CalSPEED: Measuring California Mobile Broadband - A Comparison of Methodologies

Ken Biba
Managing Director and CTO
Novarum, Inc.

Three major mobile network measurement tools: two public - the Federal Communications Commission's (FCC) tool and California's CalSPEED and the proprietary Ookla, are analyzed and their methodologies and results compared. All three, when tested at the same time, in the same place, on the same networks, offer similar results - suggesting that they are all measuring the same networks in related ways. However, a detailed examination shows that each measures in a different way, offering differing views of network quality and providing different levels of credibility.

CalSPEED

- Most rigorous, all-inclusive network measurement, assessing actual user mobile Internet experience rather than simply measuring the radio access network.
- Pro-active sampling method produces the most reliable results avoiding the selection bias and large number of samples required by crowdsourcing.
- All tests use the same two servers, making comparisons possible.
- Assesses streaming quality for UDP based Voice over Internet Protocol (VoIP) as well as TCP web access quality.

FCC

- It "primes the pump" biasing TCP measurements towards higher values than users will likely see.
- Intentionally selects test servers for lowest latency biasing results, but current set of servers are widely geographically dispersed.
- Exclusive reliance on crowdsourcing introduces selection bias and the need for a large number of measurement samples for reliable statistics.

Ookla

- Proprietary measurement algorithm biases results by discarding bottom half of upstream results and bottom third of downstream results to show highest throughput and lowest latency.
- Default selects VERY local measurement servers for the very lowest latency that effectively tests only the local radio access network.
- Exclusive reliance on crowdsourcing introduces selection bias and the need for a large number of measurement samples for reliable statistics.

CalSPEED adds the explicit methodology for creating maps of estimated service across the entire state.

1. Calibrating the Mobile Internet Experience

Each of us relies on the Internet to research school papers, to find and buy new products, to read the news and increasingly to entertain ourselves. The Internet is not only becoming our newspaper, but also our phone, radio and television. How we do our jobs, raise our families, educate ourselves and our children, interact as responsible citizens, and entertain ourselves are all influenced by the quality of the Internet service we obtain. And ever increasingly, that service is not on our desk, but in our hand wherever we go.

Knowing the quality of this service is a vital piece of our modern ecosystem in the same way we research the brand of car we drive or the type of house we own. With multiple mobile Internet providers, an independent third party assessment of this quality allows consumers and policy makers to make informed choices.

CalSPEED is an open source, non-proprietary, network performance measurement tool and methodology created for the California Public Utilities Commission with the assistance of a grant from the National Telecommunications and Information Administration. CalSPEED uses a methodology pioneered by Novarum. The software measurement system was created by a team at California State University (CSU) at Monterey Bay, led by Professors Sathya Narayanan and YoungJoon Byun. CalSPEED mapping and measurement field operations were managed by the Geographic Information Center at CSU Chico. Statisticians at CSU Monterey Bay assisted the team with detailed geographic and statistical analysis of the dataset.

While CalSPEED was initially intended to measure only mobile broadband, it has now been extended to evaluate fixed wireless and wired connections. California has used CalSPEED for two years with five rounds of measurement over the entire state collecting over 5,000,000 measurements across California of the four major mobile broadband carriers: AT&T Mobility, Sprint, T-Mobile and Verizon Wireless.

This paper describes how CalSPEED compares to two other mobile Internet measurement methodologies used by the FCC and by the commercial entity Ookla. A companion paper titled *CalSPEED: California Mobile Broadband – An Assessment* reports on what CalSPEED has discovered about California mobile broadband.

2. CalSPEED

CalSPEED is an open source network performance measurement tool that is based on an industry standard open source performance measurement tool - iPerf[1]. iPerf provides the foundational network testing engine for both the TCP and UDP protocols. CalSPEED packages this testing engine in both Windows and Android client tools to measure and record network performance. While CalSPEED was initially targeted at evaluating mobile broadband networks it can also be used to evaluate fixed wireless and wireless networks as well.

[1]iPerf is an industry-standard suite of broadband testing tools for measuring TCP and UDP bandwidth performance. See https://iperf.fr/

2.1 Capturing the End to End User Experience

CalSPEED has five core anchors that define its methodology: open source tools, assessing the full end-to-end user experience, just the facts, not just for crowds and mapping that is useful for decision makers, not just for information. Let's look at each.

Open Source. CalSPEED uses iPerf as the foundational network measurement engine for both the TCP and UDP protocols. CalSPEED packages this engine in both Windows and Android client tools to measure and record network performance.

End-to-End User Experience. A foundational assumption of CalSPEED, unique among network measurement tools, is an attempt to replicate the end to end user experience. In particular, CalSPEED recognizes that the Internet resources a typical user accesses are scattered across the entire Internet and, despite the use of content delivery networks to speed Internet performance by caching frequently accessed content, are not always "local" to the user. Many measurement tools focus on evaluating just the local radio access network - the last few miles - and not the backhaul network to the ultimate server resource used. CalSPEED instead tests the complete network path, from the client device, through the local access network, through the Internet backbone, to several server destinations.

CalSPEED emulates this user experience with two fixed servers - one physically located in Northern California and the other in Northern Virginia - both in the Amazon Web Service (AWS) cloud. CalSPEED reports performance both to each individual server and the average between them. Not only does this method measure the different local access methods, but provides a network interferometry that gives insight into the different backhaul strategies chosen by carriers. We find carrier unique 2:1 differences in end to end latency and jitter and material difference in upstream and downstream throughput between the two servers.

These differences in fundamental network performance illustrate that location matters - Internet performance delivered to the user or the Internet user experience - will vary based on where on the Internet the desired server is located. And desired servers are scattered across the Internet, not just close to every user.

CalSPEED measures a complete portfolio of network metrics including end-to-end packet latency, bidirectional TCP throughput, UDP packet loss and jitter. Appendix A describes the precise algorithm.

Just the Facts. CalSPEED does not filter any of the results -throughput, coverage, latency or other network metric. We believe that, just like the user experience where sometimes a web page fails to load, all results are valid representing the user experience.

Not Just for Crowds. Crowdsourcing is a fashionable method for collecting data at scale - but it has an inherent selection bias of only collecting data from those people who choose to use the application and at the locations where these people are. Where there is no crowd there is no data. And even where there is data, it is biased towards who collected it, why, when and where.

CalSPEED has two complementary methods of testing - the first is a structured sampling program of 1,200 locations scattered throughout California (rural, urban and tribal). In Fall 2013 or testing round four, we increased the number of test points to 1,986 to improve predictive precision of the interpolation models. Each of these locations are periodically (every six months) visited and methodically measured with CalSPEED on both the latest Android phones and a USB network device on a Windows based netbook for each of the four major carriers. The use of multiple contemporary user devices gives a good snapshot of the leading edge user experience.

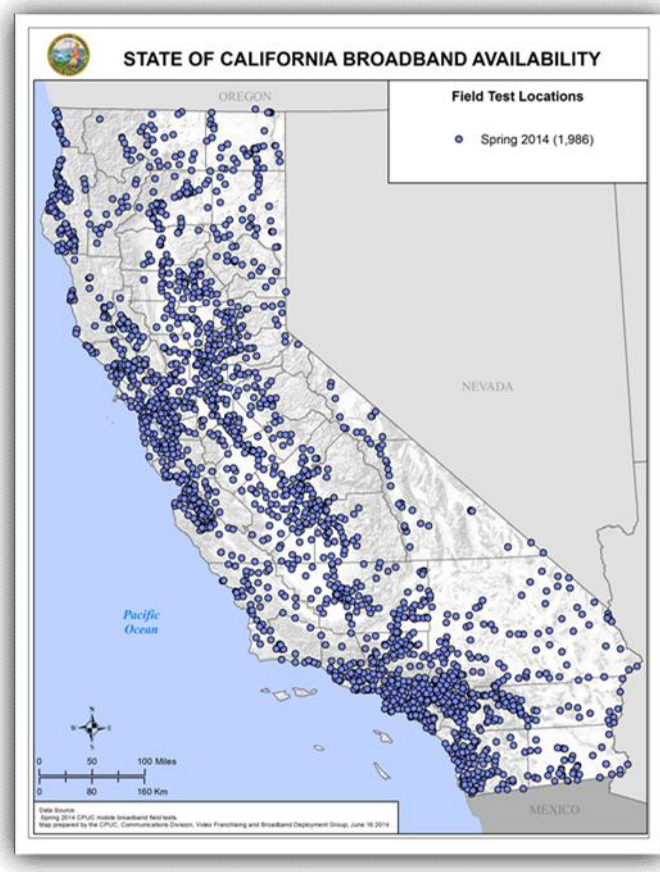

The second is the independent use of CalSPEED to provide crowdsourced data. The structured sampling program avoids selection bias of when and where measurements are made; giving a full map that covers the entire state, including places not often visited by smartphone users but which have mobile broadband service. The crowd sourced data adds further detail to areas where there are people who choose to use the test and about the range of the installed base of phones (particularly legacy mobile devices) and the performance those user devices are experiencing. The structured measurement program uses current user devices and thus gives a snapshot of the latest deployed network technology. Older user devices, with older wireless technology, will likely get slower performance in many locations.

CalSPEED explicitly samples all the major demographic groups in California - urban (37%), rural (56%) and tribal (7%). Thus, CalSPEED is able to explicitly measure the mobile digital divide.

Maps for decision-makers not just for information. We then take the measurement data and create geospatial kriging[2] maps interpolating CalSPEED measurements of (but not limited to) latency, downstream and upstream throughput, Mean Opinion Score, jitter and packet loss over the entire state.

These maps can be overlaid with other geostatistical data on population, economics, and census

[2] Kriging is an interpolation technique in which the surrounding measured values are weighted to derive a predicted value for an unmeasured location. With 1,986 testing locations, we can obtain interpolations at one kilometer resolution. Additional testing can be done in specific areas of interest at smaller resolutions. See http://support.esri.com/en/knowledgebase/GISDictionary/term/kriging for more information.

group to provide informed choices for consumers, businesses and governments. The CPUC web site uses this data to suggest what mobile service is available and at what performance at locations of the consumer's choice. The maps below estimates the coverage and mean downstream throughput of each of the major carriers in California in the Fall 2013. Similar maps are available for upstream throughput, latency, jitter and MOS to assess VoIP in the companion report CalSPEED*: California Mobile Broadband – An Assessment*.

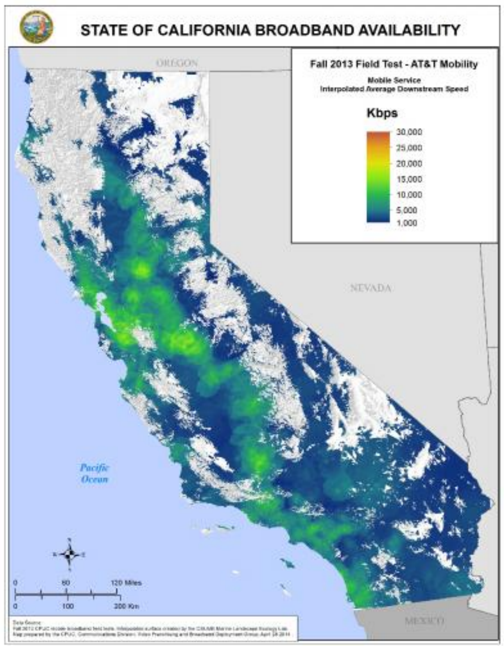

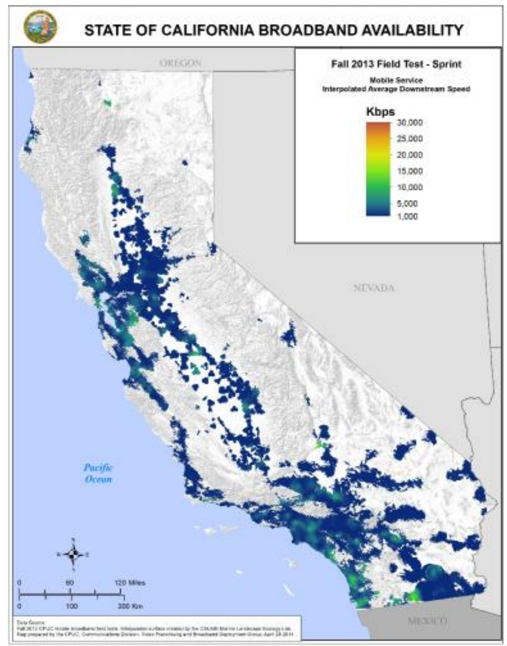

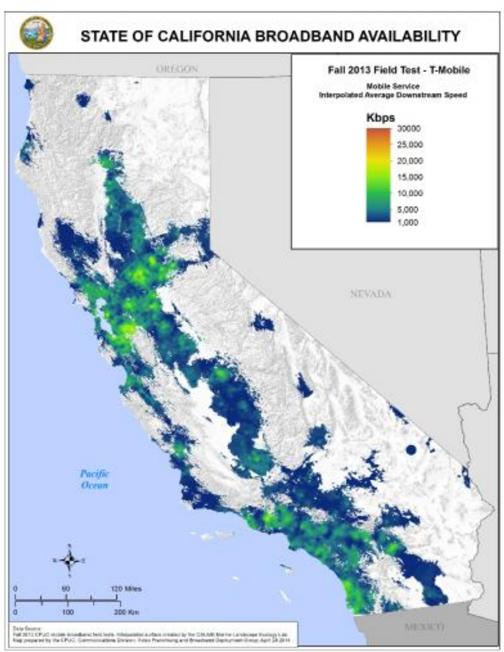

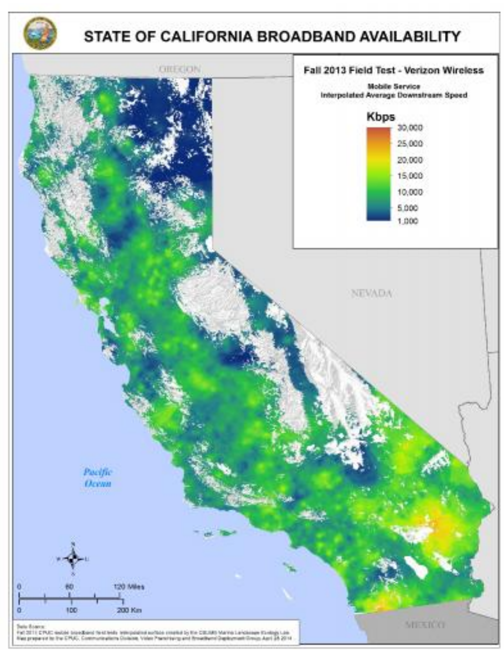

2.2 CalSPEED Methodology

CalSPEED performs the following sequence of measurements to gather its information:

1. ICMP ping to the West server for four seconds for connectivity checking. If the ICMP ping fails, CalSPEED presumes that there is no effective connectivity to the Internet and records that result.
2. iPerf TCP test (four parallel flows) to the West server - both downstream and upstream. CalSPEED uses four parallel flows to ensure that the maximum capacity is measured during the test.
3. ICMP ping to the West server for ten seconds to measure latency to the West server.
4. UDP test to the West server. CalSPEED constructs a UDP stream of 220 byte packets to emulate a VoIP connection with 88kb/s throughput. This UDP stream is used to measure packet loss, latency and jitter.
5. iPerf TCP test (4 parallel flows) to the East server to measure downstream and upstream TCP throughput.
6. ICMP ping to the East server for 10 seconds to measure latency to the East server.
7. UDP test to the East server to measure packet loss, latency and jitter with a simulated VoIP data stream.

CalSPEED does not filter any of the results - either throughput or latency. We believe that, just like the user experience where sometimes a web page fails to load on a smartphone, all results are considered valid.

CalSPEED has two complementary methods of testing - the first is a structured sampling program of 1,200 locations scattered throughout California (rural, urban and suburban). In Fall 2013 or testing round four, we increased the number of test points to 1,986 to improve predictive precision of the interpolation models. Each of these locations are periodically (every six months) visited and methodically performance measured with CalSPEED on both the latest Android phones and a USB network device under Windows for each of the four major carriers. The second is the independent use of CalSPEED to provide crowdsourced data. The structured sampling program avoids selection bias of how, when and where measurements are made, giving a full map that covers the entire state, including places not often visited by smartphone users but which have mobile broadband service. The crowd sourced data adds additional detail to areas where there are people who choose to use the test and adds additional detail about the range of the installed base of phones and the performance those user devices are seeing.

We then take the sample data and create geospatial kriging maps interpolating CalSPEED measurements of latency, throughput, jitter and packet loss over the entire state. Synthetic measurements such as Mean Opinion Score (MOS) to evaluate VoIP readiness are also mapped.

3. How Does CalSPEED Compare?

CalSPEED is one of several network performance measurement tools available free of charge to the public. We measured at the same locations, at the same time with both the latest FCC[3] and Ookla[4] test on the same smartphones, in the same locations and at about the same time[5].

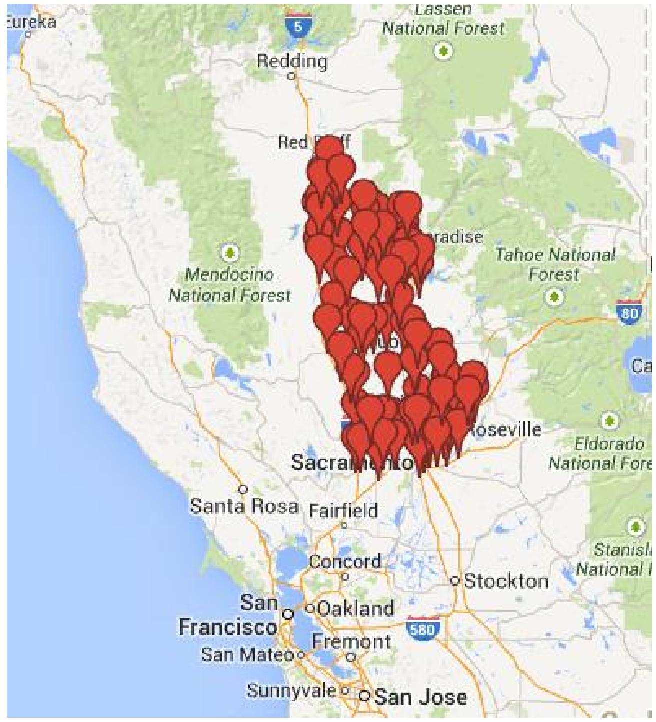

In the measurements reported here, our testing team used CalSPEED, the FCC mobile and the Ookla mobile application to measure Internet performance at a subset of the mobile broadband footprint in California, designated by 91 randomly scattered locations in Northern California. Our team also previously tested 42 of the 91 locations in the earlier four rounds of the CalSPEED measurement program. This history of measuring network performance at the same location over time is a unique feature of CalSPEED. In each location CalSPEED measured performance to two measurement servers - one on the East Coast and one on the West Coast. FCC and Ookla measurement tools each chose a variable test server at the time of the test from among each tool's portfolio of test servers.

3.1 How the Tools Compare

All three tools have more in common than differences. All are much more advanced than early broadband testing tools (and some remaining carrier performance testing tools), that merely timed the bulk file transfer of a fixed file. All of these tools measure latency as well as TCP based downstream and upstream throughput.

FCC and Ookla differ from CalSPEED in that they focus on discovering the best possible wireless edge access network performance rather than the "typical" end-to-end user experience.

Let's look at each of these tests a bit more.

3.1.1 FCC

The FCC has developed and deployed a mobile broadband measurement system[6].

[3] https://play.google.com/store/apps/details?id=com.samknows.fcc
[4] https://play.google.com/store/apps/details?id=org.zwanoo.android.speedtest
[5] The complete test suite took over 30 minutes in each location to run all the measurements.
[6] http://www.fcc.gov/measuring-broadband-america/mobile/technical-summary

Like CalSPEED it measures latency based on ICMP pings. However, the server used to conduct the test with is individually chosen based on which server from a set of possible test servers produces the fastest latency. It is possible, though uncommon, that consecutive tests in the same place and about the same time on the same carrier may test to different servers. This makes comparing measurements more difficult. In any case, the selection of the lowest latency test server biases subsequent throughput tests towards higher performance. As we will see from our actual test results, tests which are biased towards servers which are close to the client, measure less of the Internet backbone.

Throughput tests are via HTTP GET and POST layered on TCP to the selected server. Throughput tests use up to four parallel streams. However, the test preconditions streams to get past possible slow TCP startup. This precondition additionally biases the throughput test towards reporting higher throughput. The CalSPEED test has discovered that at least 10% of attempted TCP connections fail to transfer any data even for the best carrier in an urban area - with even worse performance in rural areas and for other carriers. We would expect from the test description that the FCC's TCP tests would discard failing results - further biasing the results towards a higher reported throughput.

The test is solely based on crowdsourced data, requiring users to choose to use the test. The FCC has no control over the hardware the test is run on, where and when the tests are run. It requires many more tests to get full coverage across all demographics of urban, rural and tribal. To partially compensate, the test is designed to run in background occasionally in order to provide a wider scope of tests in more locations at more times.

3.1.2 Ookla

Ookla is a commercial company that supplies a widely used mobile and fixed network performance measurement tool[7].

Unlike CalSPEED and the FCC's test it measures latency based on HTTP response time. However, like the FCC, the server it uses is chosen for each test, each time based on which server from a set of possible test servers has the fastest latency. It is possible, though uncommon, that consecutive tests in the same place and about the same time on the same carrier may test to different servers. This makes comparing measurements more difficult. In any case, the selection of the lowest latency test server biases subsequent throughput tests towards higher performance. As we will see from our actual test results, Ookla's wide network of test servers creates a bias towards using a test server VERY close to the client, measuring very little of the Internet backbone.

Throughput tests are via Flash encapsulated HTTP GET and POST layered on TCP to the selected server. Throughput tests use up to eight parallel streams. However, after collection, the test selects and discards both the top 10% and bottom 30% of downstream throughput samples. This discard biases the downstream throughput test towards reporting higher throughput. Further, the upstream test discards the slowest 50% of upstream throughput samples. The resulting reported upstream

[7] http://www.ookla.com/support/a21110547/what-is-the-test-flow-and-methodology-for-the-speedtest

throughput would appear to be biased towards a higher throughput.

Like the FCC test, the Ookla test is solely crowdsourced based, requiring users to choose to use the test. Ookla has no control over the hardware the test is run on, or where and when the tests are run. The latency filtering which finds low latency test locations should bias the throughput results by finding locations with a higher occurrence of high throughput due to low latency.

3.1.3 Summary Comparison

	CalSPEED	FCC	Ookla
Server Selection	Server locations fixed on each coast of the US. One on the Amazon Northern California and Amazon Northern Virginia.	Dynamic server location with many servers. Chosen from a set of local servers based on shortest latency at time of test.	Dynamic server location with many servers. Chosen from a set of local servers based on shortest latency at time of test.
Data Filtering	ICMP ping probe to determine if a TCP connection should be attempted. No filtering of TCP or ICMP ping results.	1. Tests run to selected lowest latency server. 2. Latency > 3000 msec discarded as lost packet	1. Tests run to selected lowest latency server. 2. Top 10% of high throughput and bottom 30% of low downstream throughput results discarded. 3. Bottom 50% of upstream throughput results discarded
TCP Throughput	Direct TCP. 20 1 second tests to each of West and East server over four (4) parallel streams. No accommodation for TCP slow start. 40 tests total for each test using both servers. No limit on data size transferred	TCP via HTTP GET and POST. 1 single connection GET, 1 multi-connection GET, 1 single connection POST, 1 multi-connection POST. Three parallel streams. Streams preconditioned to get past slow TCP startup. 60 MB daily and 20s limits	TCP via HTTP GET and POST via Flash. Up to (variable) eight (8) parallel streams are used. The fastest 10% and slowest 30% of downstream samples are discarded and the remainder averaged. The slowest 50% of upstream samples are discarded and the remainder averaged.
UDP Throughput	88 kb/s data stream. Reported from both East and West servers.	N/A	N/A
Latency	Measured twice, both from ICMP and from UDP stream	Measured from UDP	Measured from HTTP request response time
Packet Loss	Measured twice, both from ICMP and from UDP stream	Measured from UDP	N/A
Jitter	Measured from UDP stream	N/A	N/A
Intellectual Property	Fully Open Source based on de facto standard Internet measurement tools	Open source with embedded technology from Sam Knows	Proprietary

3.2 Where We Tested

We selected 91 locations in Northern California to measure - distributed randomly through northeast California - covering both rural and urban demographics. The southern edge of the footprint covers Sacramento, CA - California's sixth largest city - as the largest urban area in the measurement footprint.

3.3 How We Tested

CalSPEED normally selects two current devices from each carrier as its measurement instruments - a current model Android smartphone and a USB network device for a laptop - both devices are equipped with current technology radios that match the deployed infrastructure of the carrier. These are representative for the types of devices currently being used in the field. Not all client devices give the same performance.

For this measurement, we used only the smartphones for each carrier. Each smartphone for each carrier was configured with the three measurement applications: CalSPEED, FCC and Ookla. For each location, each application was run sequentially three times in each location on each smartphone for each carrier or a total of 36 measurements executed for each location using all three measurement tools and the four carriers.

3.4 End to End User Experience

Every CalSPEED measurement uses two fixed servers - one in San Jose, CA and one in Washington, DC, both AWS servers. All measurements are comparable since they are going to the same servers.

The Ookla measurement system filters from its collection of a large number of servers distributed across the United States to select the server with the lowest latency to conduct the measurement. Across our 91 test locations, Ookla used eight (8) different server locations varying from physically close by in Redding, CA to as distant as Reno, NV and San Jose, CA.

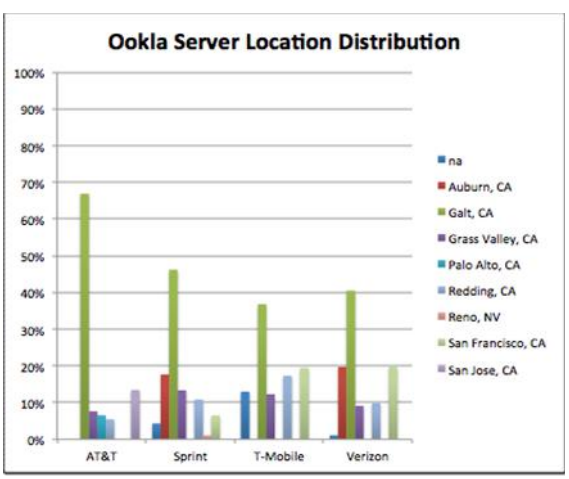

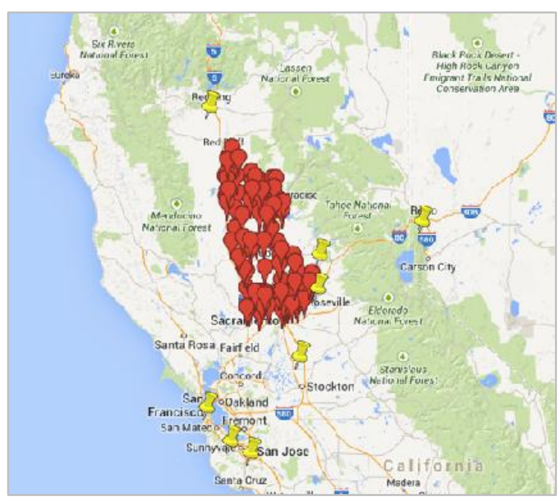

The composition of servers chosen differs between each carrier which makes comparing carriers, in a given location or across the complete set of locations, challenging since we are comparing different measurements. For example, 67% of the tests run for AT&T select the Galt, CA server while only 37% of the tests run for T-Mobile select that same Galt, CA server. Selecting a test server for the lowest latency produces a bias towards higher TCP throughput.

Now let's look at the server selection distribution for the FCC test. Across our 91 test locations, the FCC measurement selected four different servers ranging from Chicago, IL to Los Angeles, and Dallas, TX. The most commonly selected server was in San Jose, CA - approximately where the CalSPEED West server is located.

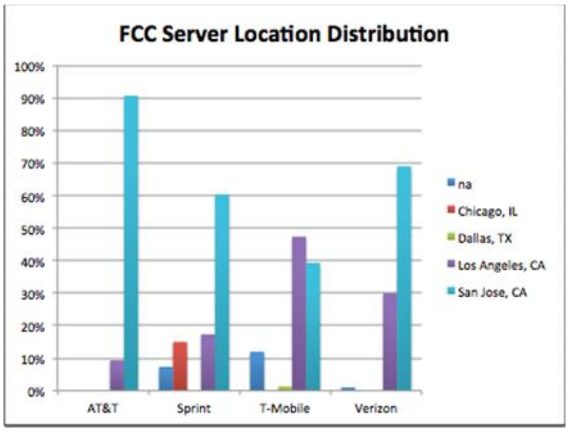

As with the Ookla tests, the distribution of server locations differs between depending on which carrier is tested making direct comparison of test results between carriers either at any given test location or among the complete set of test locations challenging. Each test independently filters to a unique set of server locations - filtering for lowest latency.

One of the challenges of comparing test results is the substantive difference in test locations and what that implies.

From the server distribution, we would expect that Ookla, on average across the footprint, will report a lower latency and a higher throughput than CalSPEED West. Similarly, we would expect that the FCC measurement would indicate a higher latency than CalSPEED West since the servers that the FCC uses appear to require traveling a greater distance than CalSPEED West. With the added latency of the Internet across the country, we would also expect CalSPEED East latency to be the slowest of all four of these measurements. As we will see, the latency and throughput measurements support most of these expectations.

The distribution of test server locations is perhaps one of the biggest differences between these measurement tools. We will see that this difference in test methodology will have implications in all the major measurements of network performance: latency, jitter, downstream and upstream throughput. And the difference between these is not the wireless access technology, but rather the backhaul choices that each carrier makes on connecting wireless cells of service into the overall Internet.

This distribution of server locations for both Ookla and FCC demonstrates the importance of measuring backhaul in assessing complete end to end user experience.

3.4 Latency

Each measurement system measures latency to different servers on the Internet, with differing wired backhaul, and so it is not surprising that they each show different latency distributions

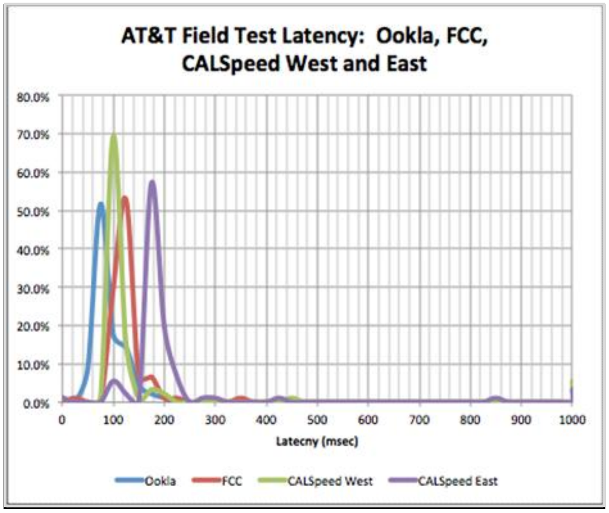

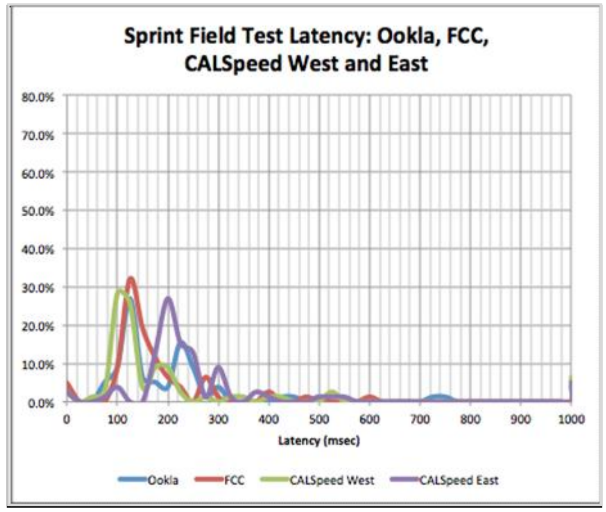

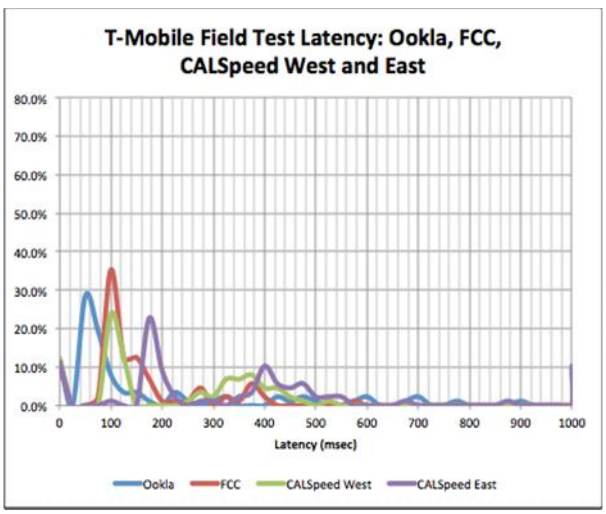

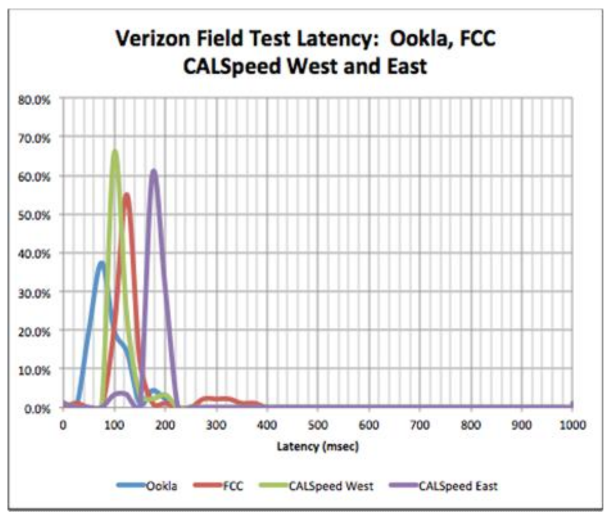

As we might expect given the geographic distribution of server sites, Ookla generally delivers the lowest absolute latency, on the order of 20% less than CalSPEED West and about half of CalSPEED East. This difference reflects the substantial difference in the length of the Internet to be traversed between the wireless access network and the server over the wired Internet backhaul.

The FCC measurement shows latencies equal to or a bit longer than CalSPEED West - again reflecting the longer distances to its measurement servers.

The latency measurements correspond to the intuitive results from the physical distribution of test servers for each test. Ookla delivers the shortest latencies, then CalSPEED West, then FCC, then CalSPEED East. This is despite the fact that Ookla should be biased slightly to longer latencies from overhead of measuring latency from HTTP response times.

Latency indirectly affects all other measurements. With longer latencies we would expect increased jitter and decreased downstream and upstream TCP throughput.

3.6 Jitter

We do not get jitter measurements from either the Ookla or FCC measurements. So only the CalSPEED measurements are relevant.

In general, jitter increases with increasing latency - such as between CalSPEED West and CalSPEED East. AT&T shows the least jitter, followed by Verizon. Both T-Mobile and Sprint show a wider range of jitter suggesting networks less able to deliver widespread, high quality VoIP services.

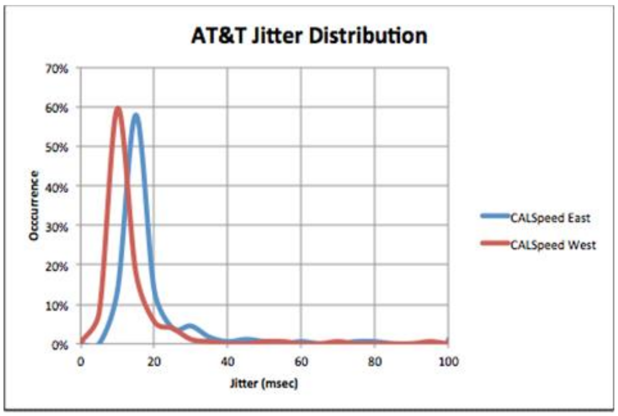

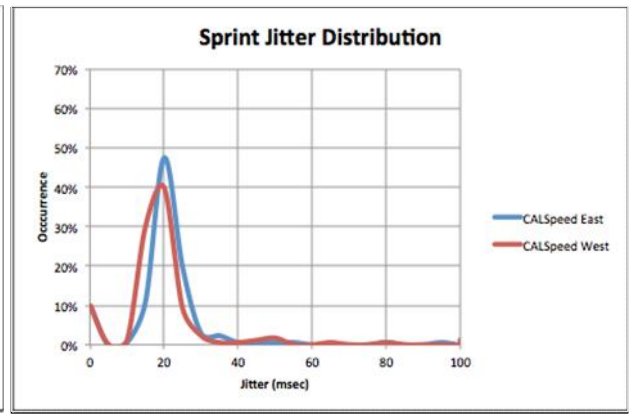

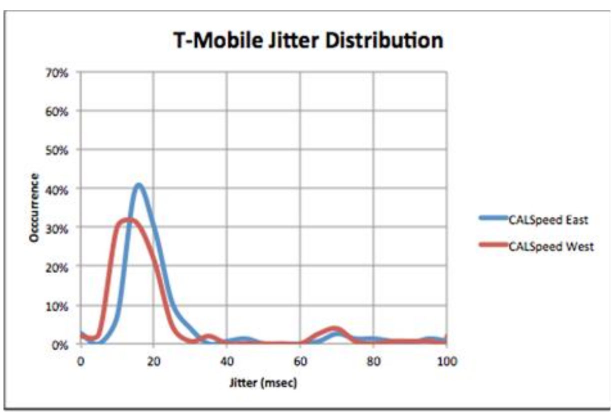

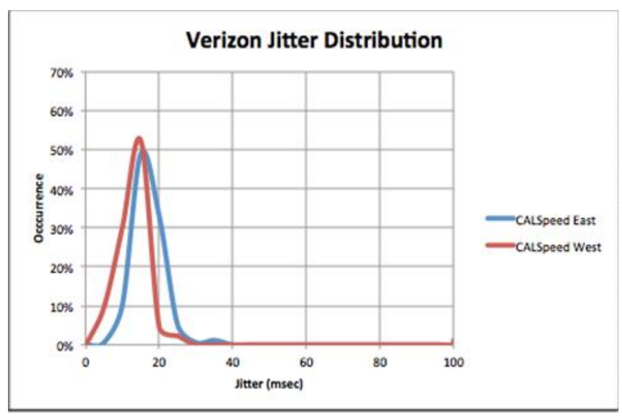

3.7 Packet Loss

Both FCC and the CalSPEED measurements give us assessments of packet loss, while Ookla does not.

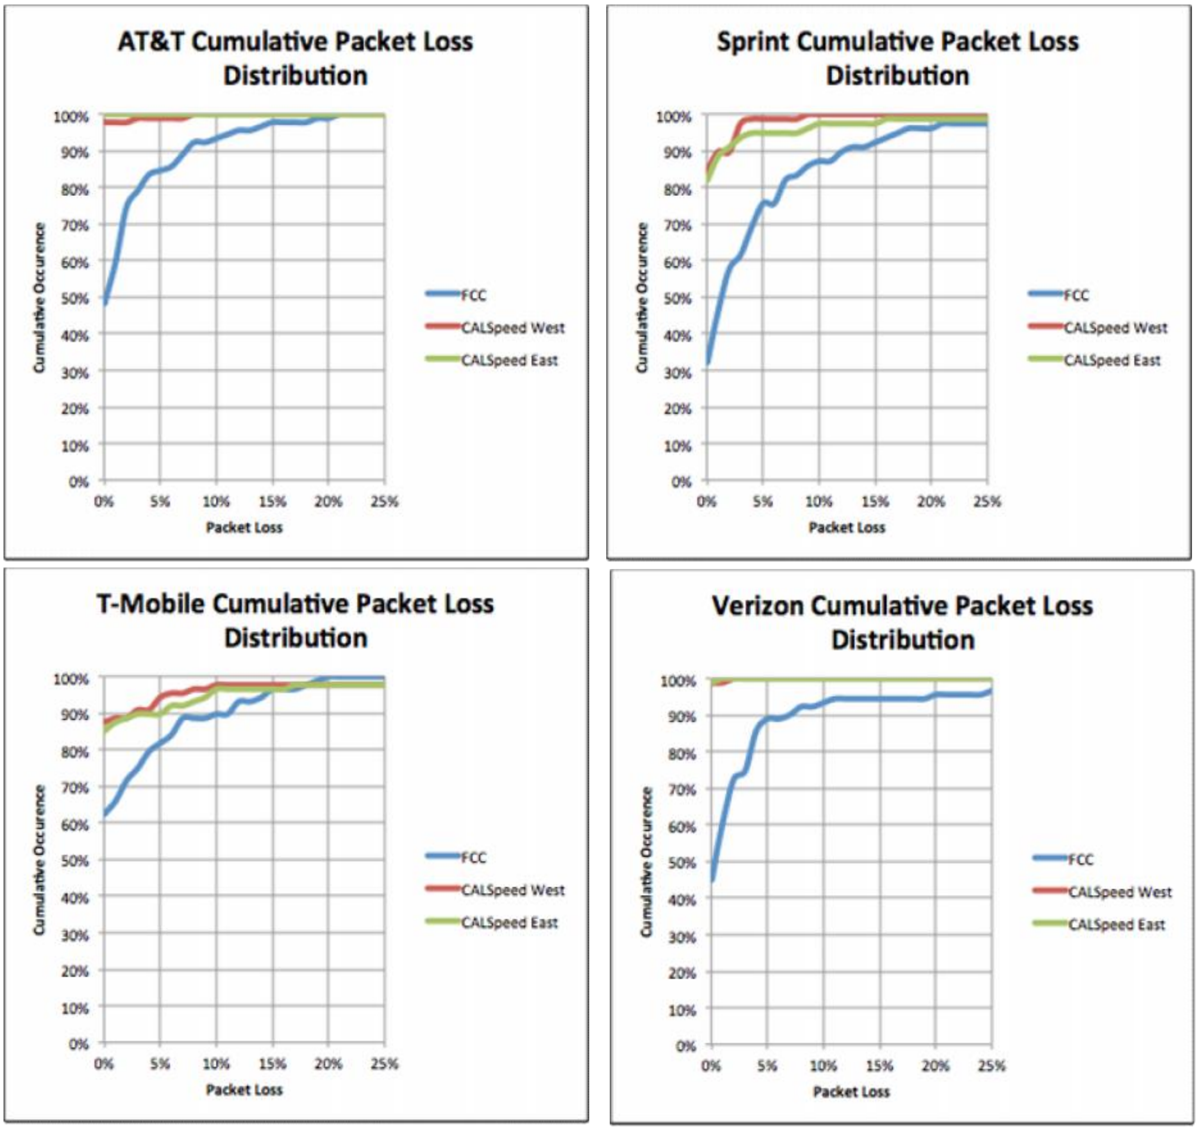

The FCC measurement shows materially higher packet loss numbers than either CalSPEED West or East. We do not have an explanation for these much higher packet loss measurements, but such high numbers would make VoIP or other real-time streaming services very difficult.

3.8 Downstream Throughput

In general, downstream TCP throughput measurements for CalSPEED West, FCC and Ookla are comparable. CalSPEED East, with the much higher latency involved with traversing the Internet between coasts, shows lower throughput more often than West, which is more apparent at very high speeds.

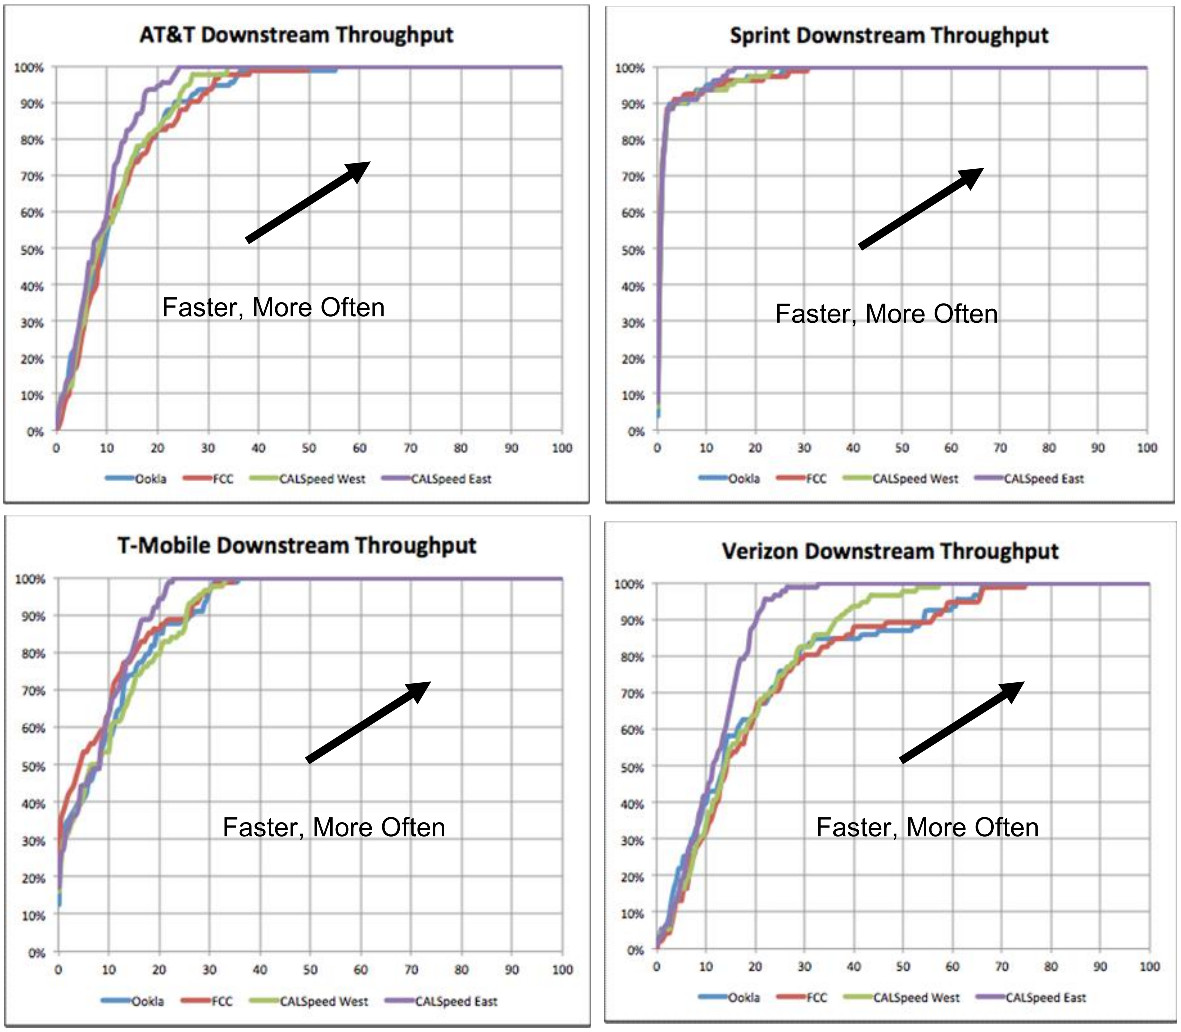

Selecting the server with the lowest latency biases the result to a higher average throughput measurement, particularly at higher performance levels. The differences between CalSPEED and FCC and Ookla are modest at low speed networks (Sprint and T-Mobile) with the advent of LTE we can see that the selection of lowest latency servers makes a difference, particularly at higher speeds. The median throughput for all these networks are quite comparable among the measurements, but diverge more sharply at high throughput.

And we can also see that at higher throughputs, the effects of the FCC and Ookla selecting high performance servers and discarding low performance TCP results (from slow TCP startup or failed connections) seems to result in a higher percentage of the sample showing higher throughput than CalSPEED West. We can infer an indirect effect of selecting the lowest latency server is selecting the highest performing servers rather than ones that represent typical user experience.

3.9 Upstream Throughput

The analysis of upstream TCP throughput shows a similar pattern to the downstream throughput. The FCC and Ookla measurements basically track CalSPEED West - showing the performance enhancement effect of testing only to local servers and do not show the effect of testing servers at a greater distance which CalSPEED East tests show.

Finding the lowest latency server biases the result to a higher throughput measurement, particularly at higher performance levels. The differences between CalSPEED and FCC and Ookla are modest at low speed networks (Sprint and T-Mobile) with the advent of LTE we can see that selecting for lowest latency servers makes a difference, particularly at higher speeds. Median throughputs are similar among the tests, but average throughputs differ materially.

Ookla and FCC both proactively filter results, albeit in different ways, with a bias towards report higher downstream and upstream throughput.

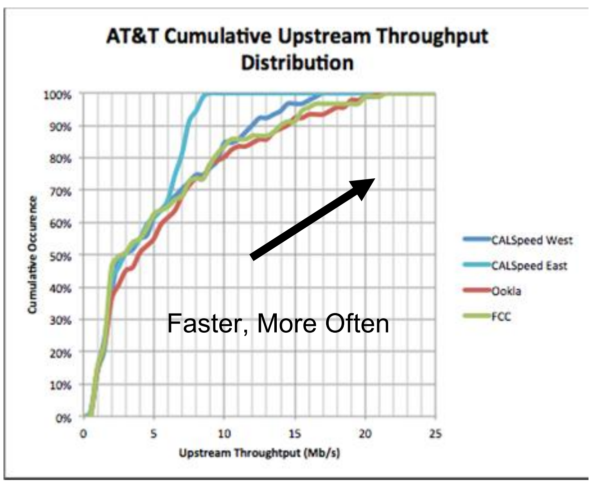

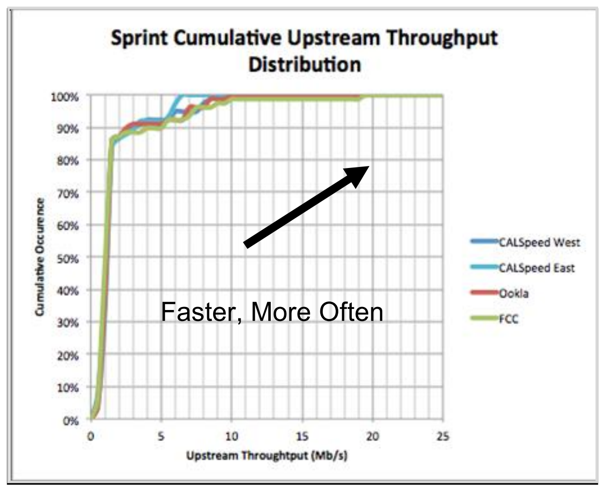

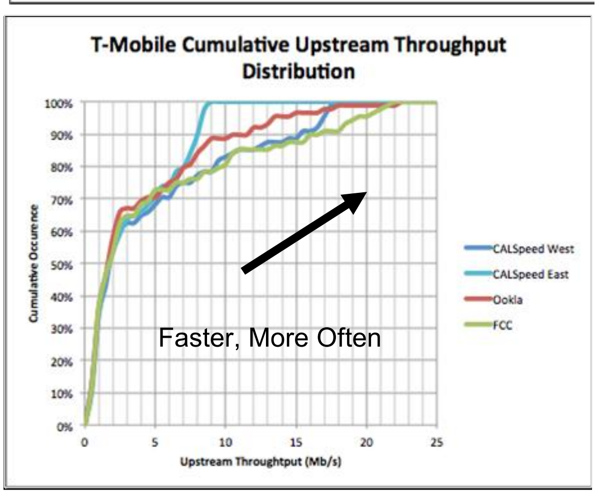

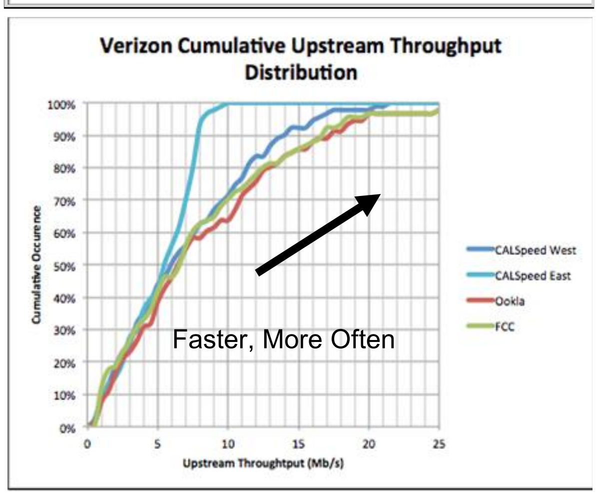

3.10 MOS

Among the three tests, only CalSPEED gives us sufficient information to calculate MOS or jitter, latency and packet loss. A MOS at 4 or above is considered acceptable quality.

The MOS distribution for this sample subset shows that only two carriers, AT&T and Verizon, have a significant footprint of VoIP capable service - at about 86% and 90% respectively. Sprint comes next closest at about 70% and T-Mobile lags far behind with about 40% of its local footprint being VoIP capable.

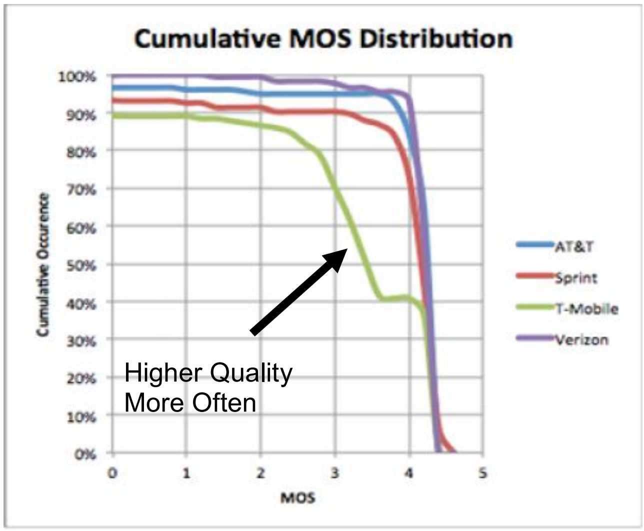

3.11 Wireless Network Technology

The charts below document the distribution of wireless access technology used by each carrier as measured by each test. All three testing methodologies give about the same result on the inventory of wireless technologies seen.

It is interesting to note the dramatic differences measured in our sample between the technologies deployed by each carrier.

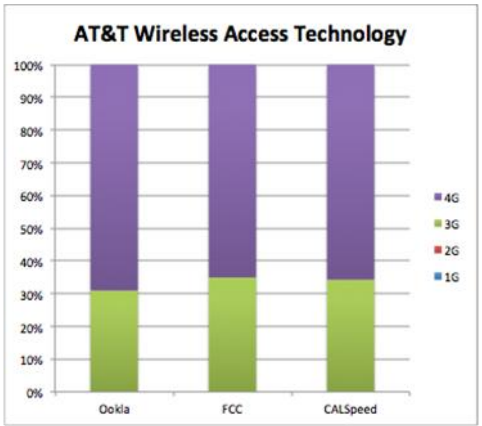

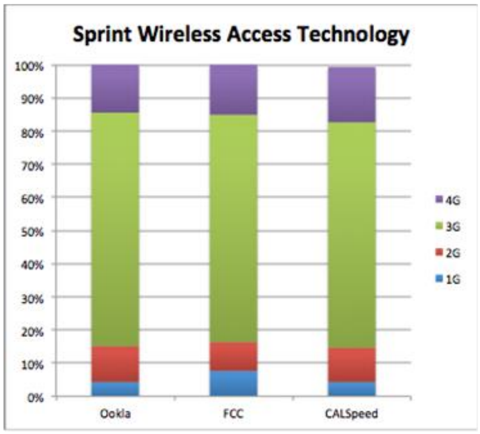

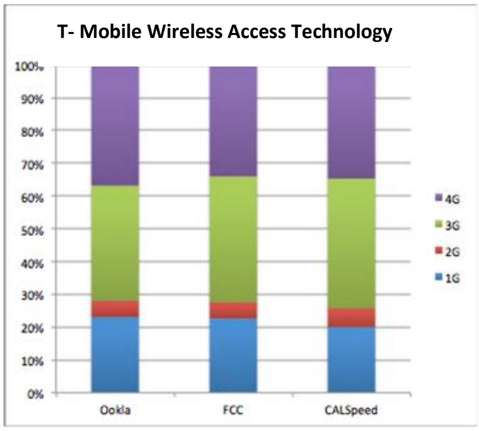

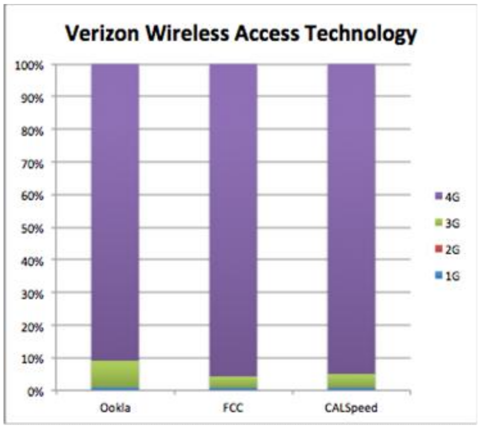

3.12 Interpolation and Mapping

CalSPEED uses kriging to establish geospatial maps of mobile broadband. This gives a visual representation of the growth and distribution of mobile broadband service. We have computed such estimated maps for a variety of measured network parameters - downstream throughput, upstream throughput, latency, and jitter - for each carrier over the announced coverage area for each carrier. An example of such a map is shown below for downstream throughput.

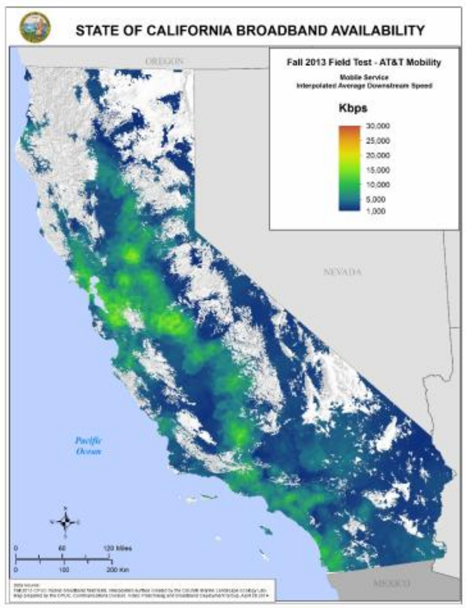

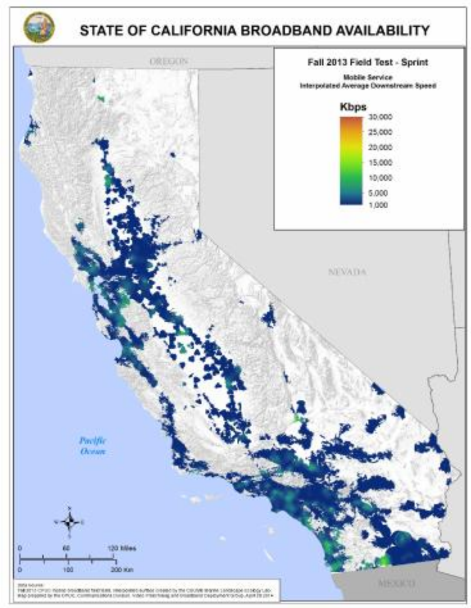

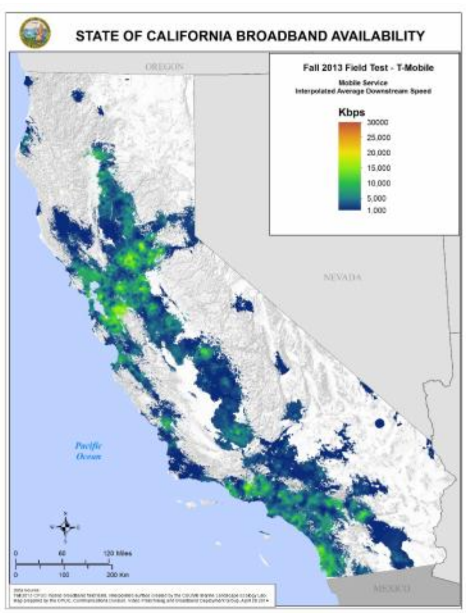

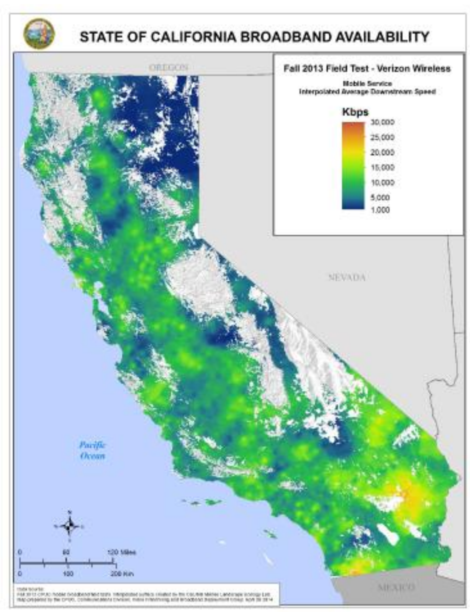

4.0 Conclusions

All three of the examined measurement systems offer similar results - suggesting that they are all measuring the same networks in related ways. However, each one, in detail, measures in a different way offering differing views.

CalSPEED	• CalSPEED's foundation is a structured measurement system that avoids selection bias and covers the entire state including urban, rural and tribal demographics. • Measurement to two fixed servers allows measurement of the effects of carrier specific Internet peering on mobile performance. • Measurement does not precondition test streams or filter results attempting to get the "best" performance • Fully open source based on de facto standard Internet measurement tools • Validated methodology includes creating interpolated kriging maps to cover the entire state not just the sampled locations. • Periodic resurvey from the same locations enables simpler comparison over time.
FCC	• The measured packet loss seems unnaturally high. • The suspected effects of the selection bias on server selection appear to bias average throughput higher that users will likely observe. • Data filtering biases result in higher reported throughput. • Crowdsourced data collection requires a very large dataset to effectively collect data covering all of the urban, rural and tribal demographics. • The wide variety of servers used adds uncertainty as how to compare measurements. • Little insight on the effects of Internet backhaul affecting performance.
Ookla	• The suspected effects of the selection bias on server selection seem to bias average throughput to be higher than users will likely observe. • Crowdsourced data collection requires a very large dataset to effectively collect data covering all of the urban, rural and tribal demographics. • Measured latency likely much smaller than users will observe. • Data filtering biases results towards higher reported throughput. • The wide variety of servers used adds uncertainty as to how to compare measurements. • Little insight on the effects of Internet backhaul affecting performance. • Proprietary.

CalSPEED adds the explicit methodology for creating maps of estimated service across the entire state.