRS.

(b) UMTS.

Fig. 5. Correlation with session size: mean and 90% percentile of RTT are computed for different session quantiles.

coupled with the qq-plot of fig. 4 - is a very powerful tool not only for inspecting correlation with size but also to quickly pinpoint bias effects due to elephants.

The previous analysis revealed that a few elephants can bias the whole RTT statistic in UMTS. Therefore it is required to filter them out in order to achieve a RTT statistics that is more representative of the state of the whole network. The problem of identifying the elephants is not trivial and has recently attracted some attention in the literature for different applications (see [15] and references therein). We used in this work a very simple heuristic approach: in each time bin we rank the sessions w.r.t. the product of the session size (i.e. number of RTT samples) times the average RTT, and we filter out the 10 top-ranked sessions. Any other approach for elephant filtering is left to future work.

In the bottom subgraph of fig. 6(b) we plot the same RTT summary statistics after filtering the elephants. As expected the volatility of RTT statistics for UMTS has been greatly reduced and the mean value has stabilized around 500 ms. Some residual 24h cycles are still evident but only for the higher percentile (95%).

For sake of completeness we plot in the bottom subgraph of fig. 6(a) the RTT statistics after the filtering also for GPRS. It can be seen that the overall curve became more regular, with clear 24h cycles (indicating dependency with global network load) and an appreciable decrease of the mean and 95% percentile values.

IEEE COMPUTER SOCIETY

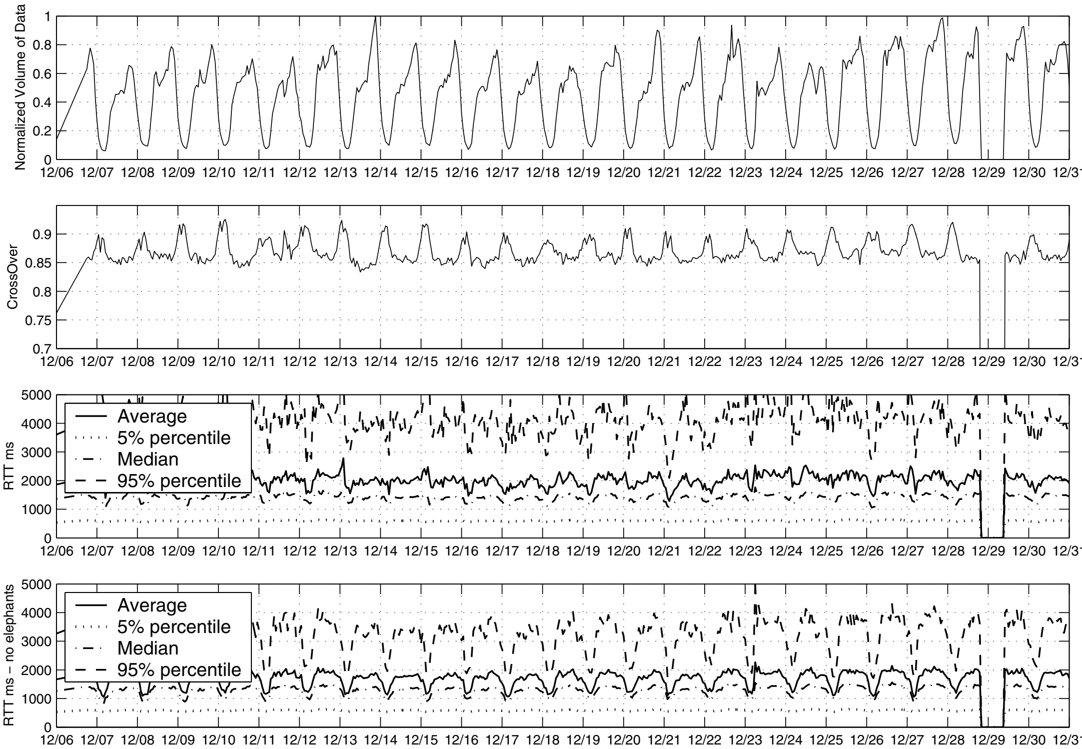

(a) GPRS.

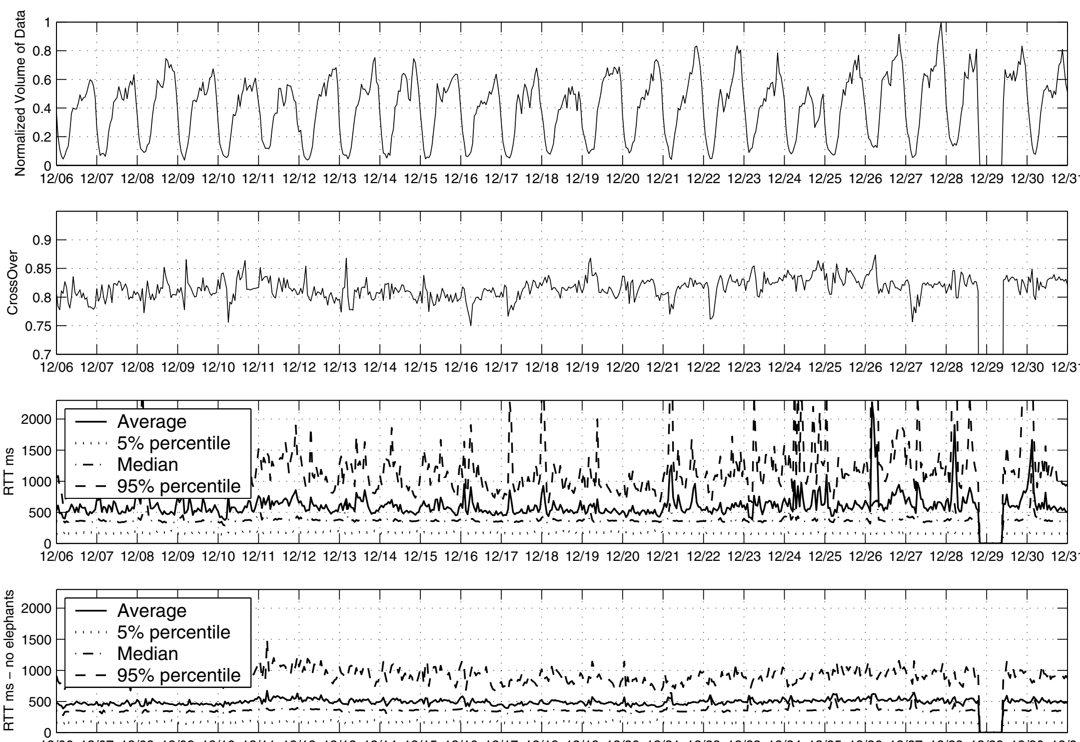

(b) UMTS.

Fig. 6. Summary statistics for 1hour time bins along 25 days. From top to bottom: 1) total traffic volume normalized to peak. 2) crossover quantile of RTT samples per session. 3) Summary RTT statistic. 4) Summary RTT statistic after elephant filtering.

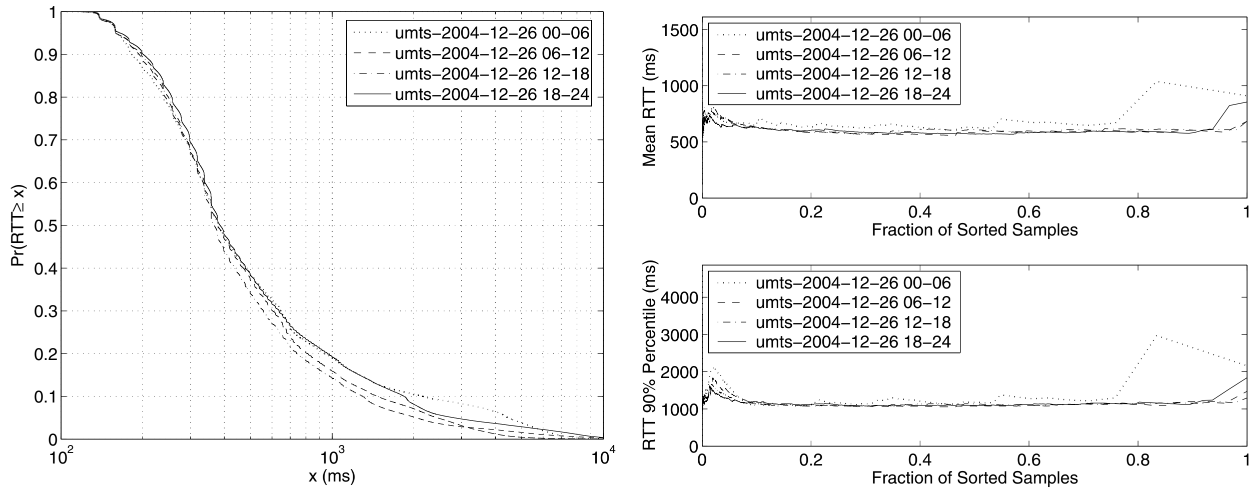

(a) Empirical CCDF of RTT in four different 6h periods.

(b) Correlation with session size: mean and 90% percentile of RTT are computed for different session quantiles.

Fig. 7. UMTS graphs for a different day (traces from 26th December 2004). It is evident the bias induced by the top sessions.

V. Summary of findings and future work

In this work we have reported some observations about TCP RTT as observed in the live traffic of an operational GPRS / UMTS network. Among other findings, we have found a strong correlation in GPRS between RTT and session size, and future research should try to reveal the cause of such correlation.

We found that in UMTS the overall RTT statistic can be biased by a few top sessions (elephants). Since the presence of biasing elephants is sporadic, this introduces a large volatility into the global UMTS RTT statistics. Notably GPRS displays a higher level of disparity, at least as far as the crossover is considered as an appropriate metric for it. On the other hand the long term RTT statistic for GPRS seems less volatile than in UMTS.

As a point for future research it would be interesting to define a synthetic metric for such volatility, in conjunction with other metrics like population size and disparity. This is a necessary pre-condition before any concrete attempt of using RTT measurements as a Key Performance Indicator of the state and performances of the whole network. In our future work we plan to complement the RTT statistics with large-scale measurement of other aspects related to TCP performances, for instance loss probability and retransmission events. As a further step we will extend our investigations to the distribution of RTTs *in space*, i.e. per-cell.

References

[1] Endace Measurememt Systems, http://www.endace.com.
[2] Tcptrace 6.6.1, available at http://www.tcptrace.org.
[3] METAWIN home page: http://www.ftw.at/ftw/research/projects.
[4] R. Pang, V. Yegneswaran, P. Barford, V. Paxson, L. Peterson, Characteristics of Internet Background Radiation, Proc. of the International Measurements Conference (IMC'04), Taormina, Italy, October 2004.
[5] W. Richard Stevens, "TCP/IP Illustrated, Volume 1, The Protocols" Addison-Wesley, 1994.
[6] W. Willinger, D. Alderson, L. Li, "A Pragmatic Approach to Dealing with HighVariability in Network Measurements" IMC'04, Taormina, Italy, 2004.
[7] S. Jaiswal, G. Iannaccone, C. Diot, J. Kurose, D. Towsley, "Measurement and Classification of Out-of-Sequence Packets in a Tier-1 IP Backbone", IEEE INFOCOM 2003.
[8] J. Aikat, J. Kaur, F. D. Smith, K. Jeffay, "Variability in TCP Roundtrip Times" in Proc. of IMC '03, Florida, USA, 2003.
[9] R. Chakravorty, J. Cartwright, and I. Pratt, "Practical experience with TCP over GPRS" in Proc of IEEE GLOBECOM, Nov. 2002.
[10] P. Benko, G. Malicsko and A. Veres, "A Large-scale, Passive Analysis of End-to-End TCP Performance over GPRS" INFOCOM 2004.
[11] H. Jiang and C. Dovrolis, "Passive Estimation of TCP Round Trip Times" ACM Computer Communication Review, August 2002.
[12] A. Broido, Y. Hyun, R. Gao, kc claffy, "Their share: diversity and disparity in IP traffic", PAM'04, April 2004, Antibes Juan-les-Pins, France, April 2004.
[13] S. Floyd, J. Mahdavi, M. Mathis and M. Podolsky, "An extension to the selective acknowledgement (SACK) option for TCP", IETF RFC 2883, July 2000.
[14] V. Jacobson, R. Braden and D. Borman, "TCP Extensions for High Performance", IETF RFC 1323, May 1992.
[15] K. Papagiannaki, N. Taft, S. Bhattacharya, P. Thiran, K. Salamatian, C. Diot, "On the Feasibility of Identifying Elephants in Internet Backbone Traffic", SPRINT Technical Report (http://ipmon.sprint.com/pubs_trs/trs/RR01-ATL-110918.pdf).
[16] F. Vacirca, T. Ziegler, E. Hasenleithner, "Large Scale Estimation of TCP Spurious Timeout Events in Operational GPRS Networks", COST-279 Technical Document, TD(05)003 (submitted).
[17] J. Laiho, A. Wacker, T. Novosad, "Radio Network Planning and Optimisation for UMTS", Wiley, 2002.

IEEE COMPUTER SOCIETY