An Empirical Characterization of Cellular Network Performance

Anand Seetharam, Peter Walker
Computer Science Program
California State University Monterey Bay
aseetharam, pwalker@csumb.edu

***Abstract*—With the proliferation of smartphones and the growth of cellular technology (e.g., 4G, LTE), individuals are constantly accessing the Internet on the go. But, there is limited understanding of the variation in cellular network performance (i.e., throughput, latency, packet loss) for various wireless service providers. In this paper, we perform an empirical characterization of cellular network performance based on a large scale measurement study of the major wireless service providers in California. We describe the system and software application (CalSPEED) used to measure various network parameters such as TCP download/upload throughput, latency, packet loss and jitter for the different wireless service providers. Six rounds of measurement is conducted for a period of approximately three years in California, (referred to as field tests) starting from spring 2012. During each field test, personnel drove with mobile clients to different parts of the state and collected data by executing the software. Data is collected over approximately 1500 locations in the state during each field test[1]. Our data analysis demonstrates quantitatively that *i)* cellular network performance varies considerably with location and is poor overall - a large number of locations in our study report low network throughput, high latency and high loss rates *ii)* all cellular providers are not the same - network performance varies considerably among the different providers *iii)* on average, cellular network performance is improving over time *iv)* cellular network performance depends on a number of factors such as download/upload, server location and type of device used.**

I. Introduction

The stupendous growth of smartphones and cellular technology (e.g., 4G, LTE) has resulted in a more connected world than ever before. In our increasingly mobile world [1], it is of prime importance to have a clear understanding of the connectivity, coverage and performance provided by cellular providers, so that individuals can select a cellular provider that matches their needs. For example, an individual's job might require her to travel across the state and she would want to find a cellular provider that works best across the different regions she visits. Currently, most providers have maps regarding their coverage and connectivity on their websites [2], but there is no simple way for individuals to compare among the different providers, primarily due to the lack of available data.

In this paper, we perform an empirical characterization of the cellular network performance in the state of California using data obtained from a large scale measurement study. The data used in this paper measures the network performance of different cellular providers over 1500 locations for a period of approximately three years in the state of California. In this paper, we describe the system for collecting these network measurements and our experimental results. Our analysis quantitatively demonstrates that cellular network performance is improving over time. We also observe that device type, service type (upload/download) and server location, all impact performance.

To the best of our knowledge, this is the first work reporting results for a measurement study conducted over such an extensive scale. The authors in [3], [4], [5], [6] characterize the performance of 3G and 4G networks, but the experimentation is much smaller in scale (primarily in an university setting with measurements collected over a short duration of time). Our work is closest to [6], where the authors investigate the characteristics of three mobile ISPs in Singapore. Based on their measurements, they conclude that packets tend to arrive in bursts. They also observe large variations in instantaneous throughput. The authors in [7] study the impact of network protocols and application behavior on user level performance over LTE. In [8], Sprout - an end-to-end transport layer protocol that achieves high throughput and low delays in cellular networks is proposed.

In contrast to all the above mentioned works, cellular network measurement data used in this paper is collected over a large geographical region for an extended period of time. We investigate the evolution of network performance (throughput, latency, loss) over time and study the impact of various factors (e.g., device type, server location) on network performance. We acknowledge that some of the results in this paper are already available in public documents [9], [10] and in [11]. The main difference between this paper and [9], [10], [11] is that this paper provides an overview of the system architecture used for obtaining network measurements and presents a holistic view of cellular network evolution in the state. This paper makes the following contributions.

- We describe the system and software used to conduct extensive cellular network measurements throughout the state of California (Section II). The system (Figure 1) consists of mobile clients (Windows laptops, Android phones) and two servers - one on the west coast (in Northern California) and the other on the east coast (in Northern Virginia). It also has a database to store all

[1]The data used in this paper is open source and was collected via joint collaboration between California State University Monterey Bay (CSUMB), California State University Chico, California Public Utilities Commission and Novarum Inc. The authors of this paper were involved in data collection and analysis as part of the collaboration. The CalSPEED application was mainly designed at CSUMB and many individuals were involved in data collection.

measurement data. The software application (CalSPEED) measures TCP upload and download throughput, latency, loss and jitter for four wireless service providers - AT&T, Sprint, T-Mobile, and Verizon. CalSPEED internally uses network measurement tools such as iperf and ping to determine overall network performance and then sends this measurement data to the database.

- Data used in this paper is collected via six rounds of extensive measurements (each called a field test) for a period of approximately three years starting from spring 2012 (Section III). The first three field tests are carried out in 1200 locations while the last three fields tests are carried out in 1986 locations (the locations are chosen such that most regions in the state are covered). Our results based on these measurements show that average network performance (throughput, latency) is improving over time. We also observe that all network providers are not the same; some providers (Verizon) have significantly better performance when compared to some other providers (Sprint). Additionally, overall network performance is poor with a large fraction of measurement locations (Table II) reporting low or no connectivity (Section IV).
- Our results also demonstrate that network performance depends on a number of factors such as download/upload, server location and type of device used (Section IV). For example, TCP download throughput is higher is in comparison to TCP upload throughput.

II. System Design

The goal of this paper is to quantify the performance of the different cellular network providers in the state of California via extensive network measurements. To meet this goal, we use network measurements collected from locations spread across the state. These measurements help us compare the performance of different cellular providers and study the evolution of cellular network performance over time. In this section, we discuss the overall system architecture for performing extensive cellular network measurements. Field personnel perform six rounds of data collection at multiple locations in the state for a period of approximately three years (additional details in Section III).

Figure 1 describes the system architecture. It consists of mobile clients (laptop, android phone) and two servers, one on the west coast (in Northern California) and the other on the east coast (in Northern Virginia); two servers located on two coasts increase server location diversity. All measurement results are directly uploaded to a database. Mobile clients (Windows based laptop with USB modems and Android smartphones) are used to collect cellular network data for four wireless service providers (AT&T, Sprint, T-Mobile, and Verizon).

To help conduct a comprehensive measurement study, the software records the time of the measurement, location (i.e., latitude and longitude), round trip time (RTT), TCP download and upload throughput, UDP jitter and datagram loss, network type (i.e., EDGE, EVDO, GPRS, HSPA, LTE, UMTS, etc), and network provider (i.e., AT&T, Sprint, T-Mobile, Verizon).

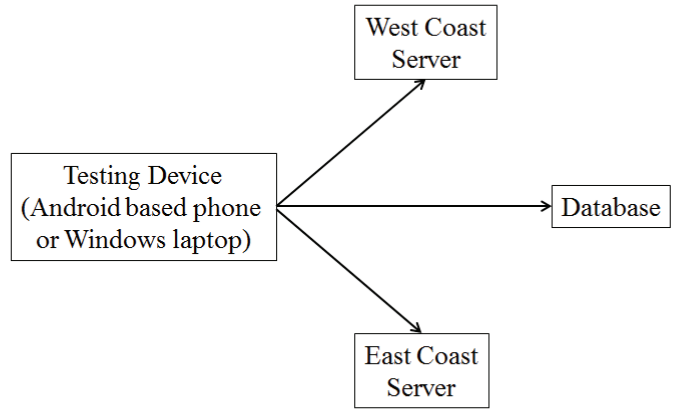

Fig. 1. System architecture for cellular network measurement

Iperf [12] is used to measure upload and download throughput of TCP connections between the mobile client and the server. Iperf uses a client-server architecture where the client requests a TCP connection to the server. Once the connection is established, iperf measures the upload and download throughput between the client and server. The TCP window size is set to 64 Kbytes (-w 64k) in iperf; four threads execute concurrently (-P 4), so as to increase the data exchanged between client and server.

Iperf can also measure data loss and jitter by sending UDP packets. For UDP measurements, the iperf client sends data to the server for a duration of five seconds. Note that no data is sent from the server to the client in case of UDP measurements. The parameters used in the UDP measurements are 220K buffer length (-l 220) and 88K bits/sec bandwidth (-b 88k).

The software uses ping to measure the minimum, maximum, and average RTT between the client and server. Ping also reports packet loss during each measurement. To accommodate for the fact that some test locations might have weak signal strength, a timeout of 2 minutes is added for each test.

When the software executes on a mobile client at a particular location for a specific wireless provider, it runs in the following order.

- GPS capture: A GPS device records the latitude and longitude of the test location.
- Connectivity check: Before conducting the suite of measurements, the software sends multiple ICMP packets. If the software does not obtain a response from the server for any of these packets, it indicates that the area has very weak signal strength or no signal for the provider and the software stops its execution. Otherwise, it moves on to the next step in the test sequence.
- Iperf TCP test: The software uses iperf to conduct two TCP upload and two TCP download measurements between the client and the east and west coast servers respectively. The -t option in iperf is set to 10 seconds which indicates the time window for data transmission. Note that even though data transmission takes place for 10 seconds, the measurement duration can be much longer especially in weak signal strength regions, primarily due

to lost TCP acknowledgments and timeouts. To account for these cases, a 2 minute timeout is added to each iperf TCP test.

- Ping test: The ping test measures the RTT to the west and east coast servers. Each ping test sends 10 ICMP packets to the server.
- Iperf UDP test: Iperf is used to measures jitter and datagram loss.
- Test result upload: Once all measurements are completed, the results are uploaded to a database.

III. Experiment Methodology

In this section, we outline the detailed experiment methodology. Field testers conducted six rounds of data collections (referred to as field tests) starting from spring 2012. The number of locations tested during a field test has been continuously increasing. The first three field tests were conducted across 1200 locations in the state, while the last three field tests were conducted across 1986 locations.

During each of these field tests, personnel equipped with four smartphones (one for each cellular provider) and a laptop (with USB modems for each cellular provider) drove to different parts of the state and collected data by executing the software. As the experiment time frame is approximately three years, data was collected across multiple devices - ASUS laptop with different USB modems for the four carriers and different smartphones (HTC One X, HTC EVO, Samsung Galaxy SII, Samsung Galaxy SIV.). At each location, the personnel execute the software in a stationary environment inside an automobile.

Details related to the sampling of the various locations is provided in [10]. The locations were selected in urban, rural, and tribal areas. For the first field test, the breakdown of 1200 locations is 34% urban, 11% tribal and 55% rural, based on Census 2010 definitions [10]. Because of the extensive scale of the experiments, multiple measurements were not conducted at the same location at different times of the same day. Additionally, there is a possibility that multiple tests are conducted for the same provider in a single location because of technical reasons or human error, especially in weak signal strength regions. For example, when a personnel starts an experiment for a provider, the software executes the ping test to the west coast server to check wireless connectivity. If the location fails the connectivity test for the provider, the tester might try multiple times before concluding that the location has “no effective service”. As a result the number of tests conducted for each provider is usually higher than the number of locations tested. Table I presents the total number of tests conducted on laptops per provider for the last two field tests (note that the number of locations tested is 1986).

Table II shows the percentage of successful TCP upload/download throughput measurements. There are multiple reasons for failed tests. One of most common reasons is that the device failed to obtain a valid signal from the provider (“no effective service”). Other possibilities include timeout due to weak signal, software error and phone issue. We observe from the data that majority of failed tests can be attributed to weak signal strength. Though there are multiple reasons

Provider	Number of Tests (5^{th} Field Test)	Number of Tests (6^{th} Field Test)
AT&T	2094	2126
Sprint	2073	2162
T-Mobile	2061	2138
Verizon	2122	2107

TABLE I
Number of tests for each provider

Provider	Successful Test Percentage (5^{th} Field Test)	Successful Test Percentage (6^{th} Field Test)
AT&T	78.9%	76.7%
Sprint	64.4%	59.5%
T-Mobile	60.2%	58.9%
Verizon	81.3%	81.5%

TABLE II
Percentage of successful tests for each provider

for failed measurements, we believe that the percentage of successful tests can be used as a meaningful indicator of a provider’s coverage. Having provided an overview of the system design and the experiment methodology, we next discuss the experimental results.

IV. Results

In this section, we report network performance results (throughput, latency and loss) using data obtained from the field tests. At the highest level, our results indicates the following, 1) Cellular network performance varies considerably with location and is poor overall - a large number of locations in our study report low network throughput, high latency and high loss rates. 2) All cellular providers are not the same - the network performance varies considerably among the different providers. 3) On average, the cellular network performance is improving over time. 4) Network performance depends on factors such as download/upload, server location and type of device used.

A. Throughput

We report throughput results obtained from the iperf TCP measurements. We use box and whisker plots to summarize the download throughput (considering all devices) for the different cellular providers for the six field tests (Figure 2). The line inside each box denotes the second quantile (median), while

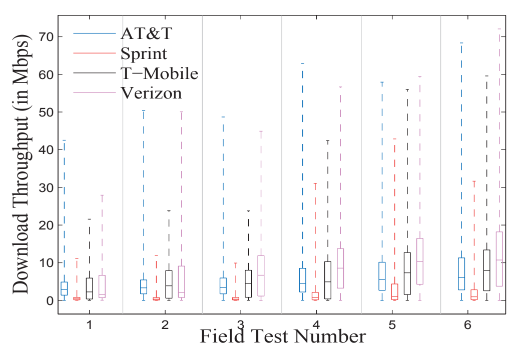

Fig. 2. Throughput Distribution

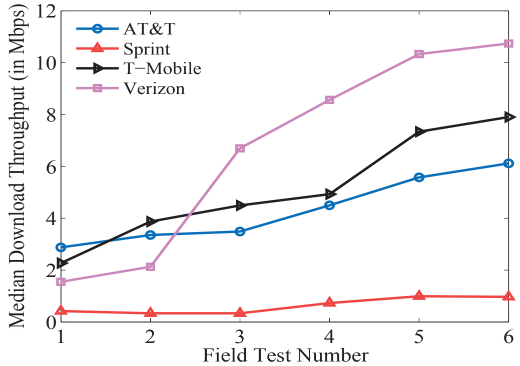

(a) Download Throughput

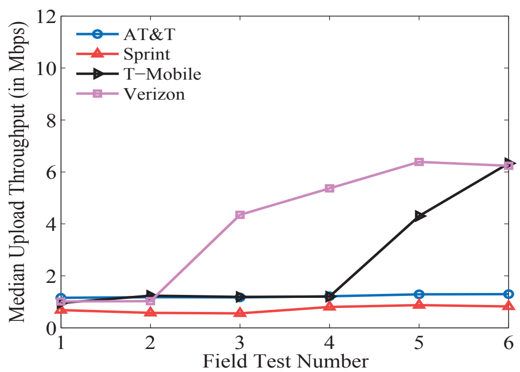

(b) Upload Throughput

Fig. 3. Download vs. Upload Throughput

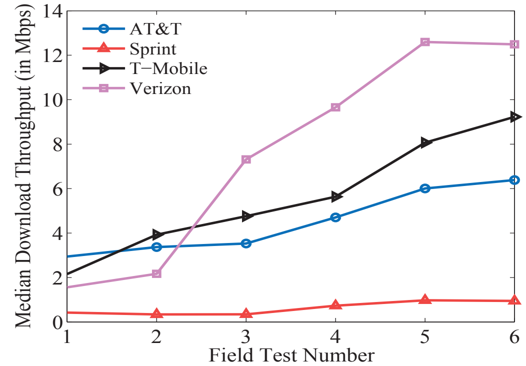

(a) West Server Throughput

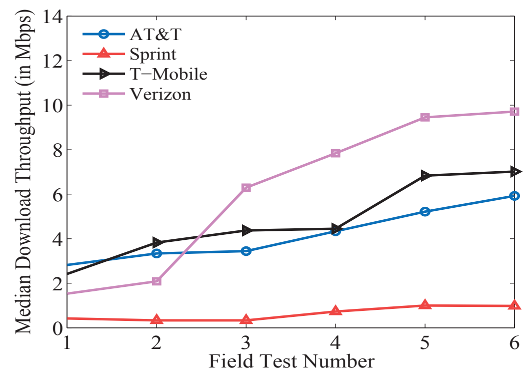

(b) East Server Throughput

Fig. 4. West Server vs. East Server Throughput

the top and the bottom lines of each box are the first and the third quantiles (25% and 75%). The ends of the whiskers are the minimum and maximum values.

From Figure 2, we observe that there is a large variation in the download throughput across the test locations. From Figure 2, it is evident that while many locations have low throughput, there are also some locations which provide very high throughput. We also observe from the figure that all providers are not the same. The performance of AT&T, Verizon and T-Mobile appear to be significantly better than Sprint. For ease of understanding, we present median throughput results for the different providers.

Download vs. Upload: Figures 3(a) and 3(b) show the variation in median download and upload throughput respectively for the different cellular providers for the different field tests. From the figures it can be easily observed that the median download and upload throughput for all providers is steadily improving over time with the improvement for Verizon being higher than others. The figures also show that T-mobile has made significant improvement in throughput during the last two field tests. The throughput of Sprint is much lower compared to the other three providers. Additionally, the upload throughput is significantly lower in comparison to the download throughput (there is a reduction of approximately 40%).

Server Location: The above results are obtained by considering experiments conducted to servers located on both east and west coasts. In Figures 4(a) and 4(b), we demonstrate the impact of server location on median download throughput. We observe from the figure that median throughput obtained for TCP connections to the west coast server is higher compared to that obtained for TCP connections to the east coast server. We attribute this lower throughput to network congestion. As reaching the east coast server requires a connection to pass over larger number of links, the probability of at least one of those links being congested increases, thereby leading to lower throughout. We obtain similar results for upload throughput, but we omit them due to lack of space; they are available in our technical report [13].

We also study the impact of device diversity (laptops, mobile phones) on performance. We observe that the throughput obtained via mobile phones is slightly higher in comparison to that obtained from laptops for both upload and download. These additional results are available in [13].

B. Latency (Round Trip Time)

We present a box and whisker plot for our RTT measurements (Figure 5). We observe that some RTT values are much higher in comparison to majority of our measurements; therefore, we limit the maximum value depicted in the box plot to 600 msec. From the box plot it is evident that the median RTT values and the RTT variation for Sprint and T-Mobile is much higher in comparison to AT&T and Verizon.

Server Location: Figures 6(a) and 6(b) show the variation in median RTT over all locations for the servers located in the

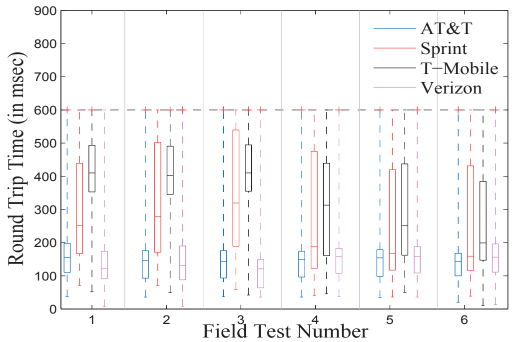

Fig. 5. RTT Distribution

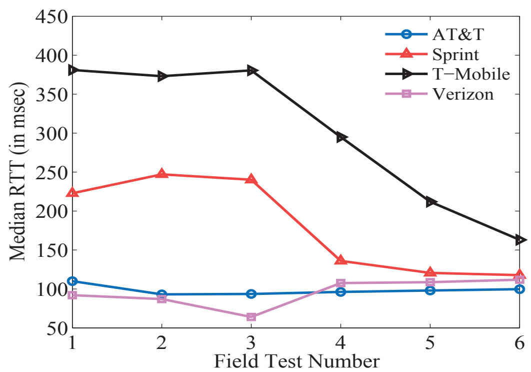

(a) West RTT

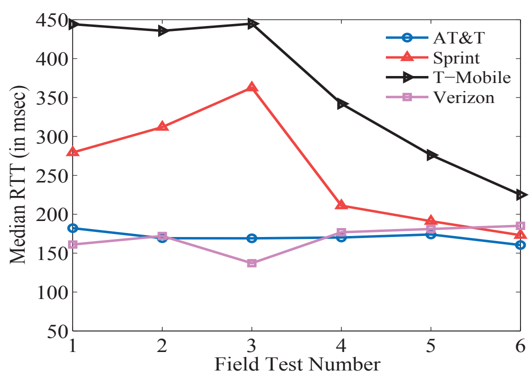

(b) East RTT

Fig. 6. Round Trip Time (RTT)

west and east coast respectively. Note that at each location, 10 ping tests were performed and the average RTT of these tests is calculated at each location; these average values are then used for determining the median. We omit the failed ping tests (i.e., time out or destination unreachable errors) from our RTT calculations.

As expected, the RTT to the west coast server is lower than that to the east coast server. The most surprising result is the performance of T-mobile. We observe from our experiments that though the throughput of T-mobile is high, its RTT is also high. Both AT&T and Verizon have low and comparable RTT. We observe that the performance of Sprint is poor (high RTT) for the first three field tests.

C. Packet Loss and Jitter

In terms of packet loss, we observe that the performance of AT&T and Verizon are comparable and better than both T-Mobile and Sprint. In case of jitter, we observe that Sprint has poor performance when compared to the other providers, which have similar performance. We omit the results due to lack of space; they are available in [13].

V. Conclusion

In this paper, we outlined the design of the software used to measure the cellular network performance (e.g., throughput, latency) of major wireless service providers in California. Extensive network measurements at approximately 1500 locations in the state were collected via field tests for a period of around three years during which time six field tests were conducted. Based on these measurements, we observed that network performance has been steadily improving over time. We also observed that there is significant difference in the performance provided by the different providers in the state. Additionally network performance is dependent on a number of factors such as download/upload, server location and type of device used.

VI. Acknowledgment

We would also like to acknowledge the effort of students and faculty at CSUMB who have been involved in the project at different points in time. We would like to acknowledge the effort of the various personnel involved in the field tests; without their help we would have never been able to collect the data. We would also like to thank the employees at CSU Chico, CPUC and Novarum, Inc involved in this project for their contribution.

References

[1] "The mobility first project," http://mobilityfirst.winlab.rutgers.edu/.

[2] "At&t coverage map," http://www.att.com/maps/wireless-coverage.html.

[3] Y.-C. Chen, E. Nahum, A. Liu, R. Gibbens, and D. Towsley, "Measuring cellular networks: Characterizing 3g, 4g, and path diversity," in *ACITA*, 2012.

[4] M. Laner, P. Svoboda, E. Hasenleithner, and M. Rupp, "Dissecting 3g uplink delay by measuring in an operational hspa network," in *PAM*, 2011.

[5] A. Elmokashfi, A. Kvalbein, J. Xiang, and K. R. Evensen, "Characterizing delays in norwegian 3g networks," in *PAM*, 2012.

[6] Y. Xu, Z. Wang, W. K. Leong, and B. Leong, "An end-to-end measurement study of modern cellular data networks," in *PAM*, 2014.

[7] J. Huang, F. Qian, Y. Guo, Y. Zhou, Q. Xu, Z. M. Mao, S. Sen, and S. Oliver, "An in-depth study of lte: Effect of network protocol and application behavior on performance," in *ACM Sigcomm*, 2013.

[8] K. Winstein, A. Sivaraman, and H. Balakrishnan, "Stochastic forecasts achieve high throughput and low delay over cellular networks," in *USENIX NSDI*, 2013.

[9] K. Biba, "Calspeed: California mobile broadband - an assessment," ftp://ftp.cpuc.ca.gov/telco/BB%20Mapping/Field%20Testing/Biba%20Mobile%20Broadband%20Assessment%209%204%2014%20filed.pdf.

[10] "California public utilities commission - field test reports," http://www.cpuc.ca.gov/PUC/Telco/bb_drivetest.htm.

[11] Y. Byun, S. Narayanan, S. Mott, K. Biba, J. Schwenkler, R. Osborn, and M. Morris, "Wireless broadband measurement in california)," in *ITNG*, 2013.

[12] "iperf," https://iperf.fr/.

[13] A. Seetharam and P. Walker, "An empirical characterization of cellular network performance," http://itcdland.csumb.edu/∼aseetharam/paper/ICNC-techreport.pdf.