California Public Utilities Commission

Measuring Mobile Broadband in California

Round 3 Mobile Field Test Results

April 2013

BROADBANDUSA
CONNECTING AMERICA'S COMMUNITIES

December 2013

Edmund G. Brown, *Governor*

CALIFORNIA PUBLIC UTILITIES COMMISSION

COMMUNICATIONS DIVISION

BROADBAND, POLICY & ANALYSIS BRANCH

Robert Osborn, CPUC
Christian Ettinger, CPUC
Ken Biba, Co-Founder and CTO, Novarum
Owen Rochte, CPUC
Principal Authors

Michael Morris
Program Supervisor
VIDEO FRANCHISING AND BROADBAND DEPLOYMENT

Rob Wullenjohn
Program Manager
BROADBAND, POLICY AND ANALYSIS

Ryan Dulin
Director
COMMUNICATIONS DIVISION

Paul Clanon
Executive Director

DISCLAIMER

Acknowledgments

We would like to thank the following organizations for making this study possible:

- The National Telecommunications and Information Administration (NTIA), who provided the grant money to do this study.

- Ken Biba, Co-founder and Chief Technology Officer of Novarum, for devising the testing methodology and for providing overall guidance on test design and results interpretation.

- The Computer Science and Information Technology program at California State University, Monterey Bay, for designing, building, and supporting the mobile testing applications used for this study.

- The Geographical Information Center at California State University, Chico, and the dedicated field testing team who performed the actual tests.

Table of Contents

Executive Summary

The California Public Utilities Commission (CPUC) is the California recipient of a State Broadband Data and Development Grant, awarded by the National Telecommunications and Information Administration (NTIA) under the American Recovery and Reinvestment Act (“ARRA”). The Grant funds certain broadband mapping and planning projects through October, 2014, including a project to both create a mobile app that can be used by the public to measure and report actual mobile broadband connection quality, to perform drive tests to be conducted by the CPUC at six month intervals to measure service quality in urban and rural areas, and on tribal lands within the state. Analysis of these results is used to help formulate state broadband policy and guide state-directed broadband spending. Pursuant to the Grant, the CPUC produces a testing summary report for each round of testing. This report is intended to fulfill that requirement.

This Report is a follow up to the “Spring 2012 Mobile Broadband Field Testing Initial Staff Report: Preliminary Findings.” It contains updated data from our second and third rounds of mobile field testing, conducted in October, 2012 and April, 2013, respectively. For each round of testing, testers drove over 35,000 miles to take measurements at 1,200 locations (34% urban, 11% tribal and 55% rural).The purpose of this Report is to present the methods, procedures, results and initial findings associated with the Commission’s mobile field testing.

In comparing results from all three rounds of testing, we made the following findings:

- The only carrier to deliver greater than advertised speeds across all census areas was Verizon, but we suspect that is due to the time gap between the availability data and the date when we performed the field tests. Sprint and T-Mobile appeared to under-deliver downstream speeds in all census areas, while and AT&T under-delivered upstream speeds in rural and tribal locations.

- Verizon continued to make dramatic improvements in upstream and downstream speeds in all census areas, particularly in tribal and rural areas, which are areas generally characterized by substandard broadband service. Improvements for the other three carriers were more incremental.

- All providers continue to make improvements in the number of locations qualifying with a Mean Opinion Score (MOS) of 4.0 or greater, which is one indicator of VoIP readiness. While the number of T-Mobile locations achieving a MOS score of 4.0 was extraordinarily low, we saw a dramatic increase in T-Mobile’s performance from October 2012 to April 2013.

1 Mobile Broadband Testing

Broadband access technologies fall into two categories: wireline and wireless. Wireless broadband access is provided in two forms: fixed and mobile. In California, mobile broadband is provided by cellular phone service providers, also known as "terrestrial mobile wireless carriers." The principal mobile wireless carriers in California are Verizon Wireless, AT&T Mobility, Sprint, and T-Mobile, although there are others that are available, including MetroPCS, and US Cellular. These carriers offer voice as well as data services, but for this study we only measured mobile broadband data service elements. The tests conducted for this report include areas where carriers offer roaming.

As defined by the FCC, mobile service qualifies as "broadband" if the downstream throughput is 768 kilobits per second or higher and upstream throughput is 200 kilobits per second or higher. For this study, we collected measurements even when they fell below this threshold. This included locations in California where no effective service was available for any mobile operator. In cases where the ping test failed, we recorded that test location for that particular provider under the category "No Effective Service." See Appendix F for a list of Abnormal Value Descriptions.

Mobile broadband field testing is part of a larger broadband mapping project funded by a grant to the CPUC from the National Telecommunications Information Administration's (NTIA) State Broadband Initiative program (SBI). The purpose of this program is to map broadband Internet access in all 50 states to facilitate efforts by the public and policy makers to increase broadband access. The CPUC collects broadband data from providers in California twice a year and provides it to the NTIA. This data feeds into the National Broadband Map as well as the feature-rich California Interactive Broadband Map (http://www.broadbandmap.ca.gov).

As part of the project, the CPUC's Communications Division is conducting mobile broadband field tests twice a year. In order to do so, the CPUC contracted with California State University, Chico to hire and manage a team of testers to perform speed and availability tests of the four major carriers (Verizon, AT&T, Sprint, and T-Mobile) at 1,200 locations throughout the state. A map of these locations may be found in Appendix C – Test Locations. The breakdown of 1,200 locations was 34% urban, 11% tribal, and 55% rural, based on Census 2010 definitions[1].

We used open source tools for our tests, and the tests were performed using commercially available smartphones and data cards. Three rounds of tests were performed at the same locations in May 2012, October 2012, and April 2013.

Summary results from the drive tests are provided in this report and results have been used to validate provider data incorporated into the California Interactive Broadband Map.[2] The CPUC has also made the

---

[1] The "urban" category includes both "urbanized areas" (UAs) of 50,000 or more people and "urban clusters" (UCs) of at least 2,500 and less than 50,000 people. "Rural" encompasses all population, housing, and territory not included within an urban area. The "tribal" category is the same as referenced in Appendix D – Test Routes.

[2] Mobile testing data that contradicts provider data is evaluated with other validation data and may result in adjustment to reported data and the display of served areas in California.

data from these tests available to the public online via the CPUC web site, in order that others can perform their own analysis.

2 Mobile Test Application

The CPUC contracted with California State University, Monterey Bay to develop an open source mobile testing application that collects mobile broadband measurements in real-world conditions. This means testing a number of different traffic types and testing them on servers physically located on opposite coasts of the United States in order to take into account each carrier's varying backhaul networks.

The application uses iPerf[3], which is an open source network testing tool that creates TCP[4] and UDP[5] (data streams and measures performance of each over the network. CSU Monterey Bay developed a Java user interface that runs on both Android and the netbook version of Windows operating systems. The data created by the test is uploaded automatically to a cloud-based database server. Details of the record format of the test results are included in Appendix F – Data Record Format.

Testers ran the testing application in a stationary environment inside an automobile. Each tester was equipped with a netbook, an external GPS antenna, and a smartphone and data card for each of the four carriers. Tests were done sequentially, first by getting a valid GPS reading from the GPS receiver connected to a netbook, then by running the tests using data cards for each carrier. Finally, the smart phone tests were performed. Testers uploaded the results at each location to the cloud-based database server. In cases where the upload failed at the tested location (generally due to inadequate speeds), the test results remained on the netbook or smart phone until the tester enters a location where sufficient network service allowed the data to be successfully uploaded.

In cases where there was no signal whatsoever, that result was also stored on the netbook or smartphone for later upload. The testing application ran ten 1-second ping tests to both west and east servers. Following that, it ran ten 1-second TCP tests to the west server and then ten 1-second TCP tests to the east server. It then repeated that sequence a second time. This was done for both upstream and downstream connections. The application then ran three 1-second UDP jitter and datagram loss tests to both west and east servers. Finally, the application ran one 5-second UDP jitter and datagram loss tests to both west and east servers.

---

[3] Iperf was developed by the Distributed Applications Support Team at the National Laboratory for Applied Network Research. The benefit of being open source is the measurement methodology can be scrutinized by anyone, as opposed to proprietary testing tools, which do not always provide sufficient description of how they measure network performance. See Appendix A – Test Application for how we customized iPerf for field testing.

[4] Transmission Control Protocol (TCP) is the protocol used to upload and download files and E-mail.

[5] User Datagram Protocol (UDP) is the protocol used for real-time streaming applications like streaming video and Voice over Internet Protocol, or VoIP.

The table below shows the number of tests per device per carrier for east and west coast servers for each test location.[6] The total number of tests done at each location is 108 tests multiplied by 2 devices (data card and phone) multiplied by 4 carriers (AT&T, Sprint, T-Mobile, and Verizon), which equals 864 tests at each of the 1,200 locations.

Test	West	East	Number of times	Total Tests
Ping	10	10	1	20
TCP Up	10	10	2	40
TCP Down	10	10	2	40
UDP 1 sec	3	3	1	6
UDP 5 sec	1	1	1	2
Total Tests per location per device per carrier				108

The image below is an example of the format of raw data for the TCP upstream and downstream tests for the west coast server:

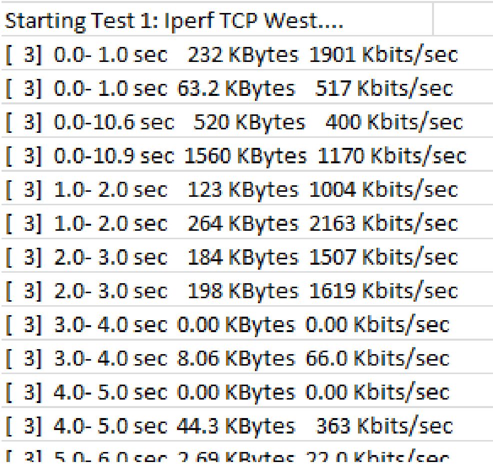

```
Starting Test 1: Iperf TCP West....
[  3]  0.0- 1.0 sec   232 KBytes  1901 Kbits/sec
[  3]  0.0- 1.0 sec  63.2 KBytes   517 Kbits/sec
[  3]  0.0-10.6 sec   520 KBytes   400 Kbits/sec
[  3]  0.0-10.9 sec  1560 KBytes  1170 Kbits/sec
[  3]  1.0- 2.0 sec   123 KBytes  1004 Kbits/sec
[  3]  1.0- 2.0 sec   264 KBytes  2163 Kbits/sec
[  3]  2.0- 3.0 sec   184 KBytes  1507 Kbits/sec
[  3]  2.0- 3.0 sec   198 KBytes  1619 Kbits/sec
[  3]  3.0- 4.0 sec  0.00 KBytes  0.00 Kbits/sec
[  3]  3.0- 4.0 sec  8.06 KBytes  66.0 Kbits/sec
[  3]  4.0- 5.0 sec  0.00 KBytes  0.00 Kbits/sec
[  3]  4.0- 5.0 sec  44.3 KBytes   363 Kbits/sec
[  3]  5.0- 6.0 sec  2.69 KBytes  22.0 Kbits/sec
```

The raw data were averaged to create a summary report for each location/device/carrier combination. Below is a calculation showing the number of records in total for a single round of testing. We publish the Summary data (averaged) on the CPUC web site. Data for the April 2013 test is now available. The raw data is available upon request.

---

[6] At a subset of locations, testers ran a separate test using an open source tool called "Glasnost" in order to measure the presence of carrier rate limiting behavior, or "traffic shaping." This paper focuses exclusively on TCP and UDP results for netbook and smartphones using both West and East coast servers. Analysis of the Glasnost results will be addressed in a later Report.

Data Type	Locations	Test records	Devices	Carriers	Total Records
Raw data	1,200	108	2	4	1,036,800
Summary data (averaged)	1,200	40	2	4	384,000
Summary data (TCP only)	1,200	8	2	4	76,800
TCP Avg Up	1,200	1	2	4	9,600
TCP Avg Down	1,200	1	2	4	9,600

Note: details on the test sequence are included in Appendix A – Test Application.

2.1 Limitations of Mobile Testing

The following caveats should be taken into consideration for any mobile testing, including the testing summarized in this report. Radio communications is not an exact science. Planning and operating networks takes into account the probability of a user's location and the forecasting of aggregate demand for a cell site, both of which may vary, depending the time of day, location, and topography of the test location, network loading and congestion, and device hardware and software limitations. All of the carriers we tested are operating at frequencies low enough that weather conditions did not affect the results.

It should be kept in mind that these test results show end user experience at a specific time and specific location and do not necessarily indicate that an end user experience at one location will be similar to that of a nearby location, or at the same location at another time.

3 Analysis of Results

The following results are based on the three rounds of testing. The next round is scheduled for April 2014. As we complete subsequent rounds of testing, we will be able to draw more conclusions about each carrier's network in terms of changes in its coverage and performance over time. Details on the data analysis are included in Appendix G – Data Processing.

3.1 Comparing Actual Speed to Advertised Speed

The FCC has done an extensive, nation-wide study of wireline and satellite broadband providers to see how provider advertised speeds compared with actual delivered speeds at the home[7]. One of the reasons for doing the study was the FCC found in a separate study done in 2012 that 80 percent of consumers did not know what speed they purchased from their internet service provider. In order to determine how accurate advertised speeds were, the FCC installed measurement equipment at over 7,000 homes[8] and collected speed data over a period of a month.

---

[7] See the FCC's "Measuring Broadband America", February 2013.

[8] See technical appendix to FCC's Measuring Broadband America report at: http://data.fcc.gov/download/measuring-broadband-america/2013/Technical-Appendix-feb-2013.pdf

The FCC found that many ISPs either meet or exceed the speeds they advertise. Specifically, during peak periods, on average, DSL service providers delivered sustained download speeds that were 85 percent of advertised speeds, cable providers delivered 99 percent of advertised speeds, and fiber-to-the-home providers delivered 115 percent of advertised speeds.

We wanted to see how the speeds measured at each of the 1,200 test locations compared to the speeds advertised by each mobile carrier for that location. For this analysis, we used the lower threshold of each carrier's maximum advertised speed for each location based on December 31, 2012 availability data. For example, for a carrier that advertises 10-25 Mb/s for a location, we look to see if the measured speed for that carrier meets or exceeds 10 Mb/s.

As shown below, Verizon's results were nearly four times those of the other three carriers, so the Y-axis scale is different than that of the other carriers. The Y-axis is shown in terms of *yx*, as for example "3.0" would mean the measured speed was 3 times the advertised speed, or 300% of the lower threshold of the advertised speed. A value of 1.0 means that the measured speed was equal to the lower threshold of the advertised speed. For comparison sake, we added a blue line to show the 1.0 level. The dotted boxes represent the range of 1 standard error from the mean for each census area.

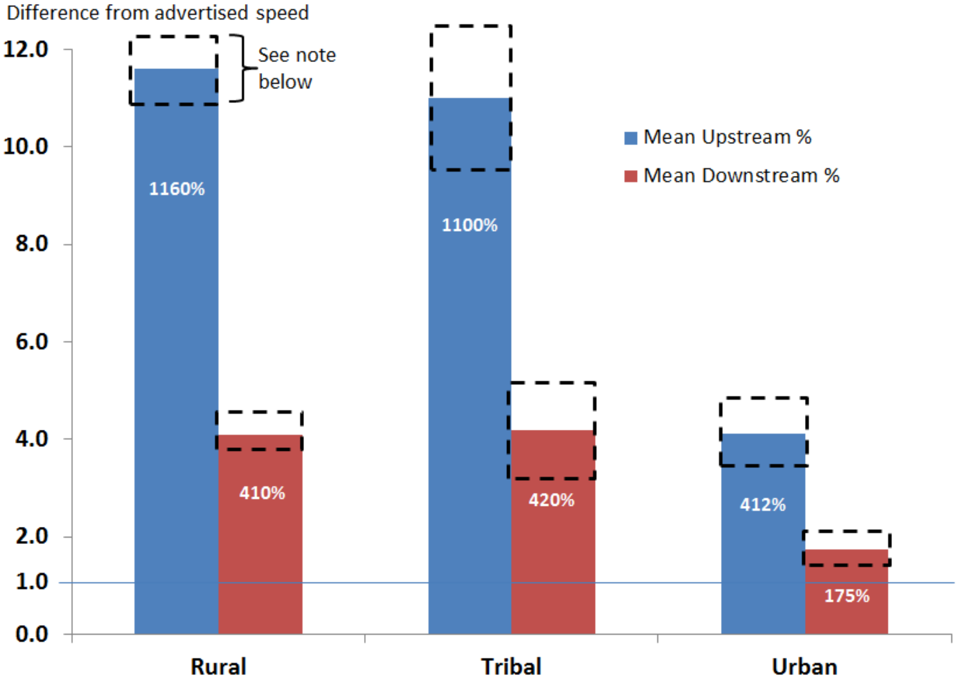

Verizon Measured Speed as Percentage of Lower End of Maximum Advertised Speed

1.0 is where measured speed = advertised speed

Note: Dotted boxes represent 1 standard error range from the mean

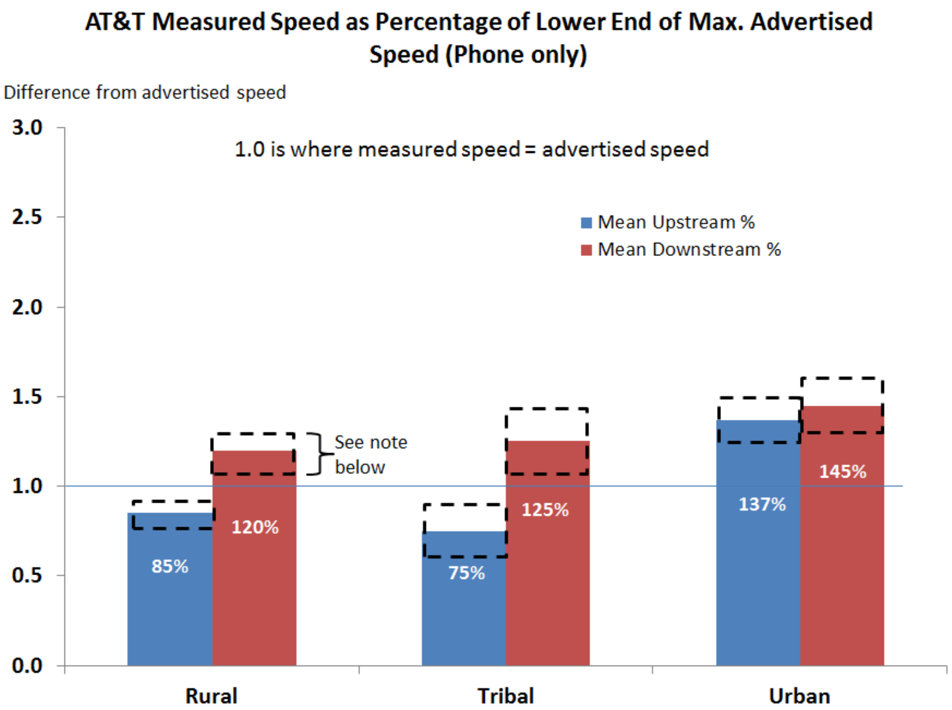

Note: Dotted boxes represent 1 standard error range from the mean

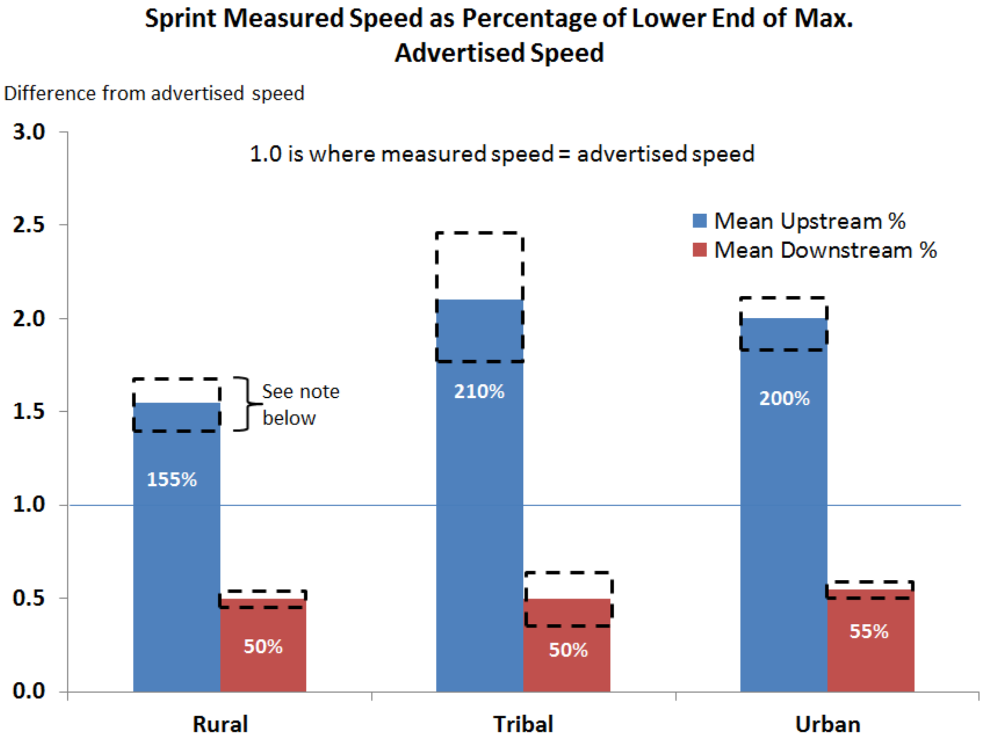

Note: Dotted boxes represent 1 standard error range from the mean

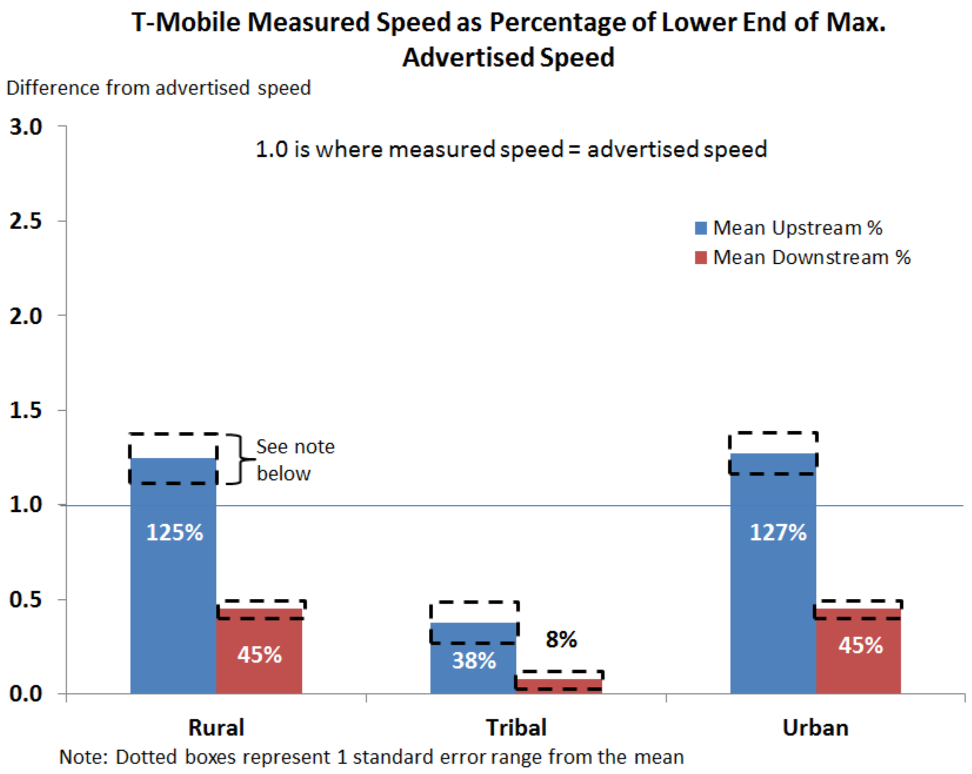

Findings

The only carrier to deliver more than the advertised speed for both upstream and downstream across all census areas was Verizon. There are several points to make about this finding. First, there is time gap between the service provider availability data (December 31, 2012) and the field testing (April, 2013). We observed from the mean and median calculations explored later in this paper that Verizon has been rapidly improving its upstream and downstream speeds – presumably because of LTE network expansion – and some of the rapid improvement in network speeds may have occurred between the time when we received availability data in December last year and performed the field test in April this year, hence the appearance of “over-delivering” upstream and downstream speeds.

The other finding is we observed Sprint and T-Mobile to have difficulty delivering advertised downstream speeds in all census areas. Part of this is due to our use of non-LTE devices for testing, but this only applies to areas where both carriers are advertising LTE. For non-LTE areas, we also observed a gap between advertised and actual speeds. We expect to see improvements in both Sprint and T-Mobile performance in the next round of field testing when we upgrade our devices to LTE.

3.2 Speed Coverage by Population

Using the third round results, we created an interpolation model that predicts downstream coverage and throughput (speed) for areas beyond the test locations themselves. The ArcGIS suite of software contains a series of interpolation tools, and for this analysis, we used the “kriging[9]” tool. We chose this tool over

---

[9] For an explanation of the kriging function and the mathematical model behind it, refer to Appendix H, Interpolation Model.

others in the suite because it interpolates a grid from a set of points, such as the mobile testing locations. We used the output of the interpolation model to estimate what percentage of the population is covered by each 1 Mb/s increment of downstream throughput for each carrier.

Using ArcGIS and Census 2010 population data, we calculated the percentage of population covered by the interpolated speed at every location in California. Using the CPUC "served" downstream threshold of 6 Mb/s, we looked at what percentage of the more than 37 million Californians were predicted to be covered at that speed.

Verizon Interpolated Downstream Speeds and Predicted Population Coverage (Apr. 2013)

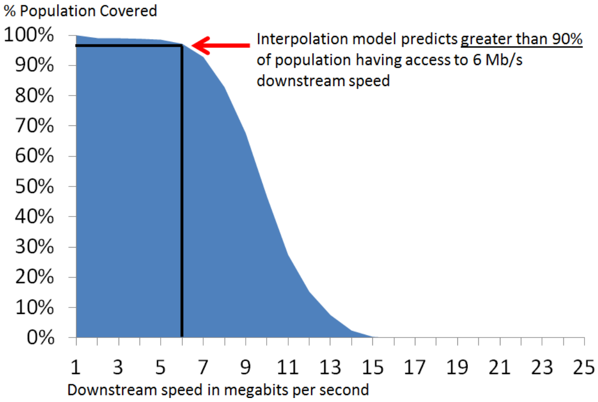

AT&T Interpolated Downstream Speeds and Predicted Population Coverage (Apr. 2013)

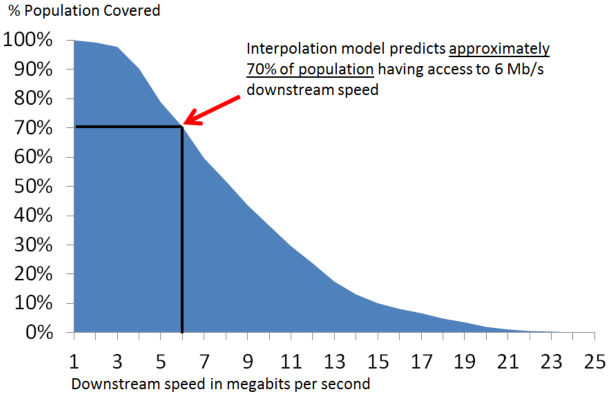

For Verizon and AT&T, both of whose LTE networks were used in this round of testing, the interpolation model predicted greater than 95% of California's population having access to up to 6 Mb/s downstream for Verizon and approximately 70% for AT&T. We also observed that while AT&T's maximum predicted speed was higher than Verizon's (over 24 Mb/s, covering less than 1% of the population), Verizon's model predicted more population coverage for the lower speeds.

For T-Mobile and Sprint, we observed no population having access to 6 Mb/s downstream based on the interpolation model prediction.

The analysis showed that all prediction models were significant predictors (p-value < 0.05) and differences between observed and predicted speeds were found to be statistically insignificant (p-value > 0.05). Please refer to Appendix H – Interpolation Model for calculated p-values for each carrier's predicted speeds.

Findings

Verizon's advertised coverage claims to cover more than 95% of the population with up to 6 Mb/s, is validated by the predicted coverage model. However, AT&T's advertised coverage also claims to cover more than 95% of the population with up to 6 Mb/s, however the predicted coverage is only 70% of the population.

3.3 Locations Meeting or Exceeding Combined 6 Mb/s / 1.5 Mb/s

One use of this data allows us to analyze to what extent mobile broadband can be a substitute for fixed wireline broadband. The threshold we used in this analysis is combined 6 Mb/s downstream and 1.5 Mb/s upstream, which is used by the FCC and the CPUC[10] as a lower bound of adequate service for today's uses. Locations with both up and down speeds below this benchmark may be eligible for a state CASF infrastructure grant.

At of the time of field testing, the only LTE equipment we used was from AT&T and Verizon. T-Mobile has recently upgraded many locations to LTE, and Sprint has also begun to upgrade to LTE. LTE results for those carriers are not included in this analysis but will be in the next. Moreover, in analyzing results from this round, we noticed that AT&T's data card was underperforming significantly compared to the smartphone (see the 22% gap in Urban_3 for AT&T, below). We later discovered that four of the eight data cards used for testing had not been provisioned to support LTE. For that reason, we separated out data card (in blue) and phone (in red) for this analysis.

The calculations below included only locations that fell within each carrier's coverage or roaming area[11]. The blue bars show the percent of locations served in each census area for each round of testing. Example: blue bar showing 6% for Rural_3 means that 6% of the rural locations (not total locations) within AT&T's coverage area had combined 6 Mb/s downstream and 1.5 Mb/s upstream speeds or better for the data card. We observed the opposite problem with Verizon and are investigating why there was a gap of as much as 25% of served locations between the data card and the phone.

Note that we did not include a graph for Sprint, because the percent of locations meeting or exceeding the 6 Mb/s downstream and 1.5 Mb/s upstream threshold was 1% or less.

---

[10] FCC 11-161, REPORT AND ORDER AND FURTHER NOTICE OF PROPOSED RULEMAKING, establishes a benchmark of 6 megabits per second downstream and 1.5 upstream for broadband deployments in later years of Connect America Fund Phase II. The same benchmark is used for certain California Advanced Services Fund grant applications.

[11] More discussion of what defines the coverage and roaming areas may be found in Section 3.8.

Chart 1: % Locations where Downstream ≥ 6 Mb/s and Upstream ≥ 1.5 Mb/s

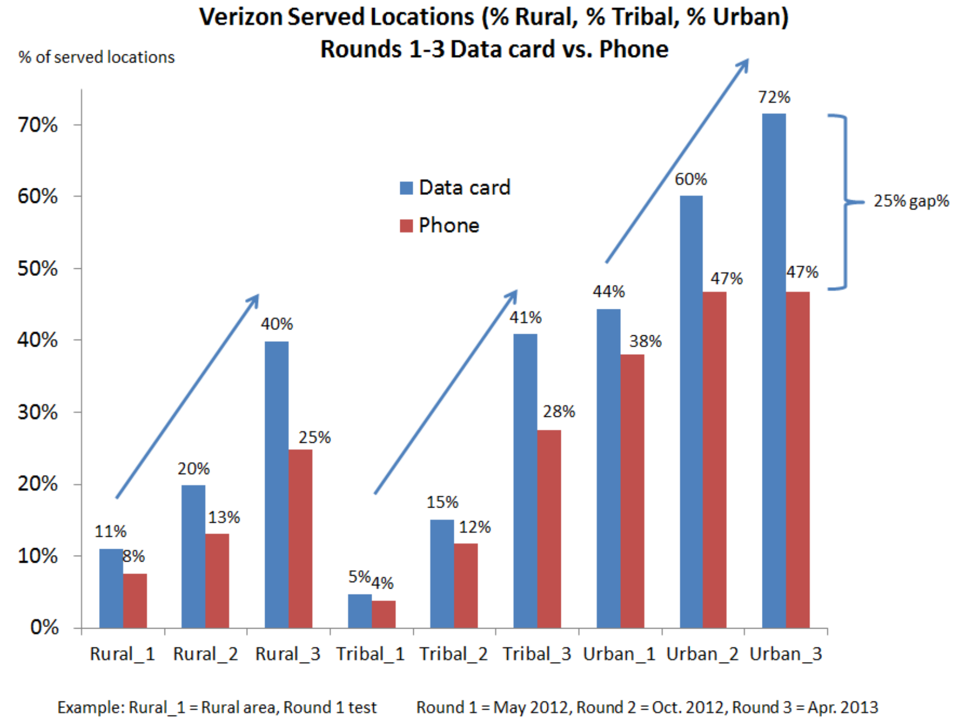

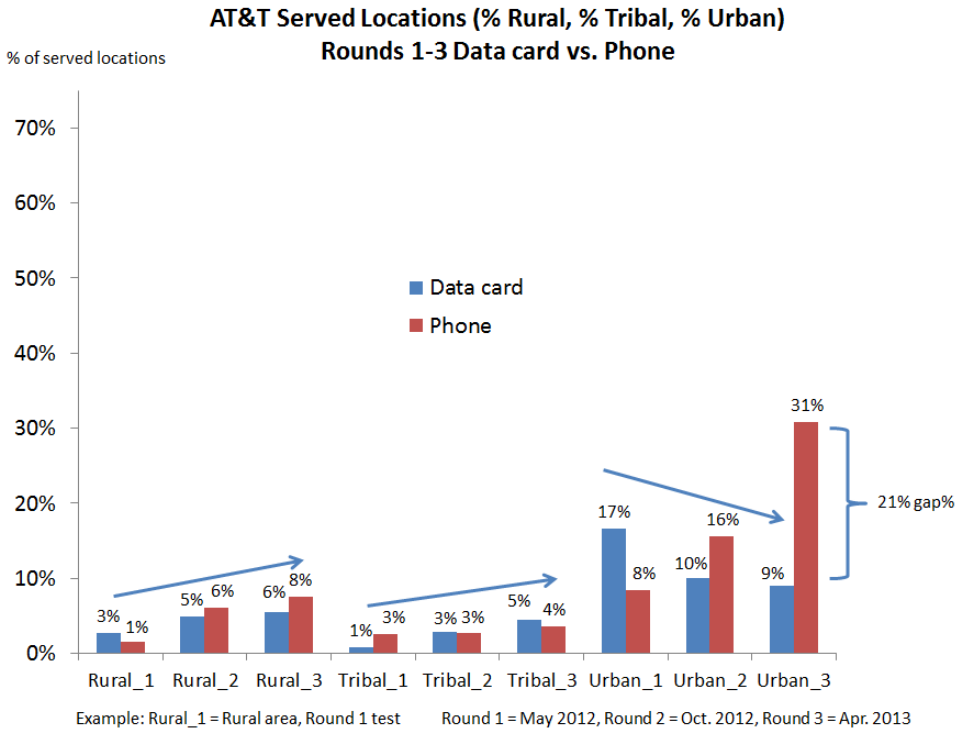

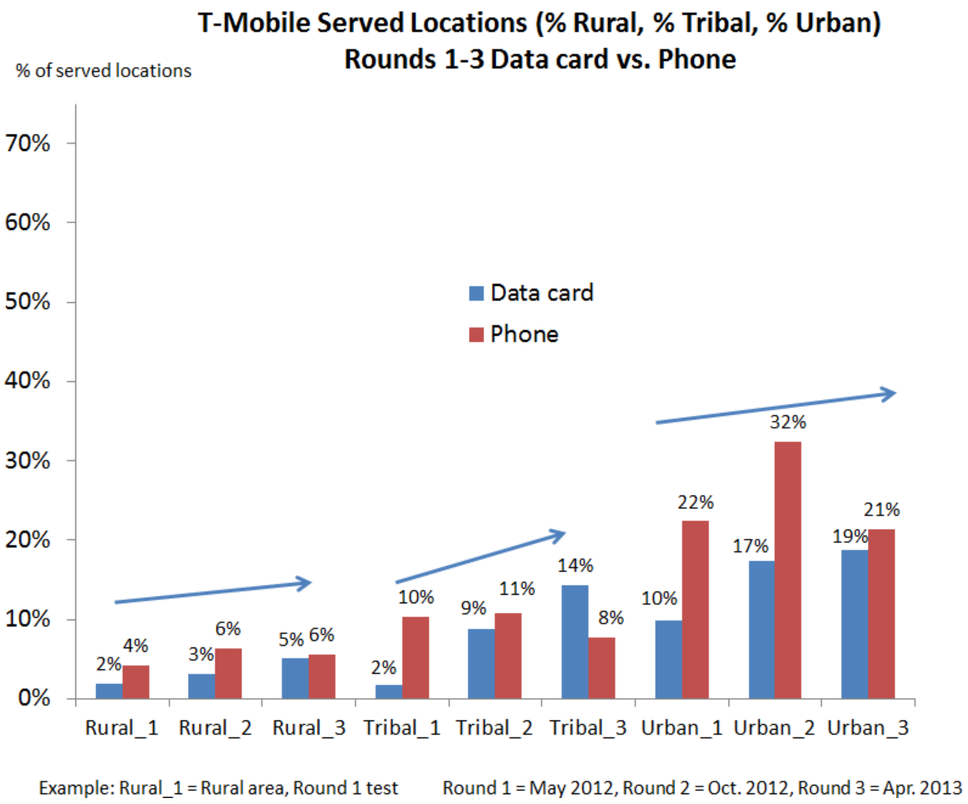

Findings

Of the four carriers, Verizon showed the greatest improvement in providing combined 6 Mb/s downstream and 1.5 Mb/s upstream or greater in all census areas from rounds 1-3, reaching as high as 72% of all urban locations and 40% of rural locations in round 3. We observed a slight improvement with AT&T, and T-Mobile’s results were effectively unchanged across all three rounds. Particularly noteworthy is Verizon coverage of 40 percent or higher in rural and tribal locations. It will be interesting to track what percentage Verizon reaches in tribal and rural in the 4th round. Can they reach 50 percent? They have more than doubled their covered locations, rural and tribal areas, between rounds 2 and 3.

3.4 Mean Downstream for Rural, Tribal, and Urban Locations

When comparing average throughput by location type, we found it was important to look at both the mean (statistical average) and median (middle point) to see how the carriers differed. As we will see below, the overall mean for all carriers is higher than the overall median. The size of the gap between the two is driven by the degree of outliers.

Notes on the graphs: The definitions for “urban” and “rural” in this analysis relies on the Census 2010 definition[12] rather than the method used for location selection referenced in Appendix D – Test Routes. The mean and median calculations included only locations that fell within each carrier’s coverage and roaming areas, as described in Section 3.8. The arrows embedded within each bar indicate the percentage change in the second round of testing compared to the first. The overall mean and median values are not

---

[12] See footnote 1 in Section 1

weighted; they are simply the overall mean and overall median for all of the locations in each carrier's coverage area.

The latency values in the charts below are averages for phone and data card, as well as for east and west servers. The purpose of showing latency with mean downstream speeds is to give a more complete picture of the true user experience. In some cases, latency improved significantly from the first round of testing to the third. In other cases, it remained effectively unchanged. The key point here is the question of how high is overall latency. For all census areas, AT&T's latency fell below 200 milliseconds by the third round of testing. Results for Verizon are similar. Both carriers have large LTE networks, and one of the benefits of LTE is low latency. Sprint and T-Mobile, despite improvements in some census areas, still show average latency by round 3 to be between 300-600 milliseconds. We expect those latency averages to drop in the next round of testing after we upgrade to LTE devices for those carriers.

As mentioned in the previous section, we observed a significant gap in AT&T's data card performance (3,000 Kb/s, or 3 Mb/s for the Urban_3 mean value), and this was attributable to half of the data cards being incorrectly provisioned.

Chart 3: Mean Downstream Throughput by Census Area, Device, and Average Latency

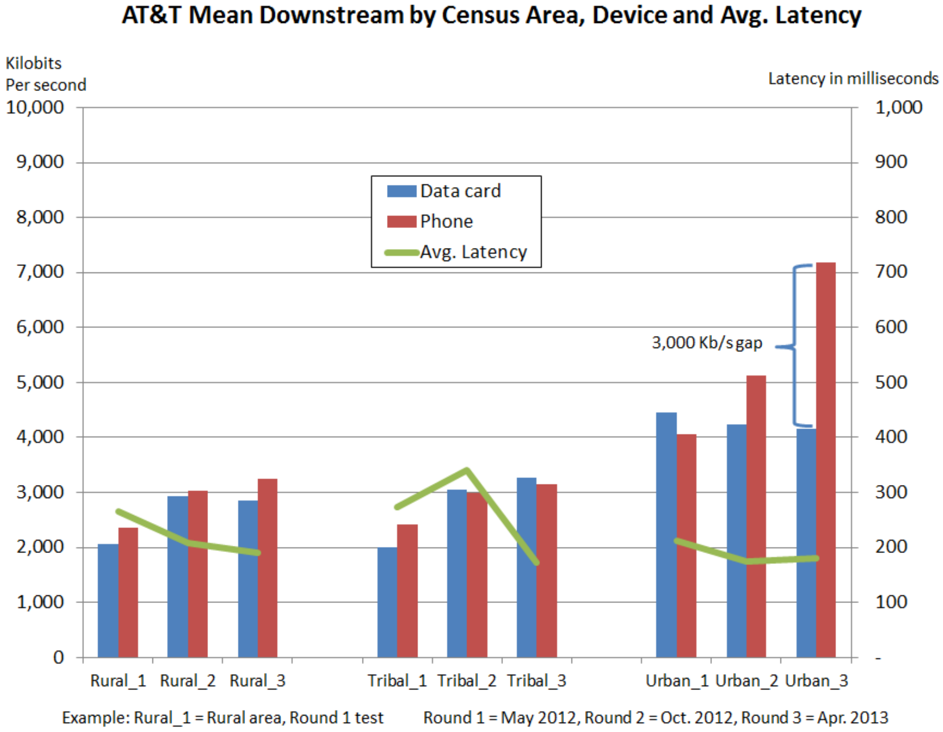

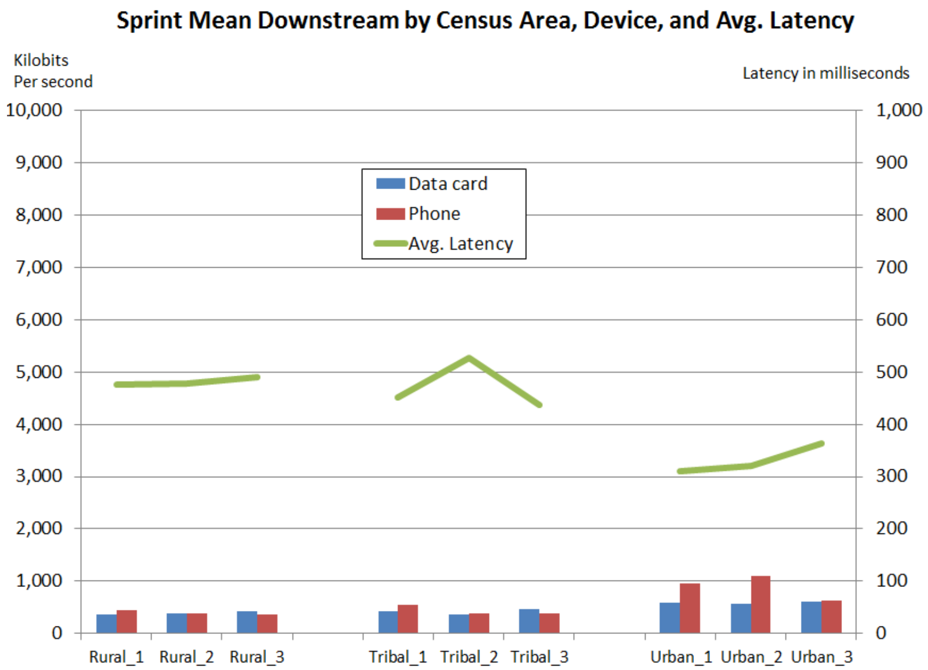

Example: Rural_1 = Rural area, Round 1 test Round 1 = May 2012, Round 2 = Oct. 2012, Round 3 = Apr. 2013

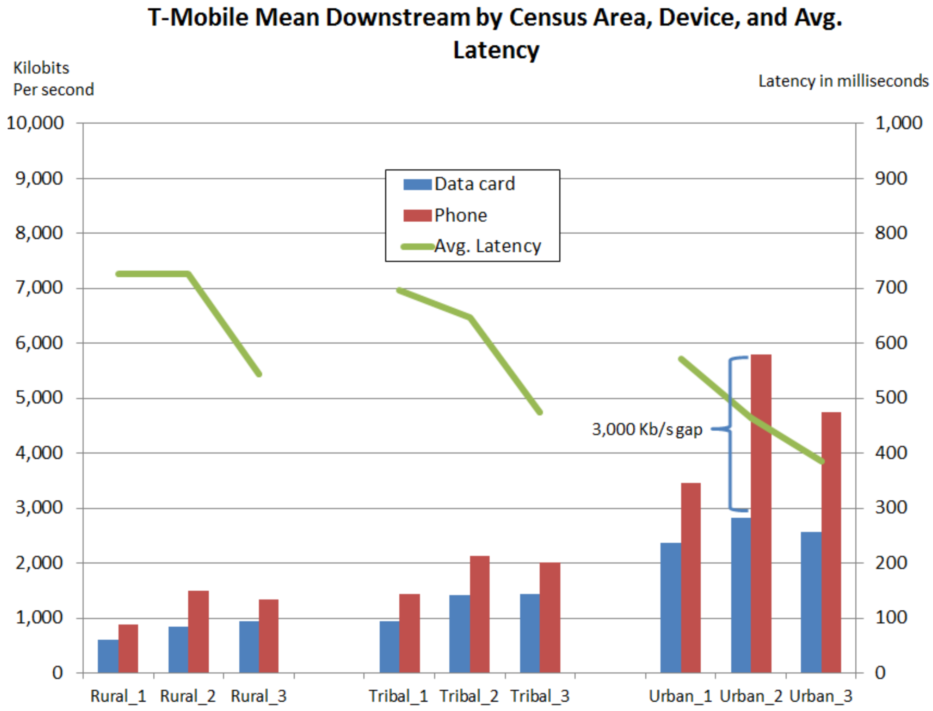

Example: Rural_1 = Rural area, Round 1 test Round 1 = May 2012, Round 2 = Oct. 2012, Round 3 = Apr. 2013

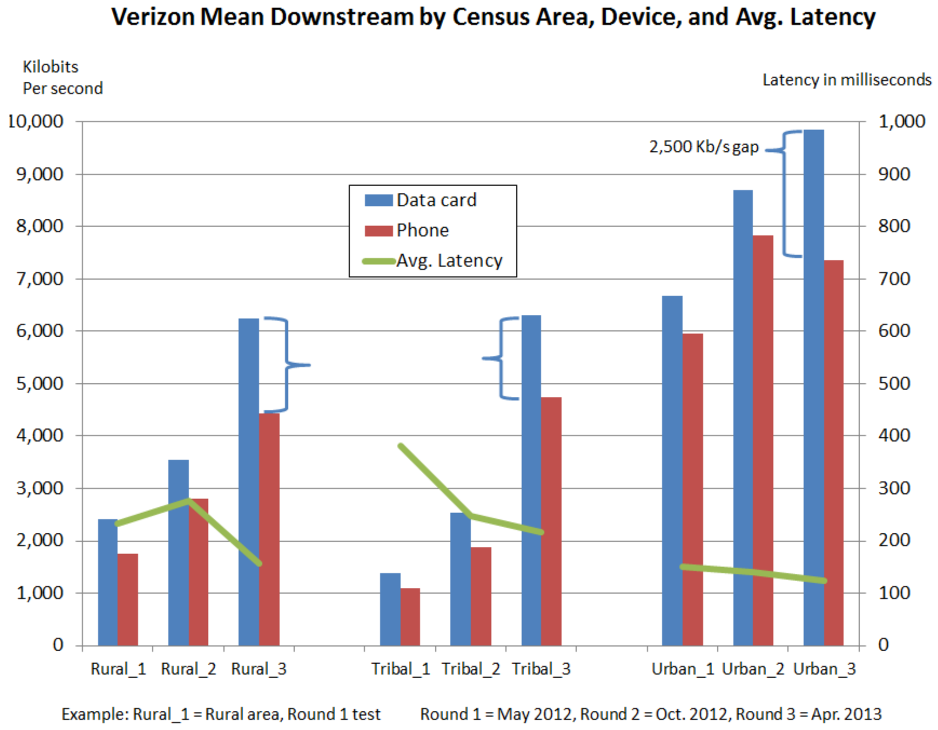

Findings

With the exception of Verizon's Round 3 phone results, we noticed a significant, steady improvement in downstream speeds and overall improvement in latency for both phones and data cards in all census designations. We suspect the discrepancy in phone results for Round 3 is attributable to limitations with the phone's technology, and we plan to address that in the next round of testing by upgrading devices for all carriers. For AT&T, improvement in downstream speed was most visible in the urban areas. As mentioned earlier, we discovered a problem with some of the AT&T data cards, and effect of this problem is shown in the gap between data card and phone results for Round 3 urban. For T-Mobile and Sprint, downstream speed improvements were negligible. The only clear improvement seen in the two carriers was the reduction in latency for T-Mobile.

3.5 Median Throughput Comparisons

For simplicity sake, we compared overall median upstream and downstream values for all four carriers together. The graphs below show the median values of the mean upstream and downstream throughputs for both data card and phone results for each round of testing across all census areas.

Median Downstream Throughput by Location Type

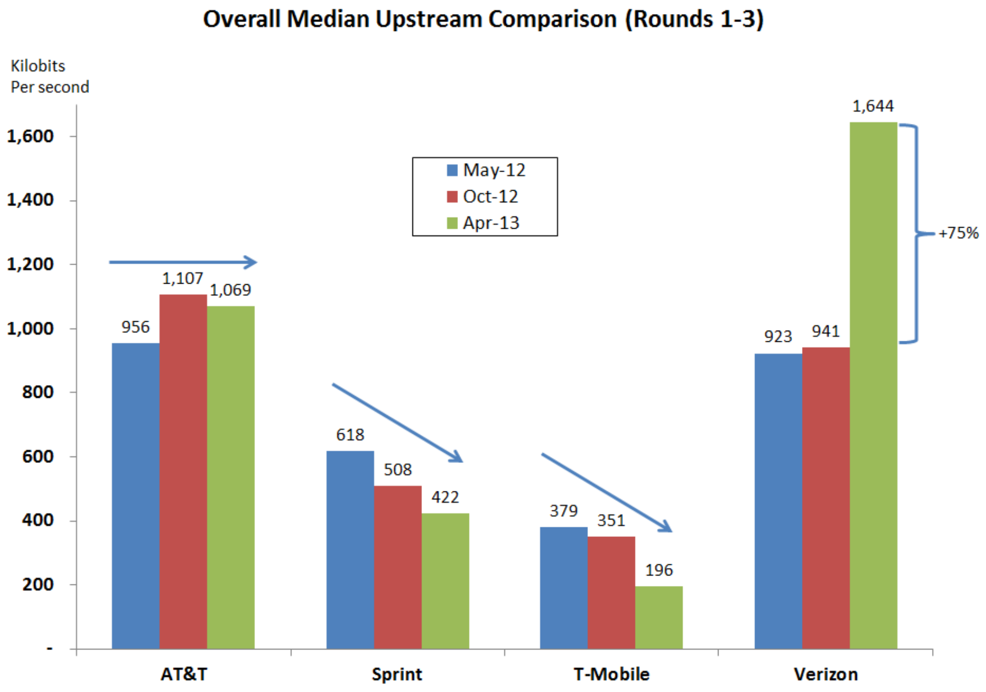

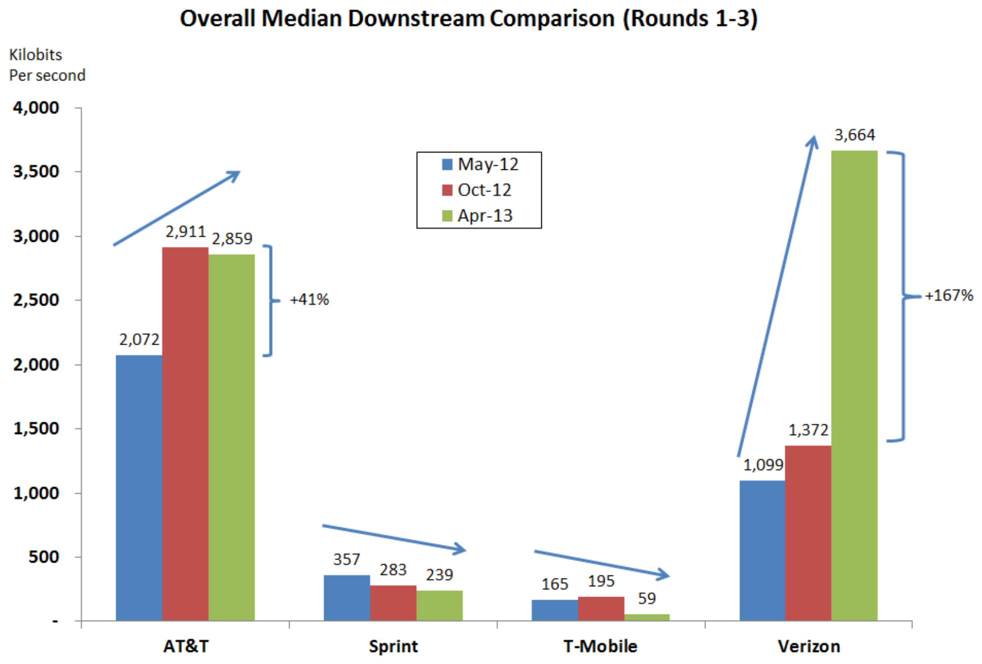

Findings

In the third round of testing, Verizon's overall median for both upstream and downstream surpassed that of AT&T. We attribute that improvement to Verizon's continued deployment of LTE network in rural and tribal areas. We observed a 41% improvement in AT&T's median downstream throughput from May 2012 to October 2012, and the main factor there was upgrading to LTE devices, however, we did not observe a significant improvement in median upstream throughput. Both Sprint and T-Mobile median upstream and downstream values appeared to decrease – more in T-Mobile's case more than Sprint's. This may be attributable to T-Mobile launching the iPhone and the associated increase in traffic congestion with the iPhone's introduction. Verizon's solid performance in median throughput shows the improvement is not guided by outliers. It is an improvement that we can see across the board.

3.6 Mean Opinion Score Estimates

MOS (mean opinion score) is a metric used to subjectively measure voice quality for calls made using voice over internet protocol (VoIP). Normally, a tester assigns a number, from 1 to 5 to a call. However, it is possible to approximate the mean opinion score using three factors that significantly affect call quality: latency, packet loss, and jitter. Packet loss is the percentage of packets that never made it from the caller to the computer you're calling and back again. If we send out 100 packets and only receive 97 back (3 failed to return), then we have 3% packet loss. Latency is the average (mean) time it takes a packet to get from your computer to the computer you're calling and back again, measured in milliseconds. Jitter is a measurement of how much latency changes during a call. Low jitter indicates a good connection, whereas a high jitter indicates network congestion.

Using latency, packet loss, and jitter, we can calculate something called an "R" value, which can easily be converted to a mean opinion score. One MOS point is equivalent to approximately 20 R-value points, although the relationship is not linear. Below is a table showing the relationship between R-value, call quality, and MOS. More information on how we calculated the R-value and MOS is included in Appendix I.

R-value	User Satisfaction	MOS
90 – 100	Very satisfied	4.3 – 5.0
80 – 90	Satisfied	4.0 – 4.3
70 – 80	Some users satisfied	3.6 – 4.0
60 – 70	Many users dissatisfied	3.1 – 3.6
50 – 60	Nearly all users dissatisfied	2.6 – 3.1
Less than 50	Not recommended	1.0 – 2.6

cutoff

We calculated the MOS for each test location, device, and carrier combination and determined how many locations met or exceeded a MOS of 4.0, which corresponds with voice call quality of "Satisfied" to "Very Satisfied." A MOS of 4.0 or greater is one way of quickly assessing an operator's VoIP readiness.

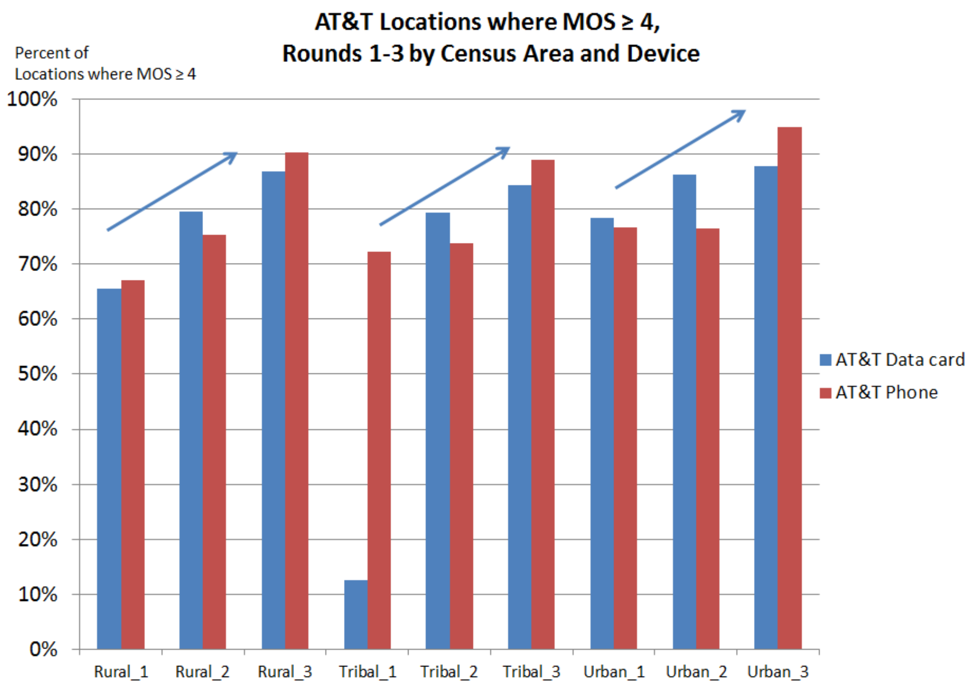

AT&T Locations where MOS ≥ 4, Rounds 1-3 by Census Area and Device

Example: Rural_1 = Rural area, Round 1 test   Round 1 = May 2012, Round 2 = Oct. 2012, Round 3 = Apr. 2013

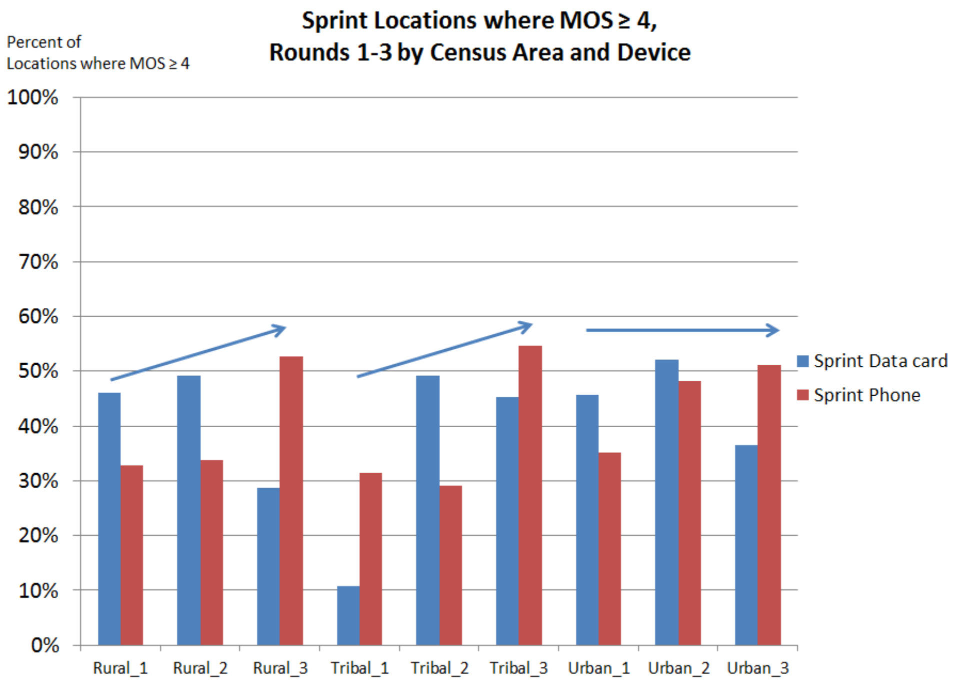

Sprint Locations where MOS ≥ 4, Rounds 1-3 by Census Area and Device

Example: Rural_1 = Rural area, Round 1 test   Round 1 = May 2012, Round 2 = Oct. 2012, Round 3 = Apr. 2013

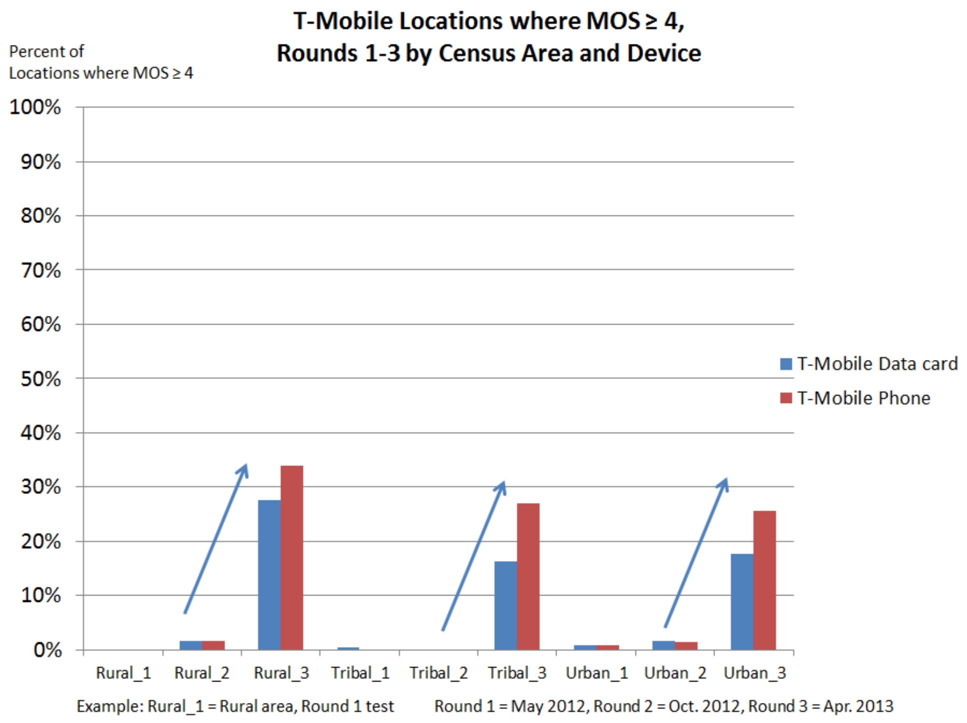
T-Mobile Locations where MOS ≥ 4,
Rounds 1-3 by Census Area and Device
Percent of Locations where MOS ≥ 4
100%
90%
80%
70%
60%
50%
40%
30%
20%
10%
0%
Rural_1
Rural_2
Rural_3
Tribal_1
Tribal_2
Tribal_3
Urban_1
Urban_2
Urban_3
T-Mobile Data card
T-Mobile Phone
Example: Rural_1 = Rural area, Round 1 test
Round 1 = May 2012, Round 2 = Oct. 2012, Round 3 = Apr. 2013

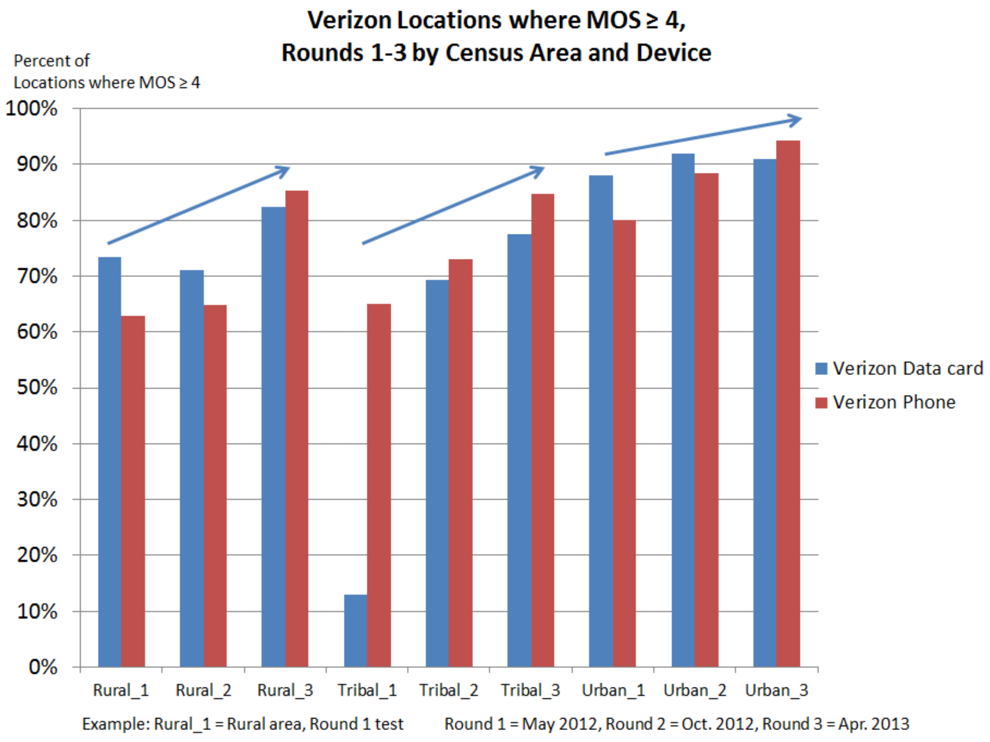
Verizon Locations where MOS ≥ 4,
Rounds 1-3 by Census Area and Device
Percent of Locations where MOS ≥ 4
100%
90%
80%
70%
60%
50%
40%
30%
20%
10%
0%
Rural_1
Rural_2
Rural_3
Tribal_1
Tribal_2
Tribal_3
Urban_1
Urban_2
Urban_3
Verizon Data card
Verizon Phone
Example: Rural_1 = Rural area, Round 1 test
Round 1 = May 2012, Round 2 = Oct. 2012, Round 3 = Apr. 2013

Findings

T-Mobile, though consistently having the lowest percentage of MOS 4 provisioning, dramatically improved their performance in the third round. T-Mobile's latency scores have dropped considerably in the third round and this may have helped increase their MOS scores because latency goes into the calculation of MOS Scores. The other providers also show steady improvement in their MOS scores. It is interesting to note that while Sprint's latency numbers are comparable to T-Mobile's, their number of locations achieving a MOS score higher than four is considerably higher. The only provider area which did not record an increase in MOS was Sprint in urban areas, which has remained level. Additionally, Sprint urban is the only provider census area that recorded an increase in latency in the third round. Aside from that Sprint urban anomaly, all providers in all areas show MOS improvement.

3.7 Tribal Areas –Locations and Number of Carriers

Tribal locations were generally tested at the entrance or near the entrance of the property. In some cases, we were able to obtain permission to test on tribal property. In most cases, the user experience was not necessarily the same as what might be found inside the tribal areas. Compared to the first round of testing, we noticed a shifting upward in the number of tribal locations meeting the 6/1.5 Mb/s threshold and away from the middle range (shown in green).

Complete test results for each of the 128 Tribal locations will be included in a forthcoming report on tribal broadband in California.

Chart 14: Tribal Locations Meeting Combined ≥ 6 Mb/s Downstream and ≥ 1.5 Mb/s Upstream Threshold (n=128)

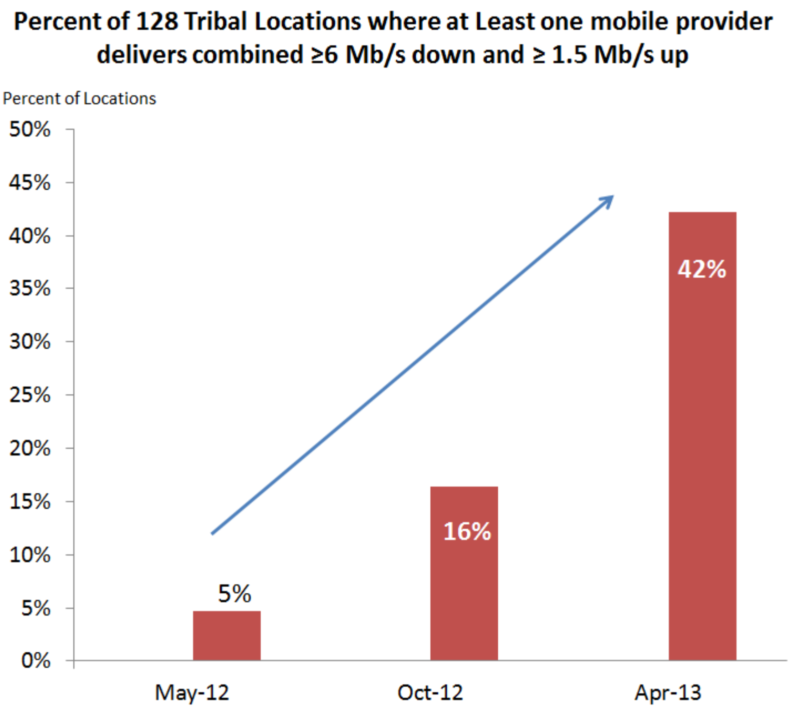

Chart 15: Number of Mobile Providers at Tribal Locations (n = 128)

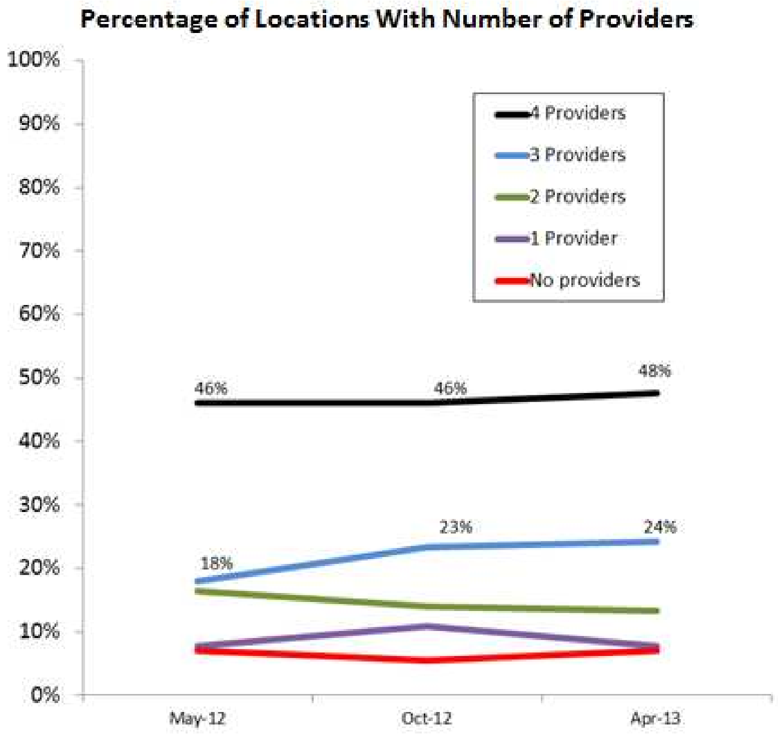

Findings

One big change is the percentage of tribal locations served went from 16 percent in round 2 to 42 percent in round 3. This huge jump is largely due to Verizon's more than doubling their speed performance on tribal lands from rounds two and three. Using the data card, Verizon itself served 41 percent of tribal locations in Round 3. While other providers have made some improvements to their tribal performance, Verizon has made dramatic improvements to their tribal performance. This has resulted in a big change in the number of tribal locations served and likely coincides with Verizon's introduction of LTE in tribal areas.

4 Conclusion

Based on the findings in this paper, we have come to the following conclusions:

- The only carrier to deliver greater than advertised speeds across all census areas was Verizon, but we suspect that is due to the time gap between the availability data and the date when we performed the field tests. Sprint and T-Mobile appeared to under-deliver downstream speeds in all census areas, while and AT&T under-delivered upstream speeds in rural and tribal locations.

- Verizon continued to make dramatic improvements in upstream and downstream speeds in all census areas, particularly in tribal and rural areas, which are areas generally characterized by substandard broadband service. Improvements for the other three carriers were more incremental.

- All providers continue to make improvements in the number of locations qualifying with a Mean Opinion Score of 4.0 or greater, which is one indicator of VoIP readiness. While the number of T-Mobile locations achieving a MOS score of 4.0 was extraordinarily low, we saw a dramatic increase in T-Mobile's performance from October 2012 to April 2013.

Appendix A – Test Application

TESTING APPLICATION IMPLEMENTATION: IPERF FOR TCP AND UDP MEASUREMENTS

The user interface is written in Java and runs on both the Android and Windows operating systems.

IPerf is a software tool to measure the network throughput and performance. It is developed using a client-server model as in Figure 1. The client software in a client device requests an either TCP or UDP data connection to the server side software. After the connection is established, the tool measures the throughput between the two ends. In the project, we set up the iPerf server software on two Amazon Linux machines at the East and West coasts, respectively.

Figure 1

In the project, the iPerf tool measures TCP upload bandwidth, TCP download bandwidth, UDP jitter, and UDP datagram loss rate. For the TCP upload, the client software makes a connection to the server side software and sends data streams to the server side for 10 seconds. After that, the server side software sends data streams from the server to the client for download bandwidth. Note that the original iPerf tool uses two separate connections for the upload and download measurements, respectively. But because many client devices and network operators do not allow for the server side software to make a new connection with the client device, we keep the connection used for the upload measurement and use it for the download measurement as well. By this technique, we can avoid the firewall blocking at the client device. As other measurement parameters, the tool uses 64.0 Kbytes window size (-w 64k), which is the amount of data that can be buffered during a connection without a validation from the receiver, and executes four threads concurrently (-P 4) at both sides, which can increase the data volumes to be exchanged between the client and server.

For the UDP jitter and datagram loss rate measurement, the iPerf sends data to the client side for either one or five seconds. Note that there's no data streams from the server to the client in the UDP measurement. As other parameters, we use 220K buffer length (-l 220) and 88K bits/sec bandwidth (-b 88k) to send data. The following shows all measurements in a single testing:

- TCP upload (10 seconds) and download (10 seconds) measurement to the West server (twice)
- TCP upload (10 seconds) and download (10 seconds) measurement to the East server (twice)
- UDP jitter and datagram loss (1 second) to the West server (three times)
- UDP jitter and datagram loss (1 second) to the East server (three times)

- UDP jitter and datagram loss (5 seconds) to the West server (one time)
- UDP jitter and datagram loss (5 seconds) to the East server (one time)

Test	Explanation
Minimum, maximum, and average packet Round Trip Time (RTT); packet loss rate during the RTT	To understand how long it takes to send & receive data and amount of data that can be lost during transmission
User Datagram Protocol (UDP) jitter & loss	Both of these measurements assess the quality of UDP. UDP is the network protocol used for streaming media such as video and voice over internet protocol (VoIP). Jitter measures the degree to which the UDP signal becomes distorted during transmission. Loss is the amount of message that gets lost during transmission.
Transmission Control Protocol (TCP), upstream and downstream throughputs	TCP is the network protocol used for E-mail. Expressed in kilobits per second, these measurements tell how quickly the user can send (upstream) and receive (downstream) TCP messages.

Appendix B – Test Equipment

EQUIPMENT USED FOR MOBILE TESTING

Equipment	Type	Technology	Quantity
Asus Netbook	Netbook	Not Applicable	8
BU-353 GPS	GPS receiver	Not Applicable	8
AT&T HTC ONE X Android	Smart phone	LTE, HSPA+, UMTS	8
AT&T USB Momentum	Data card	LTE, HSPA+, UMTS	8
Sprint HTC EVO 4G Android	Smart phone	EVO, 1XRTT	8
Sprint Sierra Wireless 250U USB	Data card	EVO, 1XRTT	8
T-Mobile Samsung Galaxy SII Android	Smart phone	HSPA+, UMTS	8
T-Mobile Jet 2.0 4G Laptop Stick	Data card	HSPA+, UMTS	8
Verizon Droid Charge 4G LTE Android	Smart phone	LTE, EVO, 1XRTT	8
Verizon Pantech Data card UML290	Data card	LTE, EVO, 1XRTT	8

Appendix C - Test Locations

The image below shows the combined maximum advertised downstream throughput for all mobile carriers submitting broadband data to the CPUC (color coded by speeds as “Maximum Advertised Downstream Speed.”). The test locations are shown in color coded points overlaid on top (color coded by location type as “Round 2 Field Testing Sample Locations.”) These same locations were used in Round 1.

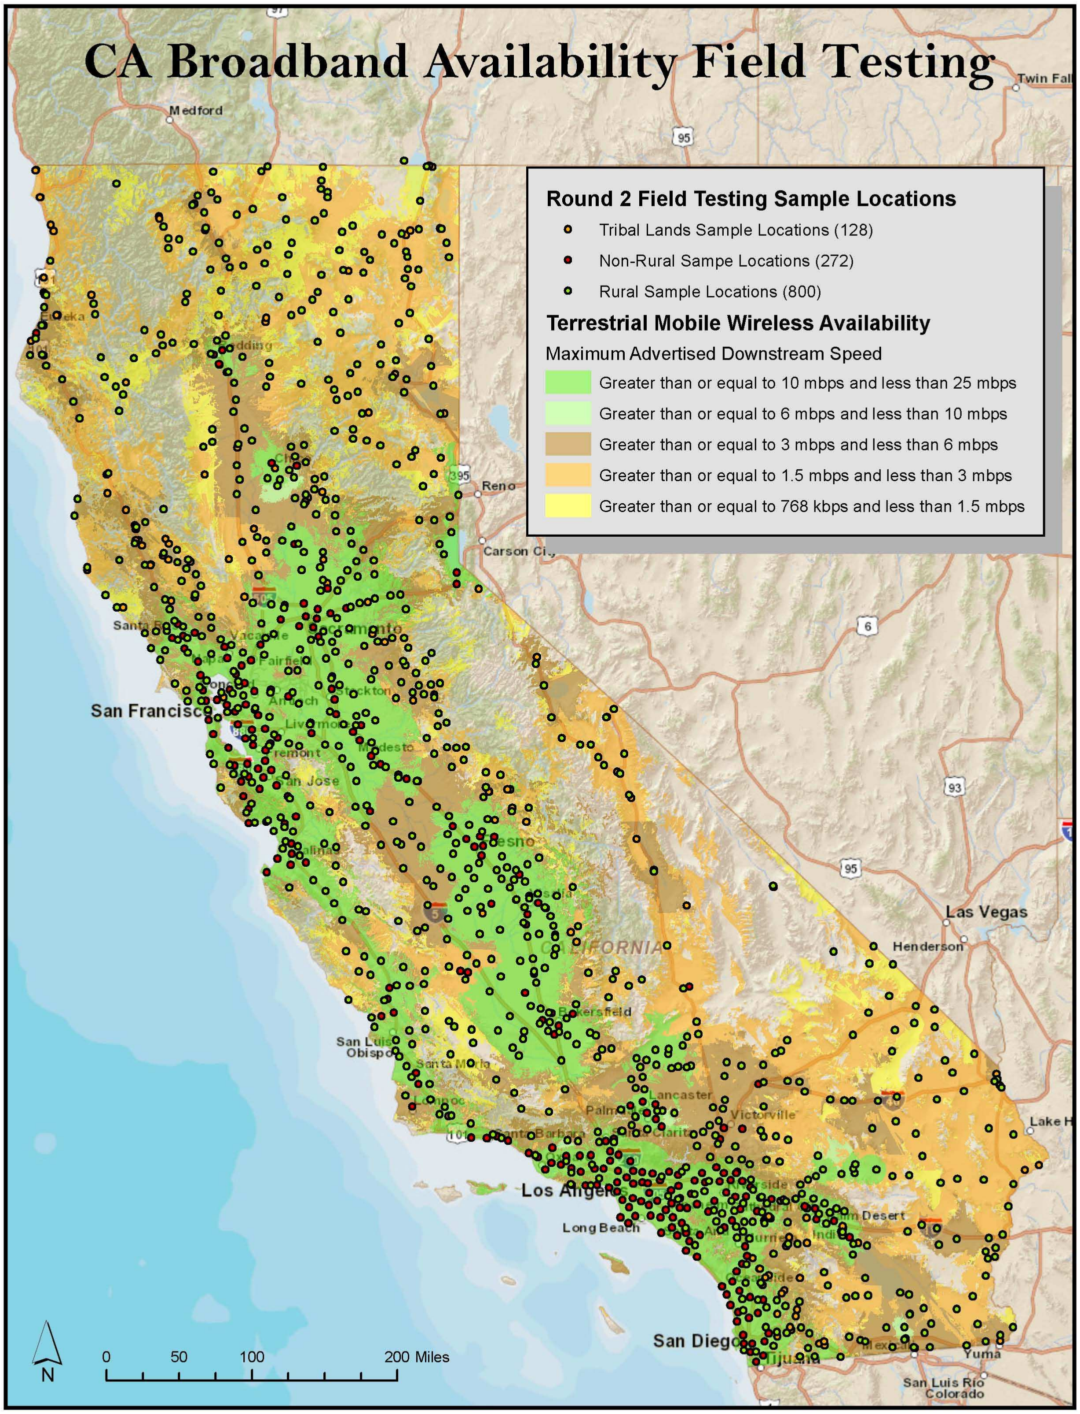

Appendix D – Test Routes

Routes were generated using ESRI's Network Analyst extension. Test locations were divided between Chico and Northridge testers based on closest location, then adjusted to make distribution equal between the two groups of testers (600 points to Chico, 600 to Northridge). Then, locations were divided into logical zones based on regional geographic considerations. Routes were designed to minimize drive time given the constraints of an 8-hour work day, idle time due to testing, a lunch break, and returning home for weekends. Routes from each zone were divided between each driver in an effort to equalize total drive time for each driver.

TRIBAL SAMPLING

Tribal Lands sampling locations were defined as a point along the California Road Network, as classified in the 2010 census road inventory, which fall within Tribal Lands. One sample location was randomly generated for each Tribal Land containing features from the California Road Network using spatial analysis software. Each point represents a location currently known to have mobile wireless broadband coverage. A total of 128 sampling locations were generated for Tribal Lands.

NON-RURAL SAMPLING:

Non-Rural sampling locations were defined as a point along the California Road Network, excluding MTFCC Class Code S1400, as classified in the 2010 census road inventory. All points in this category fell within Non-Rural Lands, and outside any Tribal Lands. All sampling locations were randomly generated at a minimum distance of five miles using spatial analysis software and represented a location known to have mobile wireless broadband coverage. A total of 136 sampling locations were generated for Non-Rural Lands.

RURAL SAMPLING:

Rural sampling locations were defined as a point along the California Road Network as classified in the 2010 census road inventory. Rural points were first allocated to 2010 Census Designated Places (2010 CDP) that were outside any Non-Rural and Tribal Lands and were randomly generated at a minimum distance of five miles using spatial analysis software. Remaining sampling locations were allocated to Rural areas not within any other lands (Tribal, Non-Rural, and CDPs) and were randomly generated at a minimum distance of ten miles using spatial analysis software in areas where no sampling locations existed from the initial Rural distribution. All sampling locations represented a location known to have mobile wireless broadband coverage. A total of 936 sampling locations were generated for Rural Lands.

Notes: 1. Minimum distances were determined by the largest distance (tested in five mile increments) the software allowed for the amount of points within the sampling area being calculated. 2. California Roads Network: Sampling locations were controlled to highway ramp (S1630), secondary roads (S1200), and local neighborhood road/rural road/city street (S1400), as classified in the 2010 census road inventory. 3. All road classes were weighted equally in the distribution process.

Appendix E - Operator Technologies & Frequencies

Carrier	2G MHz	3G MHz	4G MHz	3G Technology	4G Technology
AT&T	850, 1900	850, 1900	700, 1700/2100	GSM/HSPA+	LTE (FDD)
Sprint	1900	1900	2500, 1900 PCS, 800	CDMA/EV-DO	WiMAX (Clearwire) LTE (Sprint, FDD; Clearwire, TDD)
T-Mobile USA	1900	1700/2100, 1900	1700/2100	GSM/HSPA+	LTE (FDD) - launch in 2013
Verizon	800, 1900	800, 1900	700	CDMA/EV-DO	LTE (FDD)

Notes:

1. All frequencies listed are in megahertz (MHz)
2. Paired frequencies are listed with "/" where the first number represents the uplink frequency, and the second the downlink frequency

Acronyms

GSM	Global System (for) Mobile (Communications)
HSPA+	High Speed Packet Access Plus
CDMA	Code Division Multiple Access
LTE	Long Term Evolution
FDD	Frequency Division Duplex
TDD	Time Division Duplex

LTE DEPLOYMENT STATUS IN CALIFORNIA

During the second round of testing, only two of the four major carriers had deployed LTE in California: Verizon and AT&T. Sprint and T-Mobile have announced they plan to launch LTE, but no dates are set yet for California.

VERIZON LTE Markets

Bakersfield, Chico/Oroville, Eureka, Fresno, Los Angeles, Merced, Modesto, Oakland, Redding, Sacramento, Salinas/Monterey/Seaside, San Diego, San Francisco, San Jose, San Luis Obispo, Santa Barbara/Santa Maria, Stockton, Visalia/Porterville/Hanford, and Yuba City/Marysville.

AT&T LTE Markets

Bakersfield, Los Angeles, Modesto, Oakland, Sacramento, San Diego, San Francisco, and San Jose.

SPRINT LTE Markets

Sprint's LTE network had not been launched in California in time for this round of testing.

Appendix F – Data Record Format

DATA FIELD DESCRIPTION		
Column Name	**Example(s)**	**Description**
Tester	3	ID number of a tester who conducted the testing
Location ID	1020	Location number assigned by GIC
Date	5/30/2012	Test starting date. We use the "05/30/2012" date format.
Time	15:48:08	Test starting date.
Provider	Verizon, AT&T, T-Mobile, Sprint	Network provider.
Operator	Verizon, AT&T, T-Mobile, Sprint	Network operator.
Network	EHRPD, LTE, UMTS, HSDPA, etc.	A specific network type. At the moment, the information is available at the phone testing only. Netbook testing is always NA.
Latitude	37.323429	
Longitude	-122.036079	
Device ID	WBBDTest1 (netbook), 99000024385563 (phone)	A unique number of a testing device.
Client Type	Phone or Netbook	
ePktMin	154.062	Minimum RTT (Round Trip Time) of packets to the East server (in ms)
ePktMax	270.189	Maximum RTT (Round Trip Time) of packets to the East server (in ms)
ePktAvg	184.525	Average RTT (Round Trip Time) of packets to the East server (in ms)
ePktLoss	0	Packet loss rate (%) during the RTT test to the East server.
wPktMin	83.702	Minimum RTT (Round Trip Time) of packets to the West server (in ms)
wPktMax	218.218	Maximum RTT (Round Trip Time) of packets to the West server (in ms)
wPktAvg	118.374	Average RTT (Round Trip Time) of packets to the West server (in ms)
wPktLoss	0	Packet loss rate (%) during the RTT test to the West server.
eUDP Jitter1	14.606	Jitter for the first UDP measurement to the East server (in ms). There are total four UDP jitter measurements to the East server such as eUDP Jitter2, eUDP Jitter 3, and eUDP Jitter 4.
eUDP Loss1	0	UDP datagram loss rate for the first UDP measurement to the East server (in %). There are total four UDP data gram loss rate for each measurement such as eUDP Loss 2, eUDP Loss 3, and eUDP Loss 4.
eUDP Period1	1	UDP test period for the first UDP measurement to the East server (either 1 sec or 5 sec). There are total four UDP measurements to the East server such as eUDP Period 2, eUDP Period 3, and eUDP Period 4.

wUDP Jitter1	17.412	Jitter for the first UDP measurement to the West server (in ms). There are total four UDP jitter measurements to the West server such as wUDP Jitter2, wUDP Jitter 3, and wUDP Jitter 4.
wUDP Loss1	0	UDP datagram loss rate for the first UDP measurement to the West server (in %). There are total four UDP data gram loss rate for each measurement such as wUDP Loss 2, wUDP Loss 3, and wUDP Loss 4.
wUDP Period1	1	UDP test period for the first UDP measurement to the West server (either 1 sec or 5 sec). There are total four UDP measurements to the West server such as wUDP Period 2, wUDP Period 3, and wUDP Period 4.
wTCP_UP1	478.9	TCP upload speed to the West server at the first measurement. (in kbps). Technically, our tool has four threads to collect the TCP upload and download speed. So, the values in TCP upload and download fields are the addition of the four threads.
wTCP_DOWN1	1156.3	TCP download speed to the West server at the first measurement. (in kbps).
wTCP_UP2	567.8	TCP upload speed to the West server at the second measurement. (in kbps).
wTCP_DOWN2	889.1	TCP download speed to the West server at the second measurement. (in kbps).
eTCP_UP1	346.2	TCP upload speed to the East server at the first measurement. (in kbps).
eTCP_DOWN1	914.8	TCP download speed to the East server at the first measurement. (in kbps).
eTCP_UP2	278.5	TCP upload speed to the East server at the second measurement. (in kbps).
eTCP_DOWN2	448.2	TCP download speed to the East server at the second measurement. (in kbps).

ABNORMAL VALUE DESCRIPTION

Value	Description
no effective service	If a test failed the connectivity test (ping to our west server for four seconds), our post-processing script generates "no effective service".
ERROR: QUIT BY USER	A user quitted the testing. So, the raw should be removed for the data analysis.
NA at ePktMin, ePktMax, ePktAvg, ePktLoss, wPktMin, wPktMax, wPktAvg, wPktLoss	The testing result didn't provide the corresponding round trip time (RTT) value. In many cases, it means that the ping command had bad connections (or 100% loss.)
timeout	Our measurement tool spends 120 seconds to measure a TCP upload and download to a server. If our tool fails to get the TCP testing result within 120 seconds, it generates "timeout", which usually occurs in a weak signal location.

connect_error1	While a TCP testing, our measure tool fails to finish the TCP testing with an internal software error message "write1 failed". There are many reasons of this error. -- We are investigating the exact reason. Weak signal could be one of the reasons.
connect_error2	While a TCP testing, our measure tool fails to finish the TCP testing with an internal software error message "write2 failed". There are many reasons of this error. -- We are investigating the exact reason. Weak signal could be one of the reasons.
connect_error3	While a TCP testing, our measure tool fails to finish the TCP testing with an internal software error message "connect failed". There are many reasons of this error. -- We are investigating the exact reason. Weak signal could be one of the reasons.

Appendix G – Data Processing

In order to produce the analysis in this Report, the data were processed as follows:

- **Averaging two devices (Data card and smartphone results):** in the event either device's TCP test results showed one of the Abnormal Value Descriptions listed in Appendix F – Data Record Format, while the other device yielded numerical TCP results, the Abnormal Value Description result was treated as a zero. This meant that the average throughput for a location where one device produced a result while the other did not was effectively half of that of the device that successfully delivered a throughput greater than zero. Example: Data card throughput was 14 Mb/s, while smartphone result was "no effective service." The throughput for that location for that carrier was 7 Mb/s.
- **Elimination of duplicate "no effective service," "timeout," etc. results for a single device:** in cases where initial ping test failed, testers were instructed to run the test again up to 3 times. In cases where the ping test failed all 3 times, the results would show one or several of the Abnormal Value Descriptions listed in Appendix F – Data Record Format for all of the tests. In this case, the result for that testing location was "no effective service." In the event one of the ping tests succeeded, the TCP tests results were used and the other failed ping test results were excluded from the analysis. If some or all of the TCP tests for a location with a successful ping test showed one or several of the Abnormal Value Descriptions, the TCP results with those Abnormal Value Descriptions were treated as zero Kb/s throughput, and were included in the calculations used for this analysis.
- **GPS coordinates normalized to average:** tests were intended to be conducted for all four carriers at the exact same location. Due to minor differences in GPS readings as well as practical problems of testers sometimes being unable to conduct all tests at the exact same location due to local conditions, GPS coordinates for all tests performed at each location were averaged to a single set of LAT / LONG coordinates.
- **No GPS signal:** in cases where no GPS data was collected due to the inability of the phone to register a GPS signal, the normalized GPS coordinate was assigned to the test result based on the location ID.
- **Carrier information missing:** the test automatically recorded Provider as well as Operator identification information from the carrier's network. In cases where both Provider and Operator fields were blank, we cross referenced the unique device ID to the carrier in order to determine the correct Provider field entry.
- **Location ID entered manually by tester:** while each tester used a checklist to monitor which sites had been tested, in some cases testers entered an incorrect location ID. To correct for this, we looked at the GPS coordinates of any duplicate location IDs from different testers and assigned the correct location ID based on the LAT / LONG coordinates from the device. In very few cases, there was an incorrect location ID with no GPS coordinate, in which case the unique device ID was cross referenced to the list of sites being tested by each tester to try to determine the correct location ID for that test. In cases where this was not successful, the test result was excluded from analysis.

Appendix H - Interpolation Model

The interpolation (Kriging) models were created from the point data collected from the second test of wireless towers. A mask was used so that the surface interpolation and points used for analysis are only within the provider's service area based on the broadband availability data provided by the CPUC in the broadband service wireless GIS shapefile. The Kriging model was made using a 1 kilometer resolution. The interpolation is based on the default nearest 12 points. This produces a raster surface that only encompasses the provider's service area.

To assess the accuracy of the model surface rasters were created to symbolize standard deviation. They were based on the means and standard deviations from the predictive Kriging model. The model was also validated by interpolating raster model results to point that coincided with the original point files containing upstream and downstream data. These data were then analyzed and compared in the statistical computing software R. Observed data was plotted against predicted data and a simple linear model was developed to assess if the predicted data significantly represented the observed data (Figure 3). To further analyze differences between the predicted and observed data a paired t-test was performed.

The analysis showed that all models were significant predictors (p-value < 0.05) and differences between observed and predicted speeds were found to be statistically insignificant (p-value > 0.05). The large variation in the measured data presents some statistical issues with accurately predicting the upper extremes off the upstream and downstream speeds. The high standard deviations are due to hot spots (seemingly located around populous areas containing high densities of mobile wireless towers) of high speed that vary highly from the mean. The results from the paired t-test also show the mean difference between the observed and predicted speeds.

One way to possibly reduce the variability seen in the state wide analysis would be to run the analysis at the regional level. For example; analysis done for Southern California, Central California, and Northern California separately may lower the mean difference between observed and predicted speeds. There are other ways to divide up the state regionally but this would allow for future analysis by county and incorporating populations (as does the state wide assessment). This may lessen the mean differences seen but will likely still carry a large standard deviation as areas with high population densities tend to have far higher speeds than more rural areas.

Kriging model t-test p-values

Provider	Type	t-test p-value	Mean Difference
ATT	Upstream	0.03023	-127.9708
ATT	Downstream	0.1514	-201.1415
Sprint	Upstream	0.9553	-0.7801853
Sprint	Downstream	0.8685	3.254113
T-Mobile	Upstream	0.5873	19.16867
T-Mobile	Downstream	0.6973	37.14667
Verizon	Upstream	0.5491	-82.67628
Verizon	Downstream	0.7666	-51.23117

Appendix I – R-value and MOS Calculations

1. Effective latency. Take the average latency, add jitter, but double the impact to latency, then add 10 for protocol latencies
    a. (Avg. latency + Avg. jitter)x2+10
2. Calculate R-value. Implement a basic curve - deduct 4 for the R value at 160ms of latency (round trip). Anything over that gets a much more aggressive deduction. Deduct 2.5 R values per percentage of packet loss
    a. =IF(Effective latency<160,(93.2-(Effective latency/40)-2.5 x Avg. packet loss),(93.2-(Effective latency-120)/10-2.5 x Avg. packet loss))
3. Convert the R into an MOS value
    a. =IF(AND(ISNUMBER(R-value)),IF(AND(R-value>0,R-value<101),(1+0.035 x R-value+0.000007 x R-value x (R-value-60)*(100-R-value)),0),0)