Examining TCP Short Flow Performance in Cellular Networks Through Active and Passive Measurements

Johan Garcia
johan.garcia@kau.se

Stefan Alfredsson
stefan.alfredsson@kau.se

Anna Brunstrom
anna.brunstrom@kau.se

Department of Mathematics and Computer Science
Karlstad University, Karlstad, Sweden

ABSTRACT

In this study we examine the conditions in a current cellular network by examining data passively collected in the core of a cellular operator during a 24-hour period. More than 2 billion traffic measurement data points from over 500,000 cellular users are analyzed. The analysis characterizes the Time-of-Day (ToD) variations for traffic intensity and session length and serves as a complement to the active measurements also performed. A comprehensive active measurement campaign was completed in the HSDPA+ and LTE networks of the four major Swedish operators. We collect around 50,000 data points from stationary cellular modems and analyze the ToD variation pattern for underlying network layer metrics such as delay and throughput. In conjunction with the time-varying session size distribution obtained from the passive measurements, we then analyze the ToD impact on TCP flows of varying sizes. The ToD effects are examined using time series analysis with Lomb-Scargle periodograms and differential Bayesian Information Criterion to allow comparison of the relative impact of the network ToD effects. The results show that ToD effects are predominantly impacting longer-running flows, and although short flows are also impacted they are mostly constrained by other issues such as protocol efficiency.

CCS Concepts

•**Networks** → **Transport protocols;** *Network measurement; Mobile networks;*

1. INTRODUCTION

Cellular networks are experiencing a strong and continued growth in the amount of data traffic carried. This is fueled both by increased use of devices such as smart phones and tablets, and also by the use of stationary cellular modems used to provide broadband to the home. As the mix of devices using the network evolves, this also has an impact on the traffic patterns present in the network. Understanding current traffic characteristics are important when evaluating and tuning performance at various layers in the protocol stack.

Regarding performance, user perception is for most applications tied to the latency characteristics of the upper transport and application layers. In this work, we study performance aspects with a particular interest in the time of day (ToD) variations that occur as a consequence of the varying flow length distribution and network load over the hours of a 24-hour period. Our analysis of flow length distribution and network load is based on passive traffic measurements from over 500,000 cellular users collected at the core of a cellular operator, illustrating the prevalence of short flows and diurnal variations in load.

The passive measurement results are compared to results of long-running active measurements from the network edge. The active measurements capture network ToD effects in terms of bandwidth and round-trip time (RTT) behavior as well as ToD effects on flow completion time for flows of varying sizes. Our network edge results show that while the diurnal variations in bandwidth are significant, and in concert with the diurnal variations in load seen in the passive measurements, the diurnal variations in RTT are much less significant. This is important as the performance of short TCP flows is to a large extent dominated by the delay characteristics, rather than the available bandwidth in the network. This is confirmed by carrying out a time series analysis, using Lomb-Scargle periodograms and differential Bayesian Information Criterion, comparing the relative impact of the network ToD effects on the flow completion time for flows of varying sizes.

Taken together, our results thus suggest that a major fraction of all flows in modern cellular networks operate under conditions where they are restricted mainly by protocol design and network delay rather than by bandwidth. The performance of these flows is also less influenced by diurnal effects than one may expect.

In the next section we provide a background to the work, followed in Section 3 by the results from the passive measurements in the core. Section 4 presents results from the active measurements, followed by conclusions in Section 5.

2. BACKGROUND

A number of previous measurement studies has examined various aspects of cellular network performance. Elmokashfi et al. [2] 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. Considering research on ToD variations over longer

AllThingsCellular'15, August 17-21, 2015, London, United Kingdom

DOI: http://dx.doi.org/10.1145/2785971.2785974

time periods, Ricciato et al. [11] analyzed traffic in the core of a large 3G network, and found typical ToD patterns in delay and throughput where user activity peak during the late evening, and dip during the night. Further studies of delay and ToD effects in 3G are done by Paul et al. [10], showing similar trends. Foddis et al. [3] analyzed traffic from an LTE eNodeB on a call, frame and application level. The observed ToD and day-of-week effects were significant, but are explained by monitoring in a local business area, where LTE devices (at the time) may primarily have been used by employees rather than consumers.

Zhang et al. [14] examined HTTP traffic collected in a large cellular network, and analyzed traces on packet, flow and session levels. With regards to ToD effects, they find diurnal patterns, which also show marked differences between cellular and wireline access with regards to specific applications. For example, facebook was more used over wireline at daytime, while more used via cellular at evening and night time. Further traffic characterization by Meng et al. [9] shows that a considerable fraction of traffic to/from smartphones is generated by background traffic, i.e. traffic generated when the user is not actively using his or her phone. They identify diurnal patterns with respect to the total traffic volume as well as the fraction of background traffic, with larger fractions of background traffic being observed during off-peak hours.

Other measurements in cellular networks include work by Huang et al. [6, 7] examining the performance of several cellular technologies. Their measurements done in the core of an LTE network showed that 90% of all TCP flows were smaller than 35.9 KiB [7]. Furthermore, they also showed that many long-lived TCP transfers were unable to fully utilize the available capacity due to protocol inefficiencies. In [4], Garcia et al. examine the TCP protocol efficiency of short flows in HSDPA+ and LTE networks. Protocol efficiency contrasts the actual TCP flow completion times to an idealized case in which the transmission resources provided by the cellular network can be instantaneously and fully utilized. Zhang et al. [15] analyze flow rates and compare cellular to wireline networks. Cellular wireless networks are shown to exhibit higher variability, with the access link as the typical bottleneck.

3. PASSIVE CORE NETWORK RESULTS

The ToD characteristics of current cellular traffic were analyzed using measurements obtained from the core network of a European operator. Data was collected every 100 ms for every active user in the network using hardware also utilized for subscriber management duties. A user was considered active if there was at least one active TCP session established for that user.

3.1 Base Characteristics Time of Day Effects

Figure 1 shows the number of active users and amount of downlink and uplink data. Data is summarized in five-minute time slots. An apparent diurnal pattern is visible for both the traffic amounts and the active users, which matches with the previous work discussed in Section 2. The variation appears less pronounced for the number of active users, and can be quantified by the Coefficient of Variation (CV) where $CV = \sigma/\mu$. This gives $CV_{DL} = 0.53$, $CV_{UL} = 0.56$, and $CV_{Users} = 0.19$. The CV is thus almost three times larger for the amount of traffic than for the number of active users.

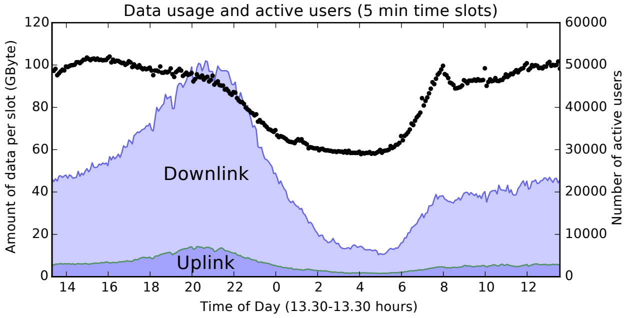

Figure 1: Data transferred and active users over ToD

3.2 User Sessions

From the recorded data it is possible to also get a characterization of user traffic patterns. To aid in such analysis it is useful to consider the observed traffic as data exchange sessions. As explicit transport layer flow data is not captured, sessions are instead used to refer to data transfers that are temporally related. One, or several parallel, TCP flows would thus be considered to be one session, under the assumption that the flows are not application-limited, i.e. have long idle periods when no data is transferred. Figure 2 illustrates the approach for splitting data transfers into sessions. The gap size parameter T_{gap} influences the session size distribution, and Figure 3 shows the session size ECDF for varying gap sizes over the complete data set. As can be seen a majority of the sessions are very small, irrespective of session gap size. We term these (<1000 bytes) as XXS sessions. A large fraction of these XXS sessions are likely background status updates for various apps that are active on the users phone. Such status updates are normally exchanged in the background, and not in response to user actions. The performance of XXS background traffic is expected to have minor influence on user perceived performance. In the following analysis a gap size parameter of 1 sec is used, similar to the threshold used to aggregate connections in [1].

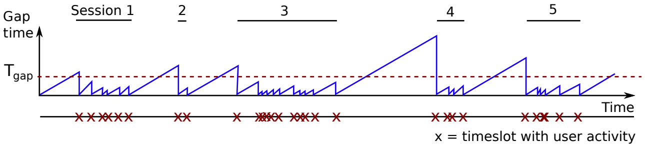

Figure 2: Splitting data transfers into sessions

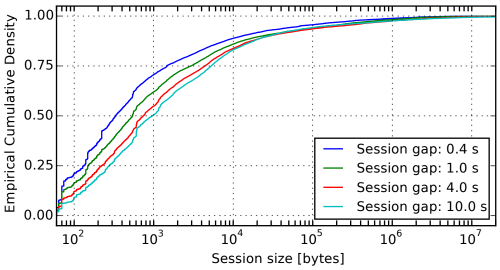

Figure 3: Session size ECDF per session gap size

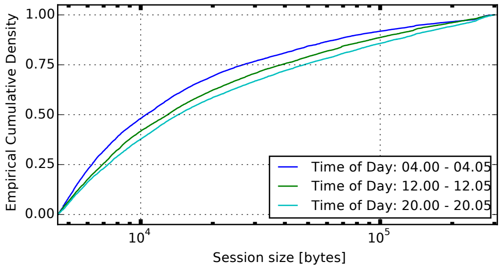

Figure 4: ECDF for 4.5 - 300KiByte sessions per ToD

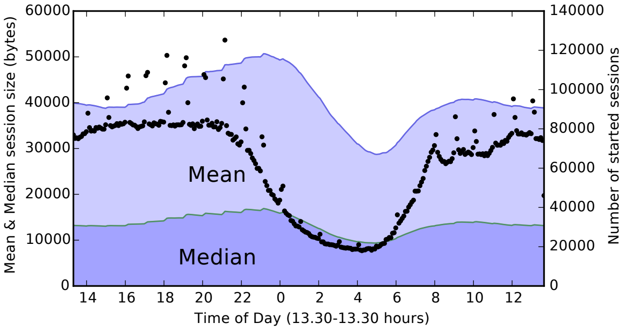

Figure 5: Session size for 4.5 - 300KiByte sessions over ToD

3.3 Session Size Time of Day Effects

User-initiated data exchange sessions are often larger than 1000 bytes, and in the following we focus on session sizes ranging from 4.5 to 300 KiBytes. Figure 4 shows the ECDF for sessions in that size range, for three different five minutes time slots. It can be observed that the session size varies in a way that is consistent with the data transfer variations observed in Figure 1.

A 24h picture of the ToD variation of the session length is provided in Figure 5, which shows Holt-Winters exponentially weighted moving averages of the mean and median session length. The figure also shows the number of sessions started. As can be seen the mean varies more than the median, reflecting that the size distribution of the session is becoming more skewed towards larger sessions sizes. The variation in the median observable in Figure 4 can also be viewed here, over the complete time range. The session sizes for an unrestricted range show a similar diurnal pattern. Considering again the data transfer ToD variations observed in Figure 1, it can now be concluded that the observed variations are due to a combination of factors. During nighttime there are fewer active users, the active users have fewer sessions per time unit, and the session sizes are smaller. Having established these passive measurement ToD characteristics, the focus in the next section is on active measurements.

4. ACTIVE ACCESS NETWORK RESULTS

The networks of the four major cellular operators in Sweden were used for an active measurement campaign during winter 2014-15. Several different types of measurement run types were used, with Table 1 describing the two used in this analysis along with the number of measurement rounds analyzed per operator.

Identifier	Description of run type	Nr of rounds
'long'	TCP throughput measurements for 20 MiB transfer	3.5G: 716, 715, 711, 593 4G: 713, 718, 710, 698
'short_flow'	Flow completion times for sizes: 1388 - 327568 bytes, (11 steps)	3.5G: 714, 713, 715, 606 4G: 713, 701, 709, 701

Table 1: Run types used in evaluation

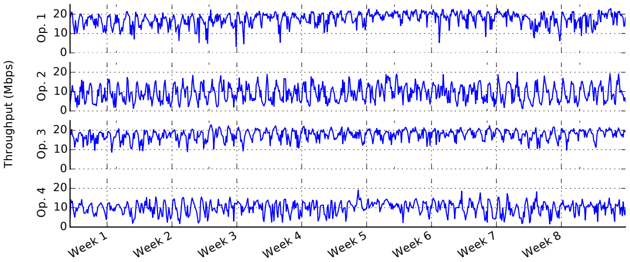

Figure 6: 3.5G Throughput data (concatenated)

The measurements were performed from a fixed location in Karlstad. Measurement data was collected on a server connected to the Swedish University backbone (SUNET), and a client with four Huawei E392 USB modems with simcards for the specific operators. Both computers were running Linux with kernel version 3.16.3. Packet traces were collected using tcpdump and processed with tcptrace. Individual TCP flows which had SYN or FIN retransmissions were not included in the data set. TCP segment offloading (TSO), TCP metric caching and other related functionality were turned off or flushed to ensure that the various optimizations that are performed in the OS networking code and in device drivers would 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.

4.1 Throughput and RTT ToD Variation

Measurements were collected at two hour intervals over a period of several months, with some interruptions. For the throughput measurements, data from the long runs were used. Figure 6 shows an overview of the measured TCP throughput over the concatenated measurement days. Some periodicity appears to be visible. The diurnal ToD effects can also be displayed using bins corresponding to the 12 measurement time slots used. Figure 7 shows the mean throughput per time slot, with bootstrapped 95% confidence intervals. The non-overlapping confidence intervals suggest that ToD variation is present in the long flow throughput of all four operators for 3.5G. For 4G the picture is not as clear.

The round trip time (RTT) measurements were in this study based on the three-way handshake (3WHS), as this has been found to be an appropriate approach [5]. The 3WHS values are extracted from the short_flow measurements. Figure 8 shows the binning of the measured RTTs. While some ToD effect can be discerned with binning, the effects appear less apparent than for the throughput. The 3WHS RTT represents the RTT observed by packets when the user has no other concurrent traffic. It is also possible to examine the RTT experienced by packets in long flows.

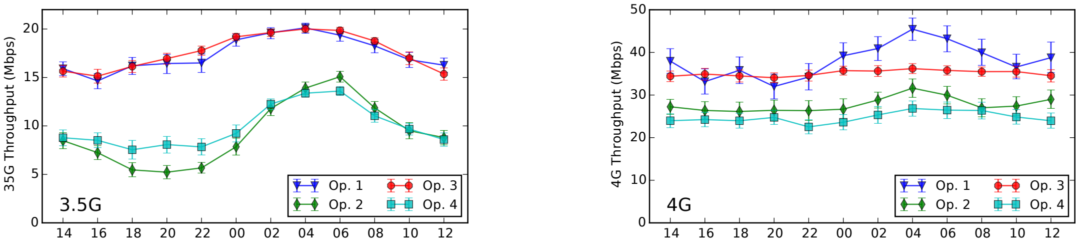

Figure 7: Throughput over Time-of-Day per operator

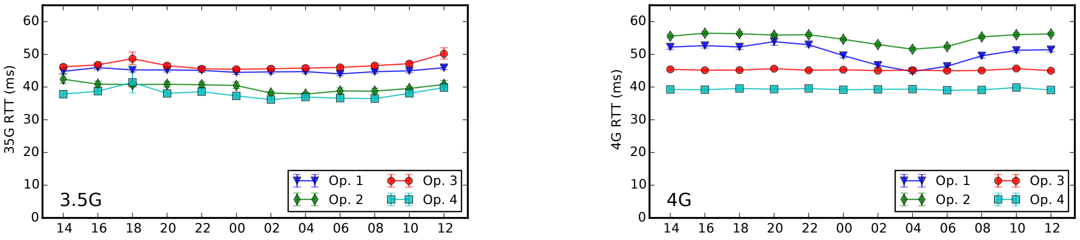

Figure 8: 3WHS RTT over Time-of-Day per operator

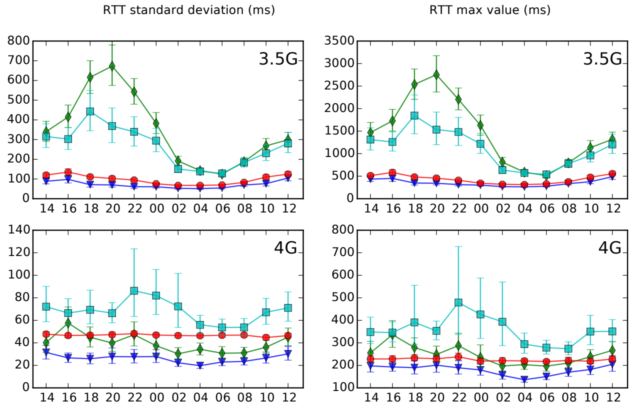

Figure 9: In-flow RTT variation for long flows over ToD

In long flows the packets may become queued, leading to a corresponding increase in the experienced RTT. Some RTT characteristics for the long flows are shown in Figure 9. As can be seen these in-flow RTT metrics exhibit a clear ToD variation for 3.5G, in a way that is consistent with the variation of throughput observed in Figure 7.

These graphs provide a complementary view from an edge device on the ToD variations. It is clear that the aggregate load variations observed in the passive measurements are also to some extent reflected in the characteristics seen by the edge device. In the following subsections we elaborate on a quantification of how large these effects can be.

4.2 Modeling of ToD Impact on Short Flows

As shown by the previously discussed session statistics, the large flows are only a very small fraction of all flows. A fraction of these very large flows are also tied to user activities that are not interactive in nature, and for such flows the performance of the flow is less crucial from a user experience point of view. In this subsection, the main focus is on examining to what extent the more frequent short (but not XXS) flows are impacted by the traffic load ToD variation.

The completion times of short flows are to a large extent dependent on the RTT, so any ToD variations in the RTT can be expected to also impact short flow completion times. Short flows are in the majority of cases not subjected to any packet loss. In this data set less than 3% of the 60989 short flows experienced packet loss. For the zero packet loss case, a simplified model similar to the one used in [13] compute the expected short flow completion time $\hat{L_S}$ as:

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

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 cap.	(bits/s)
R_S	RTT	(s)	S	Segment size	(bits)
O	Flow size	(bits)	I	Initial window	(segments)
P	Rounds w. idling	-			

The RTT value R_S used is the base network delay when the user has no concurrent connection, i.e. when the per-user buffers are empty. The parameter I represents the TCP initial window, and is set to 10 here. The variable P is the number of transmission rounds with idle time, i.e. when the received rate is constrained by the congestion window of the sender and not the available transmission resources. Using this expression it is possible to graph the relative impact of changes in RTT and throughput as a function of flow size. This is shown in Figure 10, with the mean values from the measurements used as baselines, which gives $R_S^{3.5G} = 42.6$ms, $R_S^{4G} = 47.5$ms and $C^{3.5G} = 13.7$Mbps, $C^{4G} = 31.5$Mbps. The figure shows the relative impact of a 10% change in R_S and C, i.e. what is the negative impact on $\hat{L_S}$ of a 10% decrease of C and a 10% increase in R_S. As can be seen in the figure, the completion time is much more sensitive to variation in R_S than in C, although to a lesser degree for 3.5G than for 4G. As shown in the previous subsection, the amount of ToD variation in R_S is typically lower

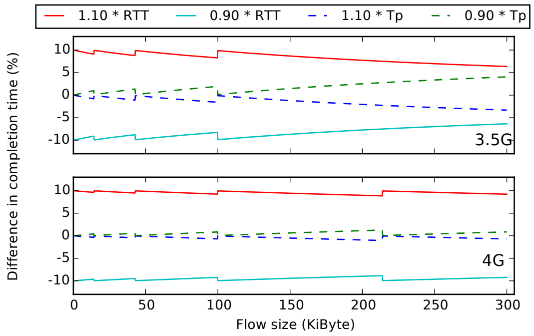

Figure 10: Relative impact of Throughput and RTT on short flow completion time

than the amount of variation in C. The net effect of these ToD variations on short flow completion times is examined and quantified in the next subsection.

4.3 Time Series Analysis

While the binning approach used in the previous section can be used to illustrate ToD effects, it has some drawbacks. The interpretation is complicated as it is a case of multiple hypothesis testing, and binning does not provide a ready quantification of the size of the observed ToD effects. Fortunately there are more powerful tools available for characterization of time series. A common method is to use the auto-correlation function (ACF) or partial-ACF (PACF) to describe the extent of cycles in the observed data. However, some data points in the time series are missing due to modem reboots etc., and the measurements are performed in sequence during each run so the exact timing of each individual measurement can vary somewhat. Consequently, the time series sampling is non-uniform which makes an ACF approach infeasible.

To allow time series analysis under these constraints, we instead use Lomb-Scargle periodograms [8, 12] as they allow non-uniform sampling of the timed events. A normalized L-S periodogram $P(\omega)$ with homoscedastic errors is computed as

$$P(\omega) = \frac{1}{\sum_j (y_j - \overline{y})^2} \left[\frac{\left[\sum_j (y_j - \overline{y}) \cos\left[\omega(t_j - \tau)\right]\right]^2}{\sum_j \cos^2\left[\omega(t_j - \tau)\right]} + \frac{\left[\sum_j (y_j - \overline{y}) \sin\left[\omega(t_j - \tau)\right]\right]^2}{\sum_j \sin^2\left[\omega(t_j - \tau)\right]} \right]$$

where $j = 1...N$, N the number of samples, y_j the jth sample and ω the angular frequency. To allow invariance to translation along the t-axis, i.e. non-uniform sampling, the τ offset is defined as

$$\tan(2\omega\tau) = \frac{\sum_j \sin(2\omega t_j)}{\sum_j \cos(2\omega t_j)}$$

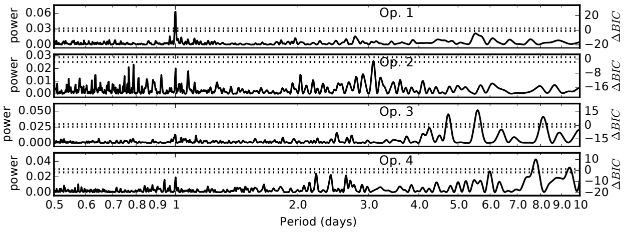

Figure 11: Periodogram for 4G Throughput

As a comparison to the binning approach, the periodogram for the throughput in the 4G network is shown in Figure 11. The left y-axis shows normalized power (0-1), and the right y-axis shows the difference in the Bayesian Information Criteria (ΔBIC). The ΔBIC is defined as $\Delta\text{BIC} = \chi^2_o - \chi^2(\omega_0) - (k_o - k_\omega) ln(N)$, where χ^2_o is for a pure noise model, $\chi^2(\omega_0)$ is for a model assuming one periodicity at $\omega_0 = 2\pi/P$ with $P = 1$ day, $k_o = 1$, and $k_\omega = 4$. The two (often overlapping) horizontal lines in the graph correspond to the 5% and 1% significance levels, established by nonparametric bootstrap resampling. As shown in Figure 11, significant periodicity at period $P = 1$ day is found for operator 1. For operators 3 and 4 there is spectral power in the vicinity of the 7 day period. Possibly this could be related to a weekday/weekend variation in the data. To explore this aspect, a more complex model with multiple periodicities would have to be employed, which is left for future work.

To quantify the presence of periodicity in the short flow completion times, a total of 88 periodograms are computed for all combinations of operators, network technologies, and flow sizes. Figure 12 shows the periodograms for one of the flow sizes, 88832 bytes. As is visible in the figure, 3.5G operators 2 and 4 had significant ToD effects for this flow size, whereas for 4G operators 1 and 2 had significant effects. The relative magnitude of these effects can be compared based on the ΔBIC values, and it can be seen that the largest effects occurred in the 3.5G network. To provide an overview of the effects for all flow sizes, Figure 13 shows the ΔBIC for all $P = 1$ day peaks which were significant at the 1% level and had positive ΔBIC, along with similar ΔBIC of the 20Mbyte long flow throughput and the short flow 3WHS

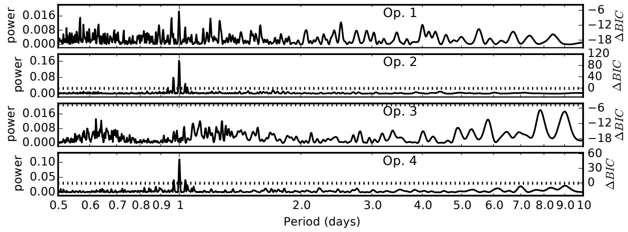

(a) 3.5G network, 88832 Byte short flow

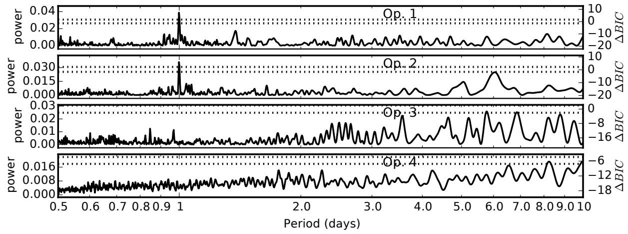

(b) 4G network, 88832 Byte short flow

Figure 12: Periodogram for short flow completion time

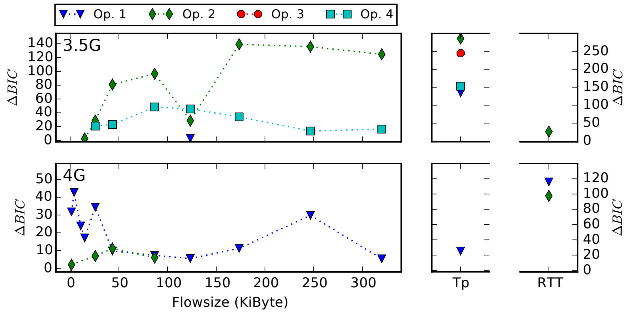

Figure 13: Differential BIC for 1 day periodicities

RTT. As is visible in the figure, there are two operators in 3.5G and two in 4G for which significant ToD effects were detected for some flow lengths. Looking at the ToD effects of throughput and RTT in relation to the effects on short flow completion times it can be observed that there is a difference between 3.5G and 4G. For 3.5G all operators which had short flow effects also had large throughput effects, and only minor RTT effects were observed. For the 4G case, the operators with flow size effects instead have RTT effects, and only a minor throughput effect is seen. These results suggest that while some short flow completion time ToD effects can be observed in both 3.5G and 4G networks, the underlying mechanisms causing the effects differ between the networks. While in total the short flow ToD variation is limited, the long flow throughput ToD variation is significantly larger.

Regarding generalizability it can be noted that these active measurements were performed at a single location, and thus the individual measured base values reflect the conditions at this location. Nevertheless, higher level observations such as the discussed difference in underlying mechanisms for 3.5G and 4G short flow ToD effects is expected to have a greater degree of generalizability.

5. CONCLUSIONS

Using large scale passive and active measurements we have examined the Time-of-Day effects present in cellular networks. In line with intuition, the passive measurements showed that the number of active users and amount of transferred traffic varied over the hours of the day. Our analysis also considered the extent of session size ToD variation, showing that the session sizes also varied significantly over the day. Using long-running active measurements, the amount of ToD variation of the base network characteristics throughput and RTT was examined. Furthermore, a more detailed time-series analysis was performed on a subset of the flow sizes. Using periodograms and differential Bayesian Information Criterion, the relative strength of ToD effects was explored. The results showed that for the 3.5G measurements the ToD effects were very strong for the long flow throughput of all four operators, but only two of the operators had significant ToD effects for short flows. For the 4G measurements, a smaller throughput variation was present only for one operator, while two operators still had short flow ToD effects. An important aspect of this work is the use of a robust approach to quantifying the amount of ToD effects, a metric which could be of considerable interest as an input to network operations and dimensioning.

Acknowledgment

The authors wish to thank Procera Networks for providing the passive data. This work was partly funded by the KK-foundation (HITS research profile), and partly funded by .SE Internetfonden.

6. REFERENCES

[1] M. Allman. Comments on Bufferbloat. *ACM Computer Comm. Review*, 43(1), Jan. 2013.

[2] A. Elmokashfi, A. Kvalbein, J. Xiang, and K. R. Evensen. Characterizing delays in Norwegian 3G networks. In *Proc. PAM*. Springer, 2012.

[3] G. Foddis, R. Garroppo, S. Giordano, G. Procissi, S. Roma, and S. Topazzi. LTE traffic analysis and application behavior characterization. In *Proc. EuCNC*, June 2014.

[4] J. Garcia, S. Alfredsson, and A. Brunstrom. A Measurement Based Study of TCP Protocol Efficiency in Cellular Networks. In *Proc. WiOpt*, 2014.

[5] J. Garcia, S. Alfredsson, and A. Brunstrom. Delay metrics and delay characteristics: A study of four Swedish HSDPA+ and LTE networks. In *Proc. EuCNC*, June 2015.

[6] 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. MobiSys*. ACM, 2012.

[7] 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.

[8] N. Lomb. Least-squares frequency analysis of unequally spaced data. *Astrophysics and Space Science*, 39(2):447–462, 1976.

[9] L. Meng, S. Liu, and A. Striegel. Characterizing the utility of smartphone background traffic. In *Proc. ICCCN*, 2014.

[10] 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.

[11] F. Ricciato, E. Hasenleithner, and P. Romirer-Maierhofer. Traffic analysis at short time-scales: An empirical case study from a 3G cellular network. *Trans. Network and Service Management*, 5(1), March 2008.

[12] J. D. Scargle. Studies in astronomical time series analysis. II - Statistical aspects of spectral analysis of unevenly spaced data. *Astrophysical Journal*, 263:835–853, Dec. 1982.

[13] M. Scharf. Comparison of end-to-end and network-supported fast startup congestion control schemes. *Computer Networks*, 55(8):1921–1940, June 2011.

[14] Y. Zhang and A. Arvidsson. Understanding the characteristics of cellular data traffic. In *Proc. ACM SIGCOMM CellNet*. ACM, 2012.

[15] Y. Zhang, A. Arvidsson, M. Siekkinen, and G. Urvoy-Keller. Understanding HTTP flow rates in cellular networks. In *Proc. IFIP Networking Conference*, 2014.