A Measurement Based Study of TCP Protocol Efficiency in Cellular 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*—This paper examines the efficiency of resource utilization with respect to short-lived TCP flows in various cellular networks. The examination is done from the vantage point of an end-user who would like to use as much as possible of the cellular transmission resources that are available at any given time, thus minimizing the delays associated with communication. Based on a comprehensive measurement campaign we first derive network characteristics with regards to base RTT, RTT under load, and average throughput. A protocol efficiency metric is introduced to capture how efficiently short TCP flows are in fact able to use the instantaneously available transmission resources in a cellular network. The measurements show that short TCP connections have low efficiency in 3.5G (HSPA+) and 4G (LTE) mobile broadband networks, and that the improved latency and throughput characteristics of 4G in relation to 3.5G nevertheless results in lower short-flow efficiency for 4G.**

I. Introduction

The inter-relationship between the different layers in the protocol stack can have a considerable impact on the end-user experience when utilizing different types of wireless networks. In this paper we use cellular network measurements in a Swedish cellular network to examine the inter-relationship between the transport layer and the underlying characteristics of the network with regards to the ability of the TCP protocol to efficiently utilize available network resources. The approach assumes that it is the cellular hop that is the constraining part of the end-to-end connection, and that short TCP flows typical in web browsing are of primary interest. Recent data collected inside the core of an LTE network shows that 90% of all TCP flows are smaller than 35.9 KiB [9].

It is possible to make an idealized model of an optimal transport protocol which is able to fully utilize the available transmission resources provided to the user by the cellular network. By contrasting such an ideal model to a model of TCP protocol behavior it is possible to derive a protocol efficiency metric. In this paper we present an analytical model for TCP protocol behavior including the effect of larger initial window, which has previously not been analytically examined. We use the collected data on underlying network characteristics of throughput and RTT to drive a Monte-Carlo simulation that provides protocol efficiency metrics over a range of flow sizes.

The results show that transport layer behavior indeed interacts with the buffers of cellular networks for the 3.5G case, and that buffer bloating [7] have a detrimental effect on both response time and protocol efficiency when multiple flows are present concurrently in 3.5G networks. On the other hand, the 4G network measurements show very little signs of buffer bloating. However, the results show that although both throughput and delay improves considerably between 3.5G to 4G, the protocol efficiency is actually worse for 4G than 3.5G. For short flows, the inability of TCP to efficiently utilize the available resources becomes even more pronounced under the improved 4G conditions than under 3.5G conditions.

In the next section we provide a background to the work, followed by a description of the measurement setup. Section IV presents results from the measurements, while section V elaborates on the protocol efficiency metric and associated results, followed by a summary and future work.

II. Background

We have previously presented work based on data collected in a measurement campaign in the 3G, 3.5G and 4G networks of a major Swedish cellular network provider [1]. The focus of that work was on the impact of different congestion controls (CCs), and the presence of buffer bloat in cellular networks. The results showed that the impact of the utilized congestion control, i.e. CUBIC, NewReno, and Westwood+, was not a major factor when there was one single flow utilizing the link. However, when shorter web flows were mixed with longer background flows the CC of the background flows was shown to have a major impact in 3G and 3.5G environments. Other measurements in cellular networks include work by Huang et al [8], [9], [10] examining the performance of several cellular technologies, as well as work by Sommers et al [16] which measure cellular access characteristics and compare to WiFi access. Other work focusing more on delay aspects are reported by Elmokashfi et al [6]. A study by Jiang et al [11] identified buffer bloating in all four major US carriers, where round trip times on the order of seconds were measured on saturated links.

In this paper the collected measurement data is used in a measurement-based modeling approach to examine TCP protocol efficiency. One of the early efforts in TCP modeling was done by Savage et al [2]. Some efforts have been made by Lee et al [12] to provide a model that includes varying TCP initial window size. Previous work on protocol efficiency has been done by Liu et al [13] which quantifies the aggregate cellular network capacity loss which occurs due to the inability of TCP to utilize all the bandwidth available. In contrast, our work focuses on protocol efficiency from the end-user perspective rather than as an aggregate. Huang et al [8] examine efficiency mainly from the perspective of increased power consumption, and conclude that the LTE network performance bottleneck for web-based applications lies mainly in the processing power of the end-user devices, as well as the underutilized network capacity resulting from protocol inefficiencies for the small object size typical in web transactions. Paul et al [14] observe that all TCP applications have significantly poorer effective bit rate compared to nominal bit rates of the underlying physical layer technology, implying that there is significant scope of protocol improvements across layers. Approaches to increase the protocol efficiency of short TCP flows include preferential treatment as suggested by Chen et al [3], larger initial windows as proposed by Dukkipati et al [4], and the proposal by Radhakrishnan et al [15] for TCP fast open. Eklund et al [5] provide a theoretical analysis of an ideal start-up scheme.

III. Measurement setup

This section describes how measurements were collected and processed, along with the metrics used for evaluation.

A. Collection of measurements

The experimental campaign was carried out in the network of one of Sweden's main cellular providers employing three different cellular network technologies (3G, 3.5G, 4G). Two measurement computers were used to perform the experiments: one server with a connection to the Swedish University backbone (SUNET) and one laptop with a Huawei E392 USB modem configured to connect using different cellular network technologies. This avoids the problem of large variations occurring due to the particular cellular communications hardware used in experiments as reported in previous work [6]. Both computers were running Ubuntu 12.04 with kernel versions 3.1.4 on the laptop and 3.2.0-27 on the server.

A set of measurements were performed to collect data both on the network characteristics for short 320 KiB flows and 3-minute long flows, as well as when short flows were competing with one or five long-lived background flows. Packet traces were collected using tcpdump and processed with tcptrace to extract transport layer metrics. 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 collection of measurement data. The initial window size was 10 segments, and SACK was enabled. Iperf was used to generate the traffic for the long-lived background flows at both end-points, whereas for the short-lived flows an Apache Web server was used along with the wget utility at the client-side to emulate Web browsing. Additional details on the data collection are provided in [1].

B. Experimental design

A wide range of measurements were made based on the run types illustrated in Figure 1. Runs A and B collect baseline information by measuring transfer characteristics for

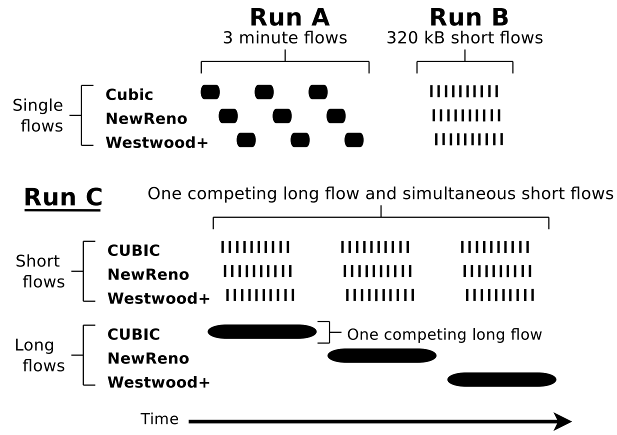

Figure 1: Configuration of measurement runs

long and short lived flows without any competing traffic. Run A uses long-lived flows to collect data for 3 minutes for each of the three different congestion control algorithms, and with three replications each. Similarly, run B collects data for the 320 KiB short flows, but with 10 replications. In contrast to runs A and B, run C involves both a long-lived background flow and short-lived flows. Additionally, a run D (not shown in Figure 1) using five instead of one concurrent long flows was also conducted. Thirteen measurement runs were concatenated into a single measurement campaign with the following run composition: ABCDABCDABCDA, further increasing the number of replications. One such campaign was performed for each of the three different network technologies. Thus, more than 1800 short flow measurements were collected in these campaigns.

C. Metrics

The collected data allows extraction of metrics at different levels in the protocol stack. At the most basic level, an estimate of raw link throughput can be provided by measuring the throughput of long running TCP flows, and end-to-end base RTT is measured by the initial TCP 3-way handshake (3WHS-RTT). As shown in [9] most of the flows transferred over cellular networks are small, making it especially relevant to consider short TCP transfers. These are more impacted by TCP's inability to fully use the available bandwidth during the initial part of the slow start. Here we use the short flow completion time as the metric for short flow performance, using a flow size of 320 KiB. This size is representative of the base HTML page of popular news sites such as Dagens Nyheter in Sweden or Huffington Post in the US.

Protocol efficiency is used as a metric to measure the ability of a TCP flow to immediately utilize the transmission resources available to the user in a cellular network. At 100% protocol efficiency, a flow would be fully utilizing the available resources directly after the shortest possible initiation of the transfer. In the typical scenario where a connection is established from a mobile client to a server and the client requests some data, the client currently will experience a delay of two RTTs before the connection has

Study	Our Results			LS LTE [9]	3GTest [10]	4GTest [8]	SpeedTest [16]		
Time	Sept - Oct 2012			October 2012	Aug-Dec 2009	Oct-Dec 2011	Feb - Jun 2011		
Location	Karlstad, Sweden			1 US Metro Area	Across U.S.	Across U.S.	NY City	Madison, US	March. UK
Type	3G	3.5G	4G	LTE Only	Four 3G ISPs	LTE	Various Cellular		
5% TCP DL	93	2601	4930	569	74 - 222	2112	108	99	28
50% TCP DL	151	3521	5830	9185	556 - 970	12740	1678	895	1077
95% TCP DL	193	16272	57771	24229	1921 - 2943	30812	12922	3485	3842
5% 3WHS RTT	237	49	36	30	125 - 182	37	21	99	34
50% 3WHS RTT	256	50	40	70	160 - 200	70	54	184	92
95% 3WHS RTT	277	70	43	467	645 - 809	127	336	773	313

Table I: Measurement results and comparison to other measurement studies (partly from [9]).

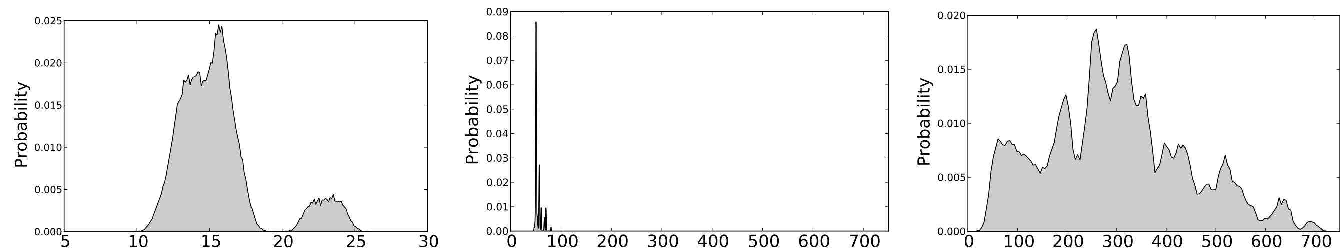

(a) Throughput (Mbps) (b) 3WHS-RTT (ms) (c) 3WHS-RTT with background flow (ms)

Figure 2: 3.5G path characteristics

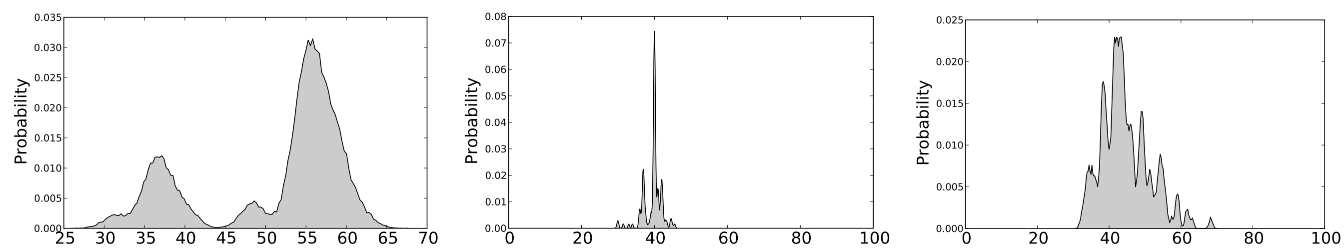

(a) Throughput (Mbps) (b) 3WHS-RTT (ms) (c) 3WHS-RTT with background flow (ms)

Figure 3: 4G path characteristics

been established and the initial data is received. For a fully efficient protocol the server would receive the request during the initial connection setup and be able to immediately send at the available throughput. The client thus would wait only one RTT to complete the connection establishment and receive the initial data, instead of two. Such idealized behavior is unfeasible given current network and protocol architectures, but does serve as a useful reference point. Efforts to increase the efficiency include proposal to allow data transfer during connection setup [15], and increased initial window sizes as discussed in [4].

IV. Measurement results

To relate the measurements obtained in our measurement campaign to other measurement studies, an overview of our measured results along with related results are presented in Table I for the two basic metrics TCP downlink throughput (kbps) and 3WHS-RTT (ms). While the numbers are useful for providing a context to our obtained results, consideration should also be made to the differences in approaches used for obtaining the reported results. Our measurements reported in Table I covers runs A and B, i.e. includes both short and long flows but with no concurrent background traffic. One aspect is that for our measurement setup, a majority of the flows were short flows, which leads to markedly lower 5% percentile and mean values, as compared to the 95% percentile which represents values from the long-running flows which have much better protocol efficiency. The LS LTE study [9] collected measurements for around 47 million flows of live user traffic from inside the core network of a US cellular operator, yielding a realistic data set but also providing less experimental control. The values reproduced in Table I are only considering flows larger than 1 MB. Results in the 3GTest study [10] are generated from around 35 K active end-user probes consisting of 20 second long TCP transfers in the networks of four US cellular operators. The 4GTest study [8] uses a similar approach, but focuses on LTE users. The SpeedTest results show data from around 114 K active end-user probes in three geographical locations using a variety of cellular networks from 2G EDGE to 4G LTE, although a majority of the data comes from UMTS/EvDO/HSPA technologies.

Comparing the results it can be observed that for the 95th percentile the 3.5G throughput results from our measurements are considerably higher than previous non-LTE studies. The 95th percentile 4G throughput results also show a markedly higher throughput. Comparisons based on the 95th percentile

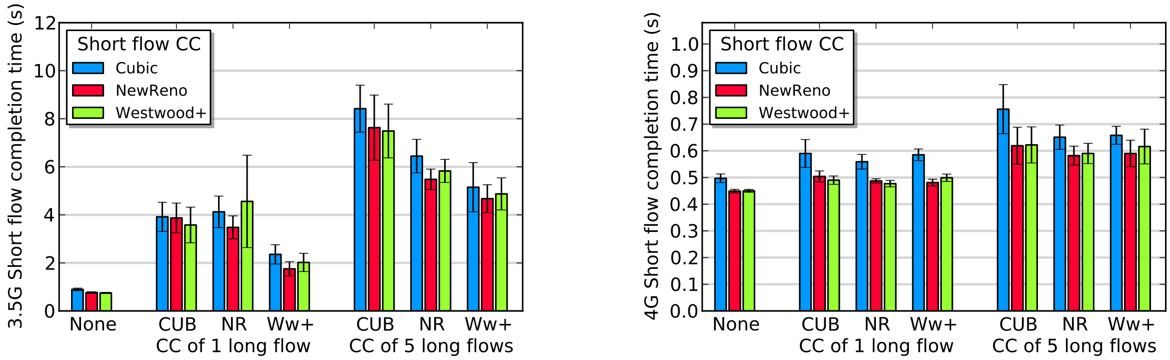

(a) 3.5G network

(b) 4G network

Figure 4: Mean response time of short flows with different background loads, with 95% confidence intervals

are expected to be less sensitive to differences in flow lengths between the measurements, since these values are likely to capture conditions close to best-case long flow characteristics, without much impact from the varying flow size distribution present in the different studies. With regards to the 3WHS-RTT, the results from our measurements show a much smaller spread between the 5th and 95th percentile than the comparable studies. One factor partly explaining this is that our results are based only on measurements containing one single flow, whereas this factor was not actively controlled in the other studies. As will be shown, adding a background flow increases our measured RTT variation, especially in 3.5G networks.

To provide a more detailed view of the measured results PDFs of throughput and RTT, created using a kernel of $s/\sqrt{n}$, are provided in Figures 2 and 3. Subfigure (a) shows the mean throughput for the A runs of long flows, whereas subfigure (b) shows the 3WHS-RTT for the runs A and B since the 3WHS-RTT metric is independent of flow lengths. The 3WHS-RTT can however be affected by an ongoing transfer of background data when new measurement flows are established. This case is shown in subfigure (c) which shows the short flow 3WHS-RTT observed for the C runs where one background flow is already running when a short flow is initiated. As can be seen by comparing Figures 2b and 2c, in the 3.5G case the 3WHS-RTT distribution is both shifted rightwards and widened when a background flow is present. This indicates that the background flow fills per-user cellular network buffers that are shared between flows, leading to increased delays for an initiating short flow. Comparing Figures 3b and 3c it is visible that the short flow 3WHS-RTT in the 4G network, however, is largely insensitive to the presence of background traffic.

Additional insights of the impact of competing background flows can be gained from our short flow completion time results provided in [1], some of which are shown in Figure 4. This figure presents the mean completion time of the short flows both with no background traffic, as well as with one or five background flows. This figure also shows data categorized by the different TCP congestion control schemes used in the experiments. Figure 4a clearly shows that there is a large difference between the mean short flow completion time when there is no background traffic versus when there is one or five background flow(s) in the 3.5G network. Also, it can be seen that the congestion control of the short flow does not have any significant impact. With regards to the background flow, using Westwood+ creates less impact on the short flow completion time. To exclude the impact of different congestion controls, only CUBIC and NewReno measurements were pooled to create the data shown in Figure 2c. For the 4G case shown in Figure 4b, it can be seen that the effect of adding a background flow has only a minor impact on the short flow completion time, which is consistent with the limited increase in 3WHS-RTT observed for Figure 3c. Further details and discussions of our measured results in relation to different congestion controls are provided in [1].

V. Measurement-based modeling and simulation

The measurements performed in the evaluated measurement campaign were done with only one flow length for the short flows. Covering a wide range of flow sizes with active measurements involves considerable expense and time, and as an alternative we use data collected in the measurement campaign as input for an analytical model that allows us to infer protocol efficiency.

A. *TCP protocol efficiency for single flow*

Previous work on protocol efficiency [13] primarily considered efficiency from the aggregate network perspective. In this work we view protocol efficiency from the vantage point of a user who would like to fully utilize the available transmission resources to minimize the response time as much as possible. It should be noted that we here are focusing on TCP protocol aspects for reducing response time, and while this is a major contributor to increased response times there are also many other factors such as DNS lookup delays, and delays associated with radio link state promotion.

To calculate protocol efficiency it is first necessary to model the TCP short flow completion time. There are several approaches for doing this, and we select a straight-forward approach that does not consider packet losses. Our 3.5G and

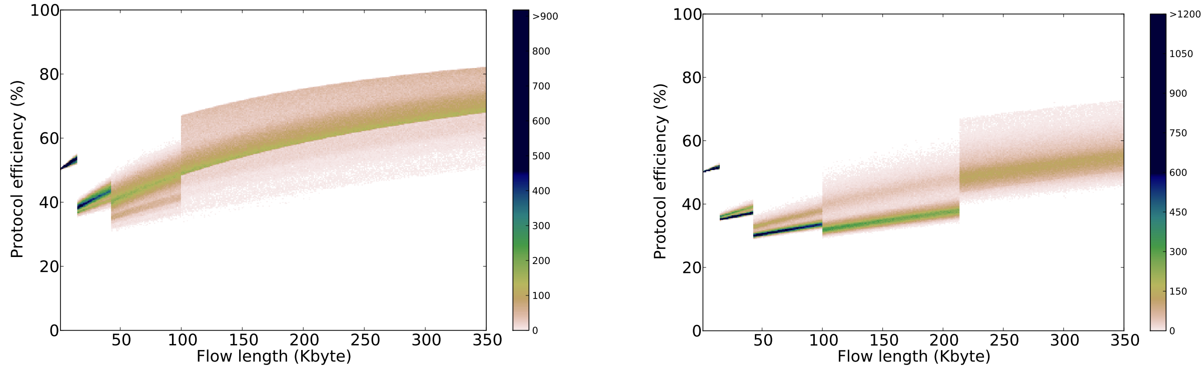

(a) 3.5G network (b) 4G network

Figure 5: Protocol efficiency as a function of flow size, no concurrent traffic

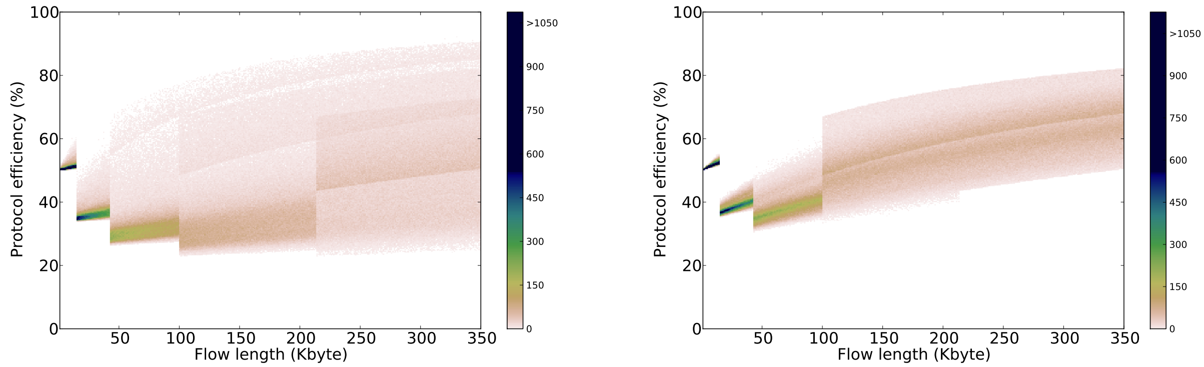

(a) 3.5G network (b) 4G network

Figure 6: Protocol efficiency as a function of flow size with one simultaneous background flow

4G measurement data show that none of the run B short flows experienced any packet loss.

Using this approach the estimated completion time is given as a function of flow size, maximum segment size, round-trip time, available capacity, and initial window size. Based on previous work [2], we formulate an extended analytic expression for the estimated short flow completion time $\hat{L_S}$ as:

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

where

$$P = min\left(\left\lfloor log_2\left(1 + \frac{R_S + S/C}{IS/C}\right)\right\rfloor, \left\lfloor log_2\left(1 + \frac{O}{IS}\right)\right\rfloor\right) \tag{2}$$

L	Flow completion time	(s)
R_S	RTT seen by short flow	(s)
O	Flow size	(bits)
C	Available Capacity	(bits/s)
S	Segment size	(bits)
I	Initial window	(segments)
P	Rounds with idling	-

The RTT value R_S used for single flows is the 3WHS-RTT when the user has no concurrent connections, i.e when buffers are empty. A novelty is the inclusion of a variable initial window size I, which is set to 10. The parameter 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 number of rounds with idle time is bounded either by the relationship between RTT, throughput and initial window size, or by the total amount of data in the short flow.

To calculate an ideal optimal protocol we use a one RTT connection establishment delay, and immediate full utilization of available transmission resources which yields an estimate of the optimal short flow completion time as $\hat{L_O} = R_S + O/C$. From this the the estimated TCP protocol efficiency is calculated as $\hat{E} = \hat{L_O}/\hat{L_S}$. We now couple these expressions to the empirically derived distributions presented in Figures 2 and 3 using the measured throughput as an estimate of the available capacity C. A Monte-Carlo approach is used to evaluate the relationship between the size of a short flow and the resulting protocol efficiency.

The results are shown in Figure 5, and it is readily observable that the protocol efficiency is worse for the 4G case than for the 3.5G case. It can be recalled from Figure 4 that in absolute terms the completion times are considerably better for

4G versus 3.5G, and these results indicate that this difference could be even larger if mitigations for protocol inefficiencies can be introduced. The step-like appearance in the left hand side of the graph comes when the size of a flow is increased so as to require a new transmission round which might incur additional idle time which in turn then decreases protocol efficiency. Somewhat surprisingly, the protocol efficiency remains low even for relatively large flows, indicating that the penalty of the extra RTT during connection establishment and the rounds with idle time takes a considerable time to amortize fully.

B. TCP protocol efficiency for multiple flows

Considering the case where there are multiple flows operating concurrently it is also possible to estimate an idealized ideal protocol behavior and the resulting protocol efficiency. In the presence of multiple flows, the optimal short flow completion time uses an assumption of ideal immediate sharing and is calculated as $\hat{L_O} = R_O + (O/C)/N$ where N is the number of concurrent flows and R_O is the unloaded RTT distribution shown in Figures 2b and 3b. Here it is necessary to note that the R_S experienced by an initiating short flow might be severely affected by a concurrent flow which has filled up shared buffers leading to a severely increased R_S, as discussed in section IV. Consequently, for this case of two multiple flows the R_S come from distributions shown in Figures 2c and 3c.

Considering the protocol efficiency for two flows as shown in Figure 6, it can be noted that the protocol efficiency shows a reduction for the 3.5G case, but an improvement in the 4G case. For the 3.5G case the decrease in optimal flow throughput is not enough to compensate the increase in RTT caused by bloated buffers being occupied by the long flow. The presence of buffer bloat is consistent with previous studies [11]. In the 4G case, the improved buffering approach in relation to 3.5G results in much lower increase of the RTT experienced by the 4G short flow. The decreased optimal rate hence leads to an overall improvement in protocol efficiency compared to the 4G single flow case.

VI. SUMMARY AND FUTURE WORK

This paper presents results from active measurements performed in the 3.5G and 4G mobile broadband networks of one Swedish operator. The study focuses on downlink throughput and RTT for both 3.5G and 4G, and relates these results to previous measurement studies. We also provide an evaluation using the TCP protocol efficiency metric. This metric shows how much TCP connection establishment overhead and slow start inefficiency contribute to increased completion times for short flows, which in turn may lead to wasted network as well as mobile battery resources. Using Monte-Carlo simulations based on the measured 3.5G and 4G characteristics, the obtained results indicate that while providing superior delays and throughput, 4G network characteristics actually result in worse TCP protocol efficiency than for 3.5G networks. The results also show that the improved queuing strategy observed in the examined 4G network lead to improved protocol efficiency when multiple flows are concurrently transferred as opposed to one single flow. For the 3.5G network our results show that the opposite is true, mainly as an effect of buffer bloat issues decreasing protocol efficiency for short flows in the presence of longer-running flows. While observant of the simplifications unavoidable in an analytical model, we believe these findings can be of value as one input when considering various network and protocol enhancement for potential further research or deployment.

We are currently collecting active measurements from all four major Swedish operators using a wider range of tests than reported here, and for future work we intend to comprehensively analyze this collected data from multiple perspectives.

ACKNOWLEDGMENT

Thanks to Giacomo del Giudice for assistance with data collection. This work was partly funded by the European Community under its Seventh Framework Programme through the Reducing Internet Transport Latency (RITE) project (ICT-317700) and 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] N. Cardwell, S. Savage, and T. Anderson. Modeling TCP latency. In *Proc. INFOCOM'00*, volume 3, pages 1742–1751. IEEE, 2000.

[3] X. Chen and J. Heidemann. Preferential treatment for short flows to reduce web latency. *Computer Networks*, 41(6):779–794, 2003.

[4] N. Dukkipati, T. Refice, Y. Cheng, J. Chu, T. Herbert, A. Agarwal, A. Jain, and N. Sutin. An Argument for Increasing TCP's Initial Congestion Window. *SIGCOMM CCR*, 40(3):26–33, June 2010.

[5] J. Eklund, K.-J. Grinnemo, and A. Brunstrom. Theoretical Analysis of an Ideal Startup Scheme in Multihomed SCTP. In *Proc. Networked Services and Applications (EUNICE'10)*, Berlin, 2010. Springer-Verlag.

[6] 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.

[7] J. Gettys and K. Nichols. Bufferbloat: Dark Buffers in the Internet. *Communications of the ACM*, 55(1):57–65, 2012.

[8] 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.

[9] 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.

[10] J. Huang, Q. Xu, B. Tiwana, Z. M. Mao, M. Zhang, and P. Bahl. Anatomizing Application Performance Differences on Smartphones. In *Proc. Mobile Systems, Applications, and Services (MobiSys'10)*, pages 165–178, New York, NY, USA, 2010. ACM.

[11] H. Jiang, Z. Liu, Y. Wang, K. Lee, and I. Rhee. Understanding bufferbloat in cellular networks. In *Proc. ACM SIGCOMM CellNet Workshop*. ACM, 2012.

[12] Y.-J. Lee and M. Atiquzzaman. Web object transfer latency over SCTP in the initial slow-start phase. *IEICE Electronics Express*, 2009.

[13] K. Liu and J. Lee. Impact of TCP Protocol Efficiency on Mobile Network Capacity Loss. In *Proc. Modeling and Optimization in Mobile, Ad Hoc, and Wireless Networks (WiOpt'13)*, pages 1–6, 2013.

[14] U. Paul, A. P. Subramanian, M. M. Buddhikot, and S. R. Das. Understanding traffic dynamics in cellular data networks. In *Proc. INFOCOM'11*, pages 882–890. IEEE, 2011.

[15] S. Radhakrishnan, Y. Cheng, J. Chu, A. Jain, and B. Raghavan. TCP Fast Open. In *Proc. Conf. Networking EXperiments and Technologies (CONEXT'11)*. ACM, 2011.

[16] J. Sommers and P. Barford. Cell vs. WiFi: On the performance of metro area mobile connections. In *Proc. Internet Measurement Conference (IMC'12)*, pages 301–314, New York, NY, USA, 2012. ACM.