OpenSignal

OpenSignal comments on Draft BEREC WORK PROGRAM 2015

OpenSignal is creating a comprehensive database of cell phone towers, cell phone signal strength readings, and Wi-Fi access points around the world. We create this database to provide insight on connectivity, adopting the philosophy that only a data driven examination of the true performance of these networks can lead to active improvements of electronic communications infrastructure. We achieve this by collecting data on mobile devices through our Android and iOS apps, which have had over 11 million downloads around the world.

We are providing comments on the Draft BEREC Work Program 2015 because the three strategic pillars of the strategy are directly aligned with OpenSignal's mission and methodology of crowdsourcing data to allow accurate analysis of the performance of networks. These comments focus on Strategic Priority 3: Empowering End Users.

About OpenSignal
The OpenSignal data is collected from real world consumer smartphones, and is recorded under conditions of normal usage. Rather than approximate the user experience, we directly measure it from the users of our smartphone application. Our application can be freely downloaded on either iOS or Android devices and constantly monitors the true network experience that users are getting on those devices. Through being located on consumer smartphones, we are able to observe the network exactly as the end user experiences it. This customer-centric approach allows us to measure the true end-to-end experience of the mobile network. We're not interested in models, simulations or assumptions – our goal is to directly measure user experience through the eyes of the users themselves.

Although operators have been monitoring how their networks perform since the very beginning, there remains a disconnect between the standard network KPIs and what surveys say about customer experience. We believe that the only way to bridge this gap is to measure the network using the customer experience as our starting point.

The OpenSignal website provides consumer-focused visualizations and analysis, based exclusively on data collected from the users of the OpenSignal application. Our website's features include coverage maps that show cellular signal strength for a given geographical area, ranking all of the networks by performance in that location, and the locations of all cell towers within that particular region.

Comments on 3.3.1 a, Feasibility study of QoS Monitoring in the context of NN

OpenSignal strongly supports the recommendation of developing best practice on a QoS Monitoring system. In assessing methodologies for monitoring Internet access and signal, Opensignal recommends consideration of crowdsourcing from users' devices.

Alternative methodologies to measuring actual user experience through crowdsourcing data, such as drive or lab testing, do not emulate real world mobile device usage. Traditional drive testing does not even attempt to measure the indoor user experience, which is where more than 50% of mobile usage occurs, according to numerous studies. Drive tests that do measure indoor conditions still lack the behavioural aspect of smartphone use, a very important dimension. Whether it's the locations users tend to spend their time or the way in which a user holds their device, simulated approaches are completely unable to model the impact of user behaviour. The on-device testing that OpenSignal employs ensures that user behavior is built into the KPIs we measure from the very beginning.

In addition, tapping into the potential of consumers allows us to achieve a scale of data analysis that is not possible by drive testing. The BEREC Work Plan discusses how it is important to evaluate "*the cost of cooperation, the complexity of the system, the legal requirements and the time constraints related to alignment among NRAs."* Through our crowdsourcing approach, we have achieved over **11 million downloads** of the OpenSignal application worldwide, with an active population of over **1.8 million devices** recording network performance metrics at any given time. We record over **150 million data points** on network performance from OpenSignal users daily. This data set is ready and available today, and growing rapidly all the time, removing time constraints of other data collection methods. Drive testing as a process is expensive and there is a limit to the number of distinct locations in which it is feasible to test, as well as the frequency with which they can be updated. While drive testing remains applicable for some use cases, we believe on-device data represents the future of customer experience monitoring.

The 2015 Work Program mentions *"cross-network (and possibly cross-border) measurements while building on existing experience"* will be taken into consideration. Crowdsourcing data is agnostic of network and borders, as we have found with our OpenSignal app, which collects data from cell phones in **over 200 countries**. This data is consistently reported, independent of the country or network, allowing for accurate benchmarking. The datasets that we have been able to build using this technique are not just for academic value, but have translated into tangible economic gains, such as with network operators who have been able to better understand their own performance at the end-user level.

For example, CSL, an operator in Hong Kong, wanted to

- Gain a deeper insight into the experience on their network from a customer-centric point of view.
- Gain an understanding of metrics they were not currently able to track such as WiFi usage behavior or the proportion of time users have no network coverage.
- Benchmark the performance of their network against their competitors whilst reducing the large capital expenditure on drive testing.

Using the data that we have gathered by OpenSignal in Hong Kong, our analysis allowed for the identification of an issue that was causing a large proportion of their users to experience poor latencies, enabling the operator to roll out a fix to their core network within 10 days. This particular problem had only been affecting users on lower end devices, which had not been included in the device subset used in the drive and lab testing, and so the existing testing methods had completely missed the problem.

The data also provided insight into certain metrics that were previously opaque such as:

- Seeing the proportion of time the average user was connected to LTE, in order to assess the success of their current LTE network rollout.
- Understanding how their unlimited data plans was impacting WiFi usage behavior.
- Seeing the proportion of their customers and their competitor's customers that were experiencing throttled data speeds.

OpenSignal has real world experience of using crowd-sourced methodology to analyze network performance to enable infrastructure improves and better overall service for the end-use. As such, OpenSignal has been invited by the International Telecommunications Union (ITU) to present on crowd-source methodology at the Quality of Service Development Group in Dubai in November 2014.

OpenSignal supports the deliverable of a BEREC Internal Report on the feasibility of QoS monitoring in the context of NN, and recommends the inclusion of a crowd-sourced approach in this assessment of monitoring options, based on its strength in reporting end-user experience, proven track record in providing valuable insights to network operators, cost-effectiveness and no required lead-time for data collection.

For more information, contact Ellie Ereira at ellie@opensignal.com or +447929854557

As an annex to these comments, we attach an OpenSignal research paper presented at the 10th IEEE International Workshop on Performance and Management of Wireless and Mobile Networks on Modelling Download Throughput of LTE Networks.

Modelling Download Throughput of LTE Networks

Joe Cainey, Brendan Gill, Samuel Johnston, James Robinson, Sam Westwood
OpenSignal
White Bear Yard
London, EC1R 5DF
Email: {firstname}@opensignal.com

***Abstract*—We report the initial findings of an investigation that attempts to model download throughput of LTE networks in the US using the results of 130,296 speedtests performed by users of the OpenSignal mobile application on Android devices. We have developed a simple model which regresses LTE download throughput on a number of independent variables. We find that signal quality and strength, temporal and network operator factors are all essential, and also find an interesting device-specific dependence. We show that none of the LTE-specific signal metrics is sufficient alone to determine download throughput, although once combined they collectively explain an appreciable fraction of the variability in the model. We obtain time-of-day results which we believe act as a proxy for network congestion, and are able to quantify the relative performance of the four largest US cellular operators.**

I. Introduction

With the advent of fourth-generation mobile telecommunications technology in the United States, affordable high-speed data connections are now widely available. Despite the high theoretical throughput of 4G standards, there is large variation in the download speed actually achieved by end users.

There are currently two competing 4G standards in use in the US. The earliest to launch, WiMAX, has largely been supplanted by Long-Term Evolution (LTE) technology. Of the biggest four network operators in the US, only Sprint currently operates a WiMAX service, and this will shortly be switched off. Therefore in this research we focus on LTE in particular, as the most relevant technology, in terms of current and future usage.

Understanding the factors that determine download throughput is valuable to both users and network operators. Users may wish to purchase devices and services that obtain the best available performance in practice, and network operators need to understand how best to extend and enhance their infrastructure. Basic studies that look at empirical estimates of average download speeds help to a certain extent, but there are significant confounding factors such as location that may mask important features.

The throughput of a mobile device can be affected by any section of the link between the data source and the end user. With sufficiently advanced smartphones, we are now able to capture data surrounding certain aspects of the link between the user and cell, allowing for explanatory models to be developed. This work represents the first steps towards such a model.

Some previous work has focussed on the development of spatial maps which characterise features of mobile network connections by location (e.g. [5], [16]). This requires vast amounts of data to perform effectively once we begin faceting by device model or by time. In this research we avoid the issue of mapping, and attempt to uncover location independent results by considering measurements of signal quality in our model. We use this to control for the fact that urban environments will typically have wider coverage and better infrastructure than rural locations, allowing us to uncover device, time and operator dependence.

We have developed a mobile application, OpenSignal [4] which collects crowd-sourced signal information from over 6 million devices world-wide, producing over 30 billion readings. This data reflects real-world usage of mobile devices over multiple years, allowing for deep insight into the nature of mobile networks and devices. We collect a wide range of values which allow us to analyse the quality of the link between the user and the cell-tower, and measures of the performance of the device. Our data is characterised by its size and richness of device and geographic variation, and represents one of the largest such datasets collected to date.

A. Existing Work

There is a growing body of literature dedicated to directly analysing download throughput data for mobile devices. Although much of what exists is restricted to data collected from small samples and geographical regions, we are beginning to observe the emergence of large and detailed datasets comparable to our own. The one unifying result, observed in almost all studies, is in the vast amount of variability observed in mobile network performance.

Many studies have looked at 802.11 WiFi networks (e.g. [11], [6]), whereas we are concerned with the performance of mobile cellular networks in particular. WiFi networks have a vastly different profile due to link-layer differences [20], and so much of the work in this area is not applicable to cellular networks.

Some studies have looked at older 3G technology in detail [10], performing deep packet-level analysis of network connections. As we discuss, LTE and 3G networks have a fundamentally different structure, and have different methods of measuring signal strength and quality which means that different approaches must be developed.

978-1-4799-3784-4/14/$31.00

The work in [8] and [7] contains an in-depth study of LTE networks, using a large data set collected by a network operator. In this work they discover that TCP connections routinely under-utilise the capacity of LTE connections, highlighting the importance of understanding and exploring the unique features of LTE for performance analysis. We attempt to look at issues from the user's perspective, rather than using data derived from the network operator.

We use a general approach of data collection that is similar to [12] and [13], which advocate using a distributed sensor network for network analysis. We do not use the same centrally-managed and minimalistic collection regime though, preferring to crowd-source the measurements, and collect as much data as is reasonable from users. [20] also conducted analysis using crowd-sourced measurements, although on a much smaller scale. They cover only Singapore, and do not include LTE analysis due to the small sample size.

Our research also bears some similarities to [15] and [16], particularly in the method of data collection. The first of these works examined the correlation between download throughput and signal strength, to see if the latter could be used as a proxy when mapping the results. The focus of our research is slightly different; we are concerned with examining and explaining the dependence of download throughput on a number of factors, including, but not limited to, signal strength. We have also included additional signal metrics that give a more complete picture of the quality of LTE connection. The second is more in-line with our aims, but has a much smaller scope, and does not delve into the interactions and correlations between the explanatory factors.

Comparisons may also be drawn between our work and that of [14], which used another similar dataset collected via [3]. The work here was limited to 15 metropolitan areas worldwide, preventing analysis of operator-level differences within the same market. Notably this dataset only contained active tests initiated by the user, while the majority of our dataset consists of background tests whose scheduling is not linked to network performance.

The diurnal patterns in our model can be compared to [16], [10] and [14], which all observe decreases in performance during the day, although to differing extents. Each of these considered the overall correlation of throughput with hour, without controlling for other confounding factors, and neglected to include the weekday/weekend variable that we consider. Similar results were observed by [17] in broadband networks.

Close analysis in [7] demonstrated that the performance bottleneck on LTE connections has moved away from the network connection, as was the case with 3G networks [9], and on to the processing power of the device. This agrees with our findings, and that of [15] and [18], which show significant device-level differences.

II. Methodology

A. Data Collection

We collect data from mobile device users who install the OpenSignal [4] application on their device. Although the application is available on iOS devices, software limitations mean that we are unable to capture the LTE signal metrics, and so this current research is limited to devices running versions of Android. Once installed, the application displays various coverage metrics and information about the current connection, a history of the user's signal, and a map displaying coverage information for the local area.

The application also offers the opportunity for a detailed speedtest to be run. Provided that the user does not opt-out, the results of these speedtests are shared with OpenSignal. In addition to this a number of background tests are run at certain intervals, and the results are again saved and shared. These speedtests allow the user to accurately gain a realistic understanding of the quality of their mobile signal in the locations most relevant to them. This offers a significant advantage over other approaches which map signal coverage by drive-testing or similar methods, which give good results for roads, but which cannot safely be extrapolated to indoor or remote locations.

Download throughput is measured by attempting to download a large (108Mb) file from a popular and widespread CDN. The test is run for a fixed period, and the quantity of data transferred is used to calculate the average throughput. Multiple connections are opened to mimic the real-world behaviour of modern browsers. The OpenSignal application also uses the Android API to access and record over 120 variables concerning the users location, connection information and device features, and this data is combined with detailed summaries surrounding the results of the speedtest. No unnecessary personal information is stored, although we are able to identify the number of unique devices that contribute results by assigning each device a unique ID.

We have collected the results of 130,296 speedtests run by users of the OpenSignal mobile application, covering the period between December 2013 and February 2014, and located within the United States. We have extracted only those tests which take place over LTE connections, and restricted to a limited subset of the mobile device models. Since we wish to include a comparison of device performance we only include a selection of the most popular devices for which we will we be able to draw significant conclusions. These speed-tests are a mix of active tests initiated by the user (25%), and passive tests run in the background at application-controlled intervals (75%).

The OpenSignal measurement platform has a similar design to those of [2] and [3], using crowd-sourced measurements to give an accurate and up-to-date measurement of the current user experience across a representative mix of locations and devices.

Our approach differs from measurement platforms which collect data by performing drivetests [5], [1], in that it is not limited by the shape of the road network, or by the number of devices carried on each test.

B. LTE Signal Metrics

It is self-evident that there will be a strong dependence of download throughput upon the quality of the connection between the mobile device and the cell tower to which it is connected. A device with little or no signal cannot possibly hope to achieve fast download rates, and therefore we have need for an adequate metric which can capture this quality.

Obtaining signal strength metrics for LTE connections is not straightforward. In the past, when dealing with 2G and 3G technologies, there has always been an unambiguous and easily available metric - the Received Signal Strength Indicator (RSSI). This metric was useful when neighbouring cells could not share frequencies, and so it was possible to directly attribute the strength of a signal contained within one frequency to a single cell. As digital modulation technologies have progressed, and network operators have developed more complex and overlapping cell topologies, this restriction has been lifted so that this metric is of vastly decreased relevance.

The RSSI is purely an indicator of the in-band received signal power, and cannot discriminate between the contributions of the currently connected cell and those adjacent, overlapping cells which are of no relevance to current network connections, and consequently to the download throughput achieved. As a result this number cannot be relied upon to provide any connection-specific information, and we believe that confusion over this issue might explain why some existing results show no correlation between throughput and signal strength.

There are three LTE-only signal metrics which are available on Android devices. These are the Reference Signal Received Power (RSRP), the Reference Signal Received Quality (RSRQ), and the Reference Signal Signal to Noise Ratio (RSSNR). The Android API reports LTE signal strength using an RSRP measurement, but RSRQ and RSSNR are not directly available and must be collected via alternative means.

The first, RSRP, appears at first glance to be the most viable candidate for use as the signal strength indicator. We have illustrated the dependence of download throughput on this value in Figure 1. This metric is primarily used for cell selection and hand-over, and consequently gives a useful indication of the strength of the current connection. However it is not without problems - RSRP can be affected by multi-path fading, while the LTE standard was designed to cope with and mitigate fading, and so the actual throughput may be high in spite of a low RSRP reading. This can be observed in Figure 1, where we see how very low values of RSRP can sometimes correspond to very quick download rates.

The second of these, the RSRQ provides extra information that can be used to determine when to perform a handover. We do see a certain amount of dependence upon this value, but we speculate that this is purely acting as a proxy for the "general connectivity" of an area - perhaps an urban/rural measure.

Finally RSSNR is a measure of the signal-to-noise ratio of the reference signal, and so provides an indication of the quality of the link. This metric may be used to assess the impact of interference upon the connection, and so in urban areas where the RSRP is likely to be high in many locations, this might be the best measure of actual signal quality.

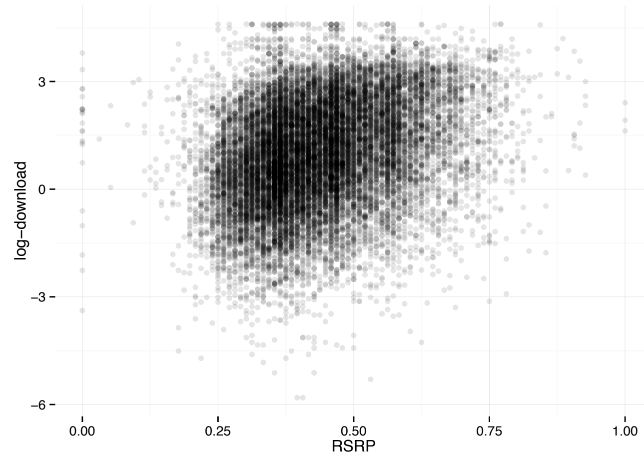

Fig. 1. Logarithm of estimated download throughput against reported values of RSRP. There is wide variation observed for almost any value of RSRP.

We hypothesise that a function of RSRP, RSRQ and RSSNR will be much more able to explain and predict the download throughput than any one variable alone. We have examined this claim by including all three variables in our linear model, and consequently see an improvement in the correlation. Figure 2 shows a plot of download throughput against the linear predictor we use. We certainly still observe significant amounts of variation in the result, but the situation is improved compared to Figure 1.

We must also understand that these values suffer from measurement error themselves, and it is likely that different models of phone will report these values to different levels of resolution. This is particularly apparent with the untransformed values of RSSNR, which take even values far more often than odd values. It could be the case that only some devices are able to measure this value precisely enough to produce fine-grained results.

III. Model

We have developed a simple linear model to analyse the dependence of download throughput on a number of available predictors. We include the available signal metrics, RSRP, RSRQ and RSSNR, the hour of the day and a weekday/weekend indicator, the device used, and the network operator that the device was connected to when the speed test was run. We allow for an operator and device interaction to capture the effect of operator-specific models of particular devices, which may run on different sets of frequencies. The weekday/weekend flag is used in conjunction with hour of day to allow for different usage patterns in the typical working week.

We take the log of download throughput as our dependent variable since it exhibits far less skew, and is consequently

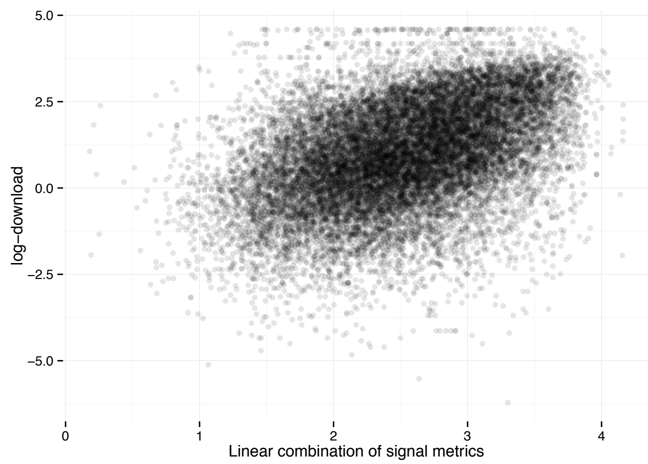

Fig. 2. Logarithm of download throughput against a linear predictor using RSRP, RSRQ and RSSNR together. Coefficients used are as in Table I.

easier to model using linear regression.

A. Normalisation and data-cleaning

To enable direct comparison between the signal metrics we have normalised them on a linear scale between zero and one, based upon their theoretical minimum and maximum attainable values. The rest of the model predictors are categorical variables, and so do not require normalisation. The temporal variables are based upon the user's local time, rather than UTC, allowing for realistic patterns to be observed.

In addition to this we clean the data by removing some obviously erroneous datapoints. A small number of results exhibit unrealistically high download speeds, which could not possibly be attained using any current network. We also discard any results that suffered a network technology change (e.g. a switch from LTE to 3G) while the speedtest was running.

B. Model Description

The model we have developed has an extremely simple structure - it is a linear model where the natural logarithm of the observed download speed is regressed upon the previously mentioned variables. The precise model is given by,

$$y_i = \alpha + \beta_1 s_{1,i} + \beta_2 s_{2,i} + \beta_3 s_{13i} + \delta_j + \nu_k + \zeta_{j,k} + \eta_{l,m} + \epsilon \quad (1)$$

where y_i denotes the log of download speed in Mbps for the ith observation, s_1, s_2, s_3 are the corresponding normalised signal metric values, and j, k, l and m indicate operator, device, hour and weekday/weekend factor membership respectively.

The model coefficients are fit by ordinary least squares regression. Visual inspection of diagnostic model plots show that that we obtain unbiased residuals, although there are signs of heteroschedasticity. The residuals also exhibit strong deviation from normality, which, while not a fatal flaw, means that we must be careful using any results which depend upon this assumption.

TABLE I
SELECTED MODEL COEFFICIENTS

	Coefficient	Standard Error
Intercept	-1.905	0.069
RSRP	0.995	0.046
RSRQ	1.607	0.033
RSSNR	1.857	0.039
Sprint	-0.730	0.061
T-Mobile	0.423	0.328
Verizon	-0.059	0.577

C. Model Verification

We have noted that we observe heteroschedasticity and non-normality of the residuals within our model, which could indicate problems and requires further verification. To check our model we have bootstrapped the regression coefficients, using 10,000 replications, and examined the bootstrap distributions of the coefficients. With such a large dataset we would expect the regression coefficients to be close to normal, and indeed the bootstrap distribution of these values does exhibit normality. Consequently we will assume that the confidence intervals reported for these coefficients, which are based upon standard theory, are indeed trustworthy.

IV. RESULTS

The model contains many too many terms to list in detail, so we extract some of the more interesting and influential findings.

A. Operators and Devices

We have collected data from only the largest four US network operators. An arbitrarily chosen operator, AT&T, is selected as the factor reference level. The coefficients assigned to the operator factors are summarised in Table I, and whose large values indicate the importance of including this variable in the model.

Note that in this work we do not attempt to characterise the network infrastructure in the vein of [19], instead we aggregate this consideration by including the network operator factor, which should capture key widespread differences in infrastructure.

The coefficients assigned to each device are on a similar scale to those assigned to each network operator. We note a particularly interesting result which shows that the Galaxy Note II has a larger positive coefficient than the Galaxy Note 3, which we speculate may be due to hardware changes in some aspects of the device.

B. Signal Metrics

The LTE signal metrics take integer values, and so cannot be directly compared to the factor variables. However as we have normalised the metrics we can compare the relative importance of each. As expected we see that the SNR ratio is a more influential predictor than the power indicator RSRP. This matches our intuition, and demonstrates that the signal strength metric

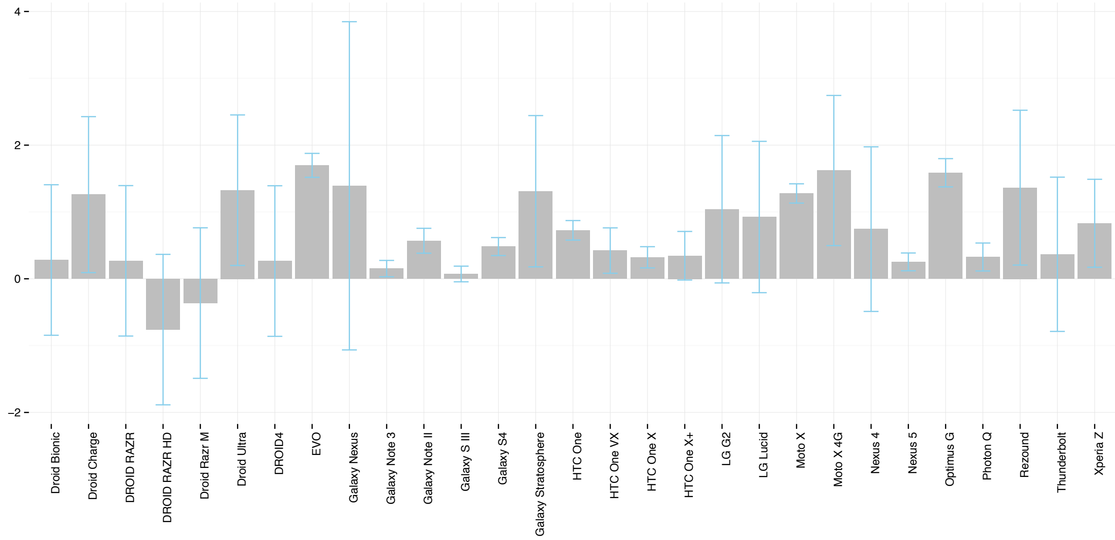

Fig. 3. Model coefficients for a selection of the most popular US devices. Factor reference level is an Motorola Atrix HD. 95% confidence intervals are indicated.

used by Android is in fact the least appropriate. Consequently we conclude that RSRP is an inadequate measure of signal, and cannot be relied upon for determining good or bad regions of LTE signal coverage.

Since the signal metrics are correlated to some extent, there is a risk that the coefficients could exhibit wild swings in response to small changes in the observed data. To check this we bootstrapped 10,000 model fits, with results indicating that this is not the case, and the standard errors reported in Table I appear to be accurate.

C. *Time of Day*

We recover a clear time of day effect, which shows the improvements to download throughput available overnight when cells are likely to be under reduced load. We introduce a weekday/weekend factor to allow for changes in behaviour that might occur around the typical working week, and indeed this is evident in the data. Figure 4 shows the model coefficients for the weekday hours, and Figure 5 for the weekend.

One interesting point is that we see a shifting of the intra-day trend between the weekdays and the weekend - the overnight increases both begin and end later, and the transition occurs more gradually over a longer period. We cannot say whether this is solely due to changes in user behaviour, operator network management, or (more likely) a combination of both. It is believed that operators react to the reduced demand overnight by switching off some cells, yet despite this reduction in capacity we still observe an improvement in throughput.

Time of day effects have been observed elsewhere [16], [12] without the weekday/weekend split, and have been calculated by averaging over all available data. By controlling for signal, device and operator we have removed potentially correlated predictors, and the remaining existence of a relatively strong and smooth pattern is particularly interesting.

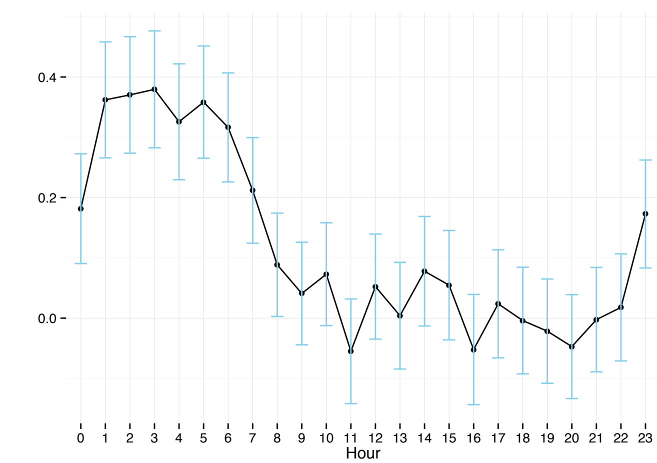

Fig. 4. Model coefficients for the weekday and hour factor interactions, with 95% confidence intervals indicated.

V. Conclusion

These results should be seen as a first step towards the development of a detailed model which can explain the relative importance of a number of predictors to download throughput. We have demonstrated a clear dependence on time of day, which has a sound a priori justification as a proxy for cell-congestion, and matches other results seen in the literature. We have also presented model coefficients for network, device and signal metrics, which indicate new and previously unknown device differences, and indicated large differences in network performance which are likely to result from variation in network backends. We have also calculated appropriate and

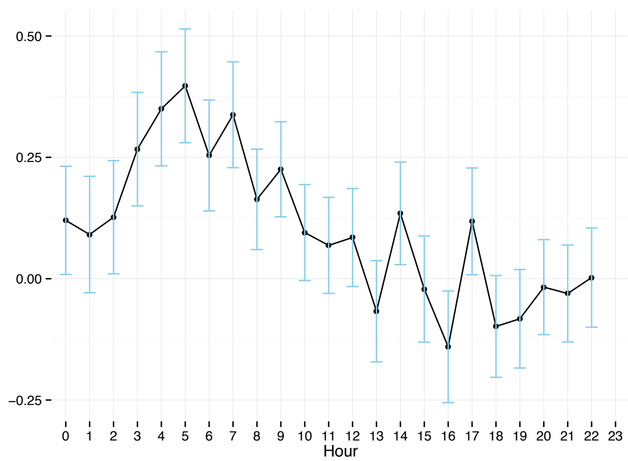

Fig. 5. Model coefficients for the weekend and hour factor interactions, with 95% confidence intervals indicated.

verifiable confidence intervals on these values. We are not aware of any other work which has determined a device-effect on throughput, and this extremely interesting, and potentially important, result will require further investigation. We suggest that these factors must all be taken into account when attempting to build explanatory models.

A. Directions for further research

We believe that there is much more that can be accomplished in this field, and many different areas remain unexplored. There are a large number of potential problems which only become tractable with large amounts of data, and the OpenSignal project is uniquely and perfectly positioned to capture this information.

The data used in this analysis covers only three months, and captures the early stages of the roll-out of LTE in the United States. In the future an increasing amount of the data we obtain will be LTE-focused, and this increase in information will allow for more detailed analysis to be carried out. In particular we envisage that finely-partitioned location-specific analysis will be possible in some regions in the near future, perhaps even on a city-by-city basis.

The United States was chosen as the basis for this research due to the large amount of LTE data available. It would be extremely enlightening to complete the same analysis in a different market, to assess whether the same factors are given the same importance. Of course the network and device-mix will differ greatly in some other markets, but there are enough similar devices available globally to make an interesting comparison.

We would also like to judge the dependence of download throughput on signal metrics in substantially different locations, with different urban/rural compositions and differing network infrastructures. We believe that more analysis is required of the relative importance of each LTE signal metric, and particularly in the accuracy of the reporting of these statistics by the operating software and manufacturers of each brand of device.

ACKNOWLEDGMENT

The authors would like to thank all of the OpenSignal mobile application users who have contributed data to this research.

REFERENCES

[1] Epitiro. [Online]. Available: http://www.epitiro.com
[2] Netradar. [Online]. Available: http://www.netradar.org
[3] Ookla speedtest.net. [Online]. Available: www.speedtest.net
[4] Opensignal. [Online]. Available: http://www.opensignal.com
[5] Rootmetrics. [Online]. Available: http://www.rootmetrics.com
[6] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris, "Link-level measurements from an 802.11 b mesh network," *ACM SIGCOMM Computer Communication Review*, vol. 34, no. 4, pp. 121–132, 2004.
[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 *Proceedings of the 10th international conference on Mobile systems, applications, and services*. ACM, 2012, pp. 225–238.
[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," in *Proceedings of the ACM SIGCOMM 2013 conference on SIGCOMM*. ACM, 2013, pp. 363–374.
[9] J. Huang, Q. Xu, B. Tiwana, Z. M. Mao, M. Zhang, and P. Bahl, "Anatomizing application performance differences on smartphones," in *Proceedings of the 8th international conference on Mobile systems, applications, and services*. ACM, 2010, pp. 165–178.
[10] X. Liu, A. Sridharan, S. Machiraju, M. Seshadri, and H. Zang, "Experiences in a 3g network: interplay between the wireless channel and applications," in *Proceedings of the 14th ACM international conference on Mobile computing and networking*. ACM, 2008, pp. 211–222.
[11] K. Papagiannaki, M. D. Yarvis, and W. S. Conner, "Experimental characterization of home wireless networks and design implications." in *INFOCOM*, 2006.
[12] S. Sen, J. Yoon, J. Hare, J. Ormont, and S. Banerjee, "Can they hear me now?: a case for a client-assisted approach to monitoring wide-area wireless networks," in *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference*. ACM, 2011, pp. 99–116.
[13] C. Shepard, A. Rahmati, C. Tossell, L. Zhong, and P. Kortum, "Livelab: measuring wireless networks and smartphone users in the field," *ACM SIGMETRICS Performance Evaluation Review*, vol. 38, no. 3, pp. 15–20, 2011.
[14] J. Sommers and P. Barford, "Cell vs. wifi: on the performance of metro area mobile connections," in *Proceedings of the 2012 ACM conference on Internet measurement conference*. ACM, 2012, pp. 301–314.
[15] S. Sonntag, J. Manner, and L. Schulte, "Netradar-measuring the wireless world," in *Modeling & Optimization in Mobile, Ad Hoc & Wireless Networks (WiOpt), 2013 11th International Symposium on*. IEEE, 2013, pp. 29–34.
[16] S. Sonntag, L. Schulte, and J. Manner, "Mobile network measurements-it's not all about signal strength," in *Wireless Communications and Networking Conference (WCNC), 2013 IEEE*. IEEE, 2013, pp. 4624–4629.
[17] S. Sundaresan, W. De Donato, N. Feamster, R. Teixeira, S. Crawford, and A. Pescapè, "Broadband internet performance: a view from the gateway," in *ACM SIGCOMM Computer Communication Review*, vol. 41, no. 4. ACM, 2011, pp. 134–145.
[18] Q. Xu, J. Erman, A. Gerber, Z. Mao, J. Pang, and S. Venkataraman, "Identifying diverse usage behaviors of smartphone apps," in *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference*. ACM, 2011, pp. 329–344.
[19] Q. Xu, J. Huang, Z. Wang, F. Qian, A. Gerber, and Z. M. Mao, "Cellular data network infrastructure characterization and implication on mobile content placement," in *Proceedings of the ACM SIGMETRICS joint international conference on Measurement and modeling of computer systems*. ACM, 2011, pp. 317–328.
[20] Y. Xu, Z. Wang, W. K. Leong, and B. Leong, "An end-to-end measurement study of modern cellular data networks," in *Passive and Active Measurement*. Springer, 2014, pp. 34–45.