<http://www.diva-portal.org>

Postprint

This is the accepted version of a paper presented at *2015 European Conference on Networks and Communications (EuCNC), June 29 2015-July 2 2015, Paris, France.*

Citation for the original published paper:

Garcia, J., Alfredsson, S., Brunstrom, A. (2015)
Delay metrics and delay characteristics: A study of four Swedish HSDPA+ and LTE networks.
In: *Networks and Communications (EuCNC), 2015 European Conference on* (pp. 234-238).
IEEE conference proceedings
<http://dx.doi.org/10.1109/EuCNC.2015.7194075>

N.B. When citing this work, cite the original published paper.

Permanent link to this version:

<http://urn.kb.se/resolve?urn=urn:nbn:se:kau:diva-37756>

Delay Metrics and Delay Characteristics: A Study of Four Swedish HSDPA+ and LTE Networks

Johan Garcia, Stefan Alfredsson, Anna Brunstrom

Department of Mathematics and Computer Science

Karlstad University, Sweden

Email: {johan.garcia, stefan.alfredsson, anna.brunstrom}@kau.se

Abstract—Network delays and user perceived latencies are of major importance in many applications in cellular networks. Delays can be measured with multiple approaches and at different protocol layers. This work involves a detailed examination of several delay metrics from a network, transport, and application perspective. The study explores base delay as well as latency under load, capturing also the effect of buffering. The examination is based on a comprehensive active measurement campaign performed in the networks of four Swedish operators. The results show that the delay captured by different metrics can vary significantly, with delay captured from the TCP three-way-handshake and adaptive ping measurements giving the most consistent results for base network delay in our measurements. As expected, when background traffic is introduced measured delay increases by an order of magnitude due to buffering in the network, highlighting the importance of also capturing latency under load when describing network performance. Finally, using an analytic model of flow completion time, we show that well-selected network measurements can provide a good prediction of higher layer delay performance.

Index Terms—Cellular networks, Delay metrics, HSDPA, LTE, 4G, Performance measurements, Completion time

I. INTRODUCTION

This study examines delay metrics at different layers in the communications protocol stack, in the context of cellular networks. It is clear that the network layer delay experienced by packets traversing the network has a large impact on the higher layer perception of a communications service. If a cellular network provider has disproportionately high network layer delays in comparison to other operators, this may have a negative impact on the user level Quality of Experience, which in turn can affect customer retention and churn.

Network layer delay characteristics is a relevant metric in many application scenarios. A particular case where network layer delay is especially significant is for short TCP flows. Their flow completion time is to a large extent dependent on the delay characteristics, rather than the underlying available transmission resources (i.e. bandwidth) in the network. Recent measurements done in the core of an LTE network showed that 90% of all TCP flows were smaller than 35.9 KiB [8]. A major fraction of all traffic thus operates under conditions where it is restricted mainly by delay rather than by bandwidth. To capture the delay dynamics at multiple layers in the protocol stack a comprehensive active measurement campaign has been performed. The campaign covered the four major Swedish cellular network operators. To examine the base network

characteristics, several different approaches to collect network delay metrics were employed. The effects on network delay of concurrently running background TCP traffic, i.e. latency under load, has also been evaluated in dedicated measurement runs. Higher layer measurements were also collected.

The results show that different delay metrics can give markedly different results, indicating the importance of careful metric selection and employing multiple metrics for verification. The study also quantifies the impact of competing traffic on delay, illustrating the importance of measuring latency under load as well as base network delay. The differences between operators in delays at higher layers are also markedly larger when measuring under load. Additionally, the measured short flow completion times are contrasted to a straight forward analytical model of flow completion time using Monte Carlo simulations. This allows for a comparison of how well variations in the measured network layer delay and bandwidth values explain the variation in flow completion times experienced at higher layers.

In the next section we provide a background to the work, followed by a description of the experimental setup. Section IV presents results from the measurements, followed by conclusions in Section V.

II. BACKGROUND

Performance measurement and delay has been of interest in the networking community for a long time. The IETF IP performance metrics (IPPM) working group has published several RFCs of relevance to delay measurements, on round-trip delay [2] and delay variation [3], [11]. In recent years, the concept of bufferbloat [6] has received increased attention. Excessively large buffers in the network causes unnecessary queuing delays when transport protocols tries to saturate the transmission capacity. While throughput-centric applications are largely unaffected by (and are a cause of) increased delay, response-centric applications such as web browsing are heavily influenced. Web browsing is characterized by many short flows to fetch page contents, and the overhead of connection setup and congestion control ramp-up are delay-bound rather than capacity-bound. The user experience of interactive real-time communication, for example video conferencing or VoIP, is also degraded by increased delay. To quantify delay under load, there is an ongoing effort within the bufferbloat community to develop a test suite for real-time response under load

(RRUL) [14]. The characteristics of the network load depend on multiple factors, e.g. on number of concurrent streams, data transfer burstiness, congestion control algorithms, transport protocol and QoS classification, many of which the RRUL test suite takes into account. Several other well known tools and their associated measurement campaigns, such as Netalyzr [10] and the SamKnows and BISMark measurements reported in [13], also evaluate multiple delay metrics, including latency under load as well as base network delay.

Delays are exhibited differently in fixed and cellular access networks, where the latter provide interesting challenges due to aspects of shared spectrum wireless transmission and more complex infrastructure. Previous work studying cellular network characteristics include Alfredsson et al. [1] which examined the 3G, 3.5G and 4G networks of one major Swedish cellular network provider. They examined the presence of bufferbloat in cellular networks and the interaction with different congestion controls, showing that the congestion control can have a significant impact on the induced delay. Using the same measurement data, Garcia et al. [5] examined the TCP protocol efficiency of short flows. Protocol efficiency contrasts the actual TCP flow completion times to an idealized case where the transmission resources provided by the cellular network could be instantaneously and fully utilized. Examining buffering, a study by Jiang et al. [9] identified buffer bloating in all four major US carriers, with round trip times on the order of seconds measured on saturated links. Other measurements in cellular networks include work by Huang et al. [7], [8] examining the performance of several cellular technologies. In particular, while the results in [8] confirm the presence of increased delay due to buffering in LTE networks they also show that many long-lived TCP transfers are unable to fully utilize the available capacity due to protocol inefficiencies. Other work focusing more on delay aspects are reported by Elmokashfi et al. [4]. They present long-term ping-based delay measurements from three Norwegian 3G Networks, with the results indicating large variations in delays both over time and between operators.

While various aspects of delay metrics and delay characteristics have been covered in the previous works described above, this paper contributes a unique coherent treatment of delay characteristics in cellular networks. Our analysis covers and relates delay at multiple communication layers, over multiple cellular technologies and for multiple operators, considering both loaded and unloaded traffic scenarios using a set of different delay metrics. An analytic model that captures the relationship between network layer delay and higher layer delay performance is also introduced.

III. EXPERIMENTAL SETUP

The networks of the four major cellular operators in Sweden were used for the measurement campaign during winter 2014-15. Different run types were used where each run type collected different types of data. An overview of the run types used in the evaluation is shown in Table I which also shows the number of measurement runs analyzed for each run type and operator.

Identifier	Description of run type	Nr of rounds
'ping'	Various types of ping measurements: 1s intervals, flood, adaptive	3.5G: 718, 719, 718, 720 4G: 719, 720, 719, 705
'long'	Long TCP throughput measurements for 20 MiB transfer	3.5G: 697, 712, 686, 588 4G: 697, 706, 697, 686
'short_flow'	Short flow completion times for sizes: 1388 - 327568 bytes, (11 steps)	3.5G: 717, 719, 718, 720 4G: 717, 720, 719, 704
'web'	Total loading time of DN.se webpage 163 objects, 0.01-233 KiB object size	3.5G: 718, 719, 718, 720 4G: 719, 720, 719, 705
'bg_ping'	'ping' with one TCP (Cubic) background flow	3.5G: 240, 242, 242, 242 4G: 242, 242, 242, 241
'bg_sf'	'short_flow' with one TCP (Cubic) background flow	3.5G: 51, 52, 52, 52 4G: 52, 52, 52, 52
'bg_web'	'web' with one TCP (Cubic) background flow	3.5G: 51, 44, 52, 49 4G: 52, 52, 52, 52

Table I: Run types used in evaluation

Two measurement computers were used to perform the experiments: one server with a connection to the Swedish University backbone (SUNET) and one client with four Huawei E392 USB modems, each modem with a sim-card for a specific operator. The modems could be configured to connect using 3.5G or 4G cellular network technologies. Both computers were running Ubuntu 12.04.2 LTS with linux kernel version 3.16.3 (installed from source), and with Cubic used as congestion control. Packet traces were collected using tcpdump and processed with tcptrace to extract transport layer metrics. Individual TCP flows which were not complete TCP connections, were of incorrect length, or which had SYN or FIN retransmission were not included in the data set. TCP segment offloading (TSO), TCP metric caching and other related functionality was turned off or flushed to ensure that the various optimizations that are performed in the OS networking code and in device drivers do not interfere with the measurements. Other TCP related variables were left at their default values. The initial window size was 10 segments, and SACK was enabled. Netperf was used to generate the traffic for the short and long flows at both end-points, whereas the web page load time tests used an Apache Web server along with the wget utility at the client-side.

IV. EXPERIMENTAL RESULTS

A. Base network characteristics

One common way to describe network characteristics is with the end-to-end delay and throughput characteristics that are observed by traffic flows in the network. Network delay characteristics can be collected in multiple ways, and one way is to use timing information from the three-way handshake (3WHS) performed by TCP at connection establishment. In this case the 3WHS values are extracted from the short_flow measurements. Delays can also be explicitly measured using ICMP ping packets, and here three variations were employed: ping messages sent every second, flooding the network with 100 ping packets per second, and adaptive ping which generates a new ping packet when the reply for the previously emitted ping is received. The delays captured by all four measurement methods are shown in Figure 1. The figure shows boxplots with the median indicated as a horizontal bar, and

(a) 3.5G network

(b) 4G network

Figure 1: Measured run averages using different RTT measurement approaches

Figure 2: Average TCP throughput for long flows

the mean as a square box. Measurement points more than 1.5 interquartile range (IQR) from the corresponding quartile are marked with plus signs. Each measurement point represents a single probing for 3WHS, the mean of 10 ping probes for 1 sec ping, and the mean of 60 ping probes for adaptive and flood ping.

As can be seen in the figure, there is a considerable difference in the results for the different measurement approaches in the 3.5G network. The 1s and flood ping approaches for operators 1 and 3 show considerably higher median and mean RTT than the other measurement approaches. This is likely related to radio resource management differences, as frequent radio state time outs may incur extra delay for 1s pings if they continuously switch from FACH to IDLE during the 1s between each ping. Such radio state timeout values can vary between operators due to different configuration settings. Here it is relevant to note that in the 3.5G network operators 1 and 3 have shared infrastructure cooperation, while operators 2 and 4 have semi-shared infrastructure. For the 4G measurements, the difference in the mean and median is not showing any considerable variation between the different approaches or operators. For the 4G network, operators 1 and 2 have shared infrastructure cooperation.

Figure 2 shows the throughput measurements from the long runs, and as expected the 4G network obtains better throughput for these long TCP flows. As can be seen, the throughput of the long flows is roughly doubled in the 4G network. As will be seen when examining web traffic below, a similar increase in performance can not be expected for shorter flows.

B. Network delays with concurrent flows

Previous studies such as [1] have shown that the presence of concurrently running TCP background traffic may cause bufferbloat and thus severely impact the delays experienced by new flows. Bufferbloat occurs when excessive buffers in the network are consistently in high occupancy, thus leading to inflated delays for all traffic. Out of the delay measurement approaches discussed in the previous sub-section, the 3WHS and adaptive ping approaches were used to evaluate the network delay effects of concurrent background traffic. The `bg_ping` and `bg_sf` runs were used to collect data, shown in Figure 3. These measurements could model a case where a mobile terminal has an ongoing long-running TCP transfer, which could for example be an app update. An alternative use case is when a cellular modem is used as a home router serving multiple users, and one user downloads a large file. As can be seen in the figure, the delays are considerably increased both in the 3.5G and 4G networks when a background TCP flow is present. While not quite as severely affected as the 3.5G network, the 4G network also displays a considerable increase in delay in our measurements.

With regards to the different measurement approaches, in the 3.5G measurements operator 2 shows a difference between adaptive ping and 3WHS for the case with background traffic (considering the mean/median). Differences in the buffering parameterization/prioritization may be one explanation for this difference.

C. Web page completion times and concurrent flows

Considering higher level delay metrics, web page completion time is a useful metric. Here, page completion time is measured by the amount of time it takes to download all web objects of a web page of the Swedish newspaper site `dn.se`. To download this site a total of 4.1 MiB must be retrieved, split between 163 different objects ranging in size from 12 bytes to 233 KiB. The web page completion time is measured using one TCP connection for each web object, and four parallel download threads. The web page completion time is the time from the start of the first flow until the completion of the flow for the last of the objects. The objects were downloaded from a page mirror located at our measurement server.

(a) 3.5G network

(b) 4G network

Figure 3: Measured RTT run averages, without and with one simultaneous background flow

Figure 4: Mean web page completion times without and with one background flow, with 95% confidence intervals

The resulting page completion times are shown in Figure 4. Without background traffic, the differences between operators and networks are minor. When comparing the cases of with and without background traffic, there is a large difference. The increase in network layer delay seen with background traffic is also strongly reflected at the application level. The individual differences in performance between the operators are also markedly larger when there is background traffic, in relation to when there is no background traffic.

D. Comparison to analytically derived flow completion times

This subsection focuses on the transport layer, which sits in between the previously discussed network and application layers. The delay metric of interest here is the flow completion time of single, relatively short, TCP flows.

Using the basic network characteristics and an appropriate analytical model it is possible to examine to what extent the actual measured completion times are consistent with the measured network characteristics. When choosing an appropriate model for the expected completion time it is important to note that only a small fraction (<3%) of all short flows in the data set experienced any packet loss at all, so a loss based TCP model is not the most appropriate. Further, considering that it is short flows that are of interest, the model must faithfully capture relevant aspects of the TCP connection establishment and slow start phases. Models of short flow completion time have previously been used as part of an analytical examination of TCP protocol efficiency in [5], and for evaluating fast

startup schemes in [12]. The model used here derives the flow completion time as a function of flow size, maximum segment size, round-trip time, available capacity, and initial window size. Larger initial TCP windows, as used in current TCP implementations, has a considerable impact on the completion times of short flows. The model includes two RTTs for the connection setup and one RTT for the connection teardown. The expression for the estimated flow completion time $\hat{L}_S$ is:

$$\hat{L}_S = 2R_S + \frac{O}{C} + P \left(R_S + \frac{S}{C} \right) - (2^P - 1) \frac{IS}{C} + R_S$$

where

$$P = \min \left(\max \left(\left\lceil \log_2 \left(\frac{R_S + S/C}{IS/C} \right) \right\rceil, 0 \right), \left\lceil \log_2 \left(1 + \frac{O}{IS} \right) \right\rceil - 1 \right)$$

L	Flow comp. time	(s)	C	Available capacity	(bits/s)
R_S	RTT	(s)	S	Segment size	(bits)
O	Flow size	(bits)	I	Initial win	(segments)
P	Rounds w. idling	-			

The RTT value R_S used for single flows is the base network delay when the user has no concurrent connection, i.e when buffers are empty. Inclusion of a configurable initial window size is provided by I , which is set to 10 in our analysis. The variable P is the number of transmission rounds with idle time, i.e when the received rate is constrained by the senders congestion window and not the available transmission resources. The value of P is bounded either by the relationship between RTT, throughput and initial window size, or by the total amount of data in the short flow.

The analytical model is used to perform straight-forward Monte-Carlo simulations of the expected completion times for short flows over a range of flow sizes. The simulations are performed using data from the 3WHS delay measurements in the short_flow runs, and the throughput values collected in the long runs, as input. The resulting analytically derived expected short flow completion times are shown as heatmaps in Figure 5.

Overlaid in the subfigures are also the actual measured short flow completion times as experienced by the same short_flow flows whose 3WHS estimation was used in the analytical model. The 11 point bars shows each measurement value collected for the 11 different flow sizes. Comparing the figures it is clear that the analytically derived estimation is reflected in the actual measured values, implying that lower level delay metrics can provide a useful indication of higher

Figure 5: Heatmaps of analytically derived short flow completion times with actual measurements overlaid.

layer performance in cellular networks. In the right part of Figure 5a it can also be seen that the greater spread in estimated completion times is reflected in the actual measured values.

Considering Figure 5b, the stepwise behavior to the left is a result of the TCP congestion window not being sufficiently large during start-up, thus forcing extra idle delay each time the flow needs to send a window of data. It can be noted that an initial window size of 10 was used for these results, and a lower window size would lead to both smaller and more steps in the expected completion times.

For the 4G network results shown in Figure 5b and 5c the number of steps present for the analytical estimation is larger than in Figure 5a, which is a result of the higher available bandwidth present in the 4G network. The higher bandwidth requires additional TCP transmission rounds before the window is sufficiently large to fully utilize the bandwidth. So while 4G networks provide higher throughput, this cannot be fully exploited by the short flows which form the majority of TCP flows carried over cellular networks.

V. CONCLUSIONS

Using large scale measurements in the 3.5G and 4G networks of four Swedish operators this study has examined delays and delay metrics at different layers. Considering the network layer delays, four different approaches to measuring network layer delays were examined, with adaptive pings and 3WHS measurements showing the most consistent results.

Short TCP flow lengths are the norm in cellular networks [8], and this study also examined the transport layer metric of short flow completion times over a range of flow sizes. The presence of concurrent traffic is an important factor for cellular network delays at all considered layers. In contrast to some earlier studies, these results show that the presence of a background flow increases network delay in 4G networks, and not only in 3.5G networks. The increase in delay is clearly evident at both the network and application layer, with web page completion times increasing up to an order of magnitude in the presence of a single background TCP flow.

Based on analytical estimation using network layer data a comparison was done to actual measured short flow completion times. The results show that network layer measurements and an appropriate analytical model can serve as a useful

predictor of the achieved transport layer performance in a cellular network context. The results also highlight a future research direction towards using statistical tests on delay distributions as a means for shared infrastructure detection across operators.

ACKNOWLEDGMENT

This work was partly funded by .SE Internetfonden. The views expressed are solely those of the authors.

REFERENCES

- [1] S. Alfredsson, G. Del Giudice, J. Garcia, A. Brunstrom, L. De Cicco, and S. Mascolo. Impact of TCP Congestion Control on Bufferbloat in Cellular Networks. In *Proc. World of Wireless, Mobile and Multimedia Networks (WoWMoM'13)*, 2013.
- [2] G. Almes, S. Kalidindi, and M. Zekauskas. A Round-trip Delay Metric for IPPM. RFC 2681 (Proposed Standard), Sept. 1999.
- [3] C. Demichelis and P. Chimento. IP Packet Delay Variation Metric for IP Performance Metrics (IPPM). RFC 3393 (Proposed Standard), Nov. 2002.
- [4] A. Elmokashfi, A. Kvalbein, J. Xiang, and K. R. Evensen. Characterizing delays in Norwegian 3G networks. In *Proc. Passive and Active Measurement (PAM'12)*, Berlin, 2012. Springer-Verlag.
- [5] J. Garcia, S. Alfredsson, and A. Brunstrom. A Measurement Based Study of TCP Protocol Efficiency in Cellular Networks. In *Proceedings of 12th International Symposium on Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt)*, 2014.
- [6] J. Gettys and K. Nichols. Bufferbloat: Dark Buffers in the Internet. *Communications of the ACM*, 55(1):57–65, 2012.
- [7] J. Huang, F. Qian, A. Gerber, Z. M. Mao, S. Sen, and O. Spatscheck. A Close Examination of Performance and Power Characteristics of 4G LTE Networks. In *Proc. Mobile Systems, Applications, and Services (MobiSys'12)*, pages 225–238, New York, NY, USA, 2012. ACM.
- [8] J. Huang, F. Qian, Y. Guo, Y. Zhou, Q. Xu, Z. M. Mao, S. Sen, and O. Spatscheck. An In-depth Study of LTE: Effect of Network Protocol and Application Behavior on Performance. *SIGCOMM CCR*, 43(4):363–374, Aug. 2013.
- [9] H. Jiang, Z. Liu, Y. Wang, K. Lee, and I. Rhee. Understanding bufferbloat in cellular networks. In *Proc. ACM SIGCOMM CellNet Workshop*. ACM, 2012.
- [10] C. Kreibich, N. Weaver, B. Nechaev, and V. Paxson. Netalyzr: Illuminating the edge network. In *Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement, IMC '10*, pages 246–259, New York, NY, USA, 2010. ACM.
- [11] A. Morton and B. Claise. Packet Delay Variation Applicability Statement. RFC 5481 (Informational), Mar. 2009.
- [12] M. Scharf. Comparison of end-to-end and network-supported fast startup congestion control schemes. *Computer Networks*, 55(8):1921–1940, June 2011.
- [13] S. Sundaresan, W. de Donato, N. Feamster, R. Teixeira, S. Crawford, and A. Pescapè. Broadband internet performance: A view from the gateway. In *Proceedings of the ACM SIGCOMM 2011 Conference, SIGCOMM '11*, pages 134–145, New York, NY, USA, 2011. ACM.
- [14] D. Täht. Realtime Response Under Load (RRUL) Test (draft .07), <http://www.bufferbloat.net/projects/codel/wiki/RRUL-test-suite>.