APPENDIX A

CalSPEED: California Mobile Broadband - An Assessment

Ken Biba
Managing Director and CTO
Novarum, Inc.

This paper analyzes the first four measurement rounds of CalSPEED covering about one and half years of measurement between Spring 2012 and Fall 2013. CalSPEED is the open source, mobile broadband measurement tool and methodology used by the California Public Utilities Commission. Analysis of CalSPEED results provides a solid foundation for seven key findings, on which this paper elaborates.

1. Calibrating the Mobile Internet Experience

Each of us relies on the Internet to research school papers, to find a job, 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 as 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. In addition to crowdsourcing, the CPUC performs structured tests at the same testing locations every six months, by a team of trained testers. The software measurement system was created by a team at California State University 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 California State University at Chico. Statisticians at CSU Monterey Bay assisted the team with detailed geographic and statistical analysis of the dataset.

California has now used CalSPEED for two years with five rounds of measurement over the entire state collecting over 5,000,000 measurements across California for the four major mobile broadband carriers: AT&T Mobility, Sprint, T-Mobile and Verizon Wireless. This paper does a deep analysis of the first four rounds of measurement, with a subsequent update paper to follow to add additional analysis for the just collected fifth round data. A companion paper reports on how CalSPEED

compares to other measurement tools, namely, the FCC's Android speed test application and Ookla's.

This paper examines what the disciplined use of CalSPEED has revealed about the coverage and quality of mobile broadband in California over 2012-14. Our analysis of the CalSPEED dataset reveals the following trends for mobile broadband in California:

- Mobile broadband, at its best, has gotten much better VERY quickly, on average.
- Averages can be deceiving, because mobile broadband performance varies widely across the state and even at individual locations.
- Not all carriers are equal - there is a substantial and growing difference in the coverage and quality of mobile broadband service between the best performing and worst performing carrier reducing the number of competitors in many areas.
- Mobile broadband service is not just local - the user experience is governed not just by radio access but also by the backbone interconnect strategies of the carriers.
- A mobile digital divide between urban and rural and tribal appears to be getting larger.
- The bulk of the state's mobile networks are not Voice over IP (VoIP) ready.
- Measured performance is generally lower than advertised performance but the difference varies by carrier.

2. CalSPEED: Capturing the End to End User Experience

Open Source. CalSPEED is an open source network performance measurement tool that is in turn based on an industry standard open source performance measurement tool - iPerf[1]. iPerf provides the foundation network measurement engine for both the TCP and UDP protocols. CalSPEED packages this engine in both Windows and Android client tools for measuring and recording mobile 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 that 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 ultimate server destinations.

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

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 Services 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. Measurement to a local server only results in an overly optimistic expectation of service quality rather than what a typical user will actually experience.

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 its results - throughput, coverage, latency or other network metric - rather it uses the results of all tests performed and recorded. We believe that just like the user experience where sometimes a web page fails to load, all results are valid in representing the user experience. Other testing systems filter results in ways that bias the results and give a more optimistic expectation of network performance than what a user will typically 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 an application where it is chosen to be used by those people who choose to use it. 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 1986[2] locations scattered throughout California (tribal, rural and urban) that are each 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 best user experience.

The second method 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 having mobile broadband service. The crowd sourced data adds additional detail to areas where

[2] Originally 1200, but later increased to improve predictive precision of the interpolation models.

there are people who choose to use the test and adds additional detail 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 the most current user devices available at the time of each round of field measurement and thus gives a snapshot of the latest deployed network technology. Older user devices, with older wireless technology still in use by many, will likely get slower performance in many locations.

Because CalSPEED samples all areas of California - urban (37%), rural (56%) and tribal (7%), analysis of its results explicitly measures the state's mobile digital divide.

Maps for decision-makers not just for information. We then take the measurement data and create geospatial kriging[3] maps interpolating CalSPEED measurements of (but not limited to) latency, downstream and upstream throughput, jitter and packet loss over the entire state.

These maps can be overlaid with other geostatistical data on population, income, ethnicity, education, and census areas to provide more 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. Appendix B has further examples of upstream throughput and latency maps for California for Fall 2013.

CalSPEED has now had four rounds of sampling in California (Spring 2012, Fall 2012, Spring 2013 and Fall 2013) and a fifth round (Spring 2014) has been completed, but not yet analyzed. In each sampling round, we have surveyed the entire state and all four of the major wireless carriers - AT&T Mobility, Sprint, T-Mobile and Verizon Wireless.

Let's examine what CalSPEED tells us about the state of mobile broadband in California.

3. The Best of Mobile Broadband is Getting Much Better, VERY Rapidly

CalSPEED was designed to support comparisons over time of network performance. One key foundation of this design is its use of standard server locations, standard test locations, standard measurement metrics and an established pattern of structured measurement.

We have tracked three major trends over time: changes in performance (throughput, latency and jitter) due to new technology and capacity deployment, changes in performance due to increases in user load, and changes in coverage as carriers deploy their footprint.

For the best carrier, Verizon, downstream throughput increased by 3x over approximately a year and a half. The other three carriers showed a slower rate of performance improvement.

---

[3] Kriging is an interpolation technique in which the surrounding measured values are weighted to derive a predicted value for an unmeasured location. See http://support.esri.com/en/knowledgebase/GISDictionary/term/kriging for more information.

3.1 Throughput

One straightforward summary measurement is the average across all measurement locations, both types of user devices[4] and both East and West measurement servers of the downstream and upstream TCP throughput. The following chart documents the change in upstream and downstream TCP throughput by carrier.

Verizon is the clear leader with continuous increases in both downstream and upstream throughput over the measurement period resulting in an over 3x increase in average downstream throughput and over 2x in upstream throughput. The three other carriers deliver both lower throughput in either direction, as well as lower rates of improvement - about 2x. Verizon appears to be increasing its performance margin over the other three carriers. Sprint delivers the lowest performance of any of the carriers throughout California.

3.2 Latency

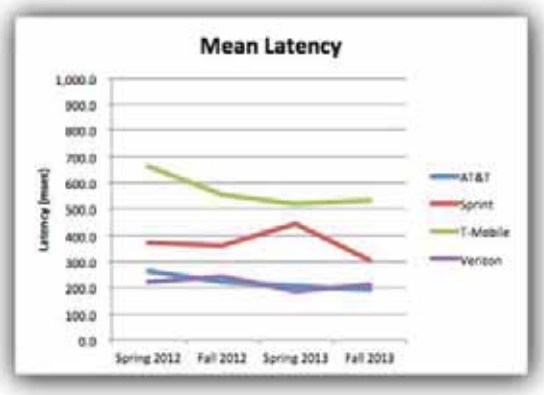

The analysis of overall average latency for each of the carriers shows a different story. The consistent trend is modestly decreasing latency over time. During the time period of these measurements, the AT&T network has slowly improved in latency (the mean latency improved from 260 to 190 milliseconds (msec), Verizon has made more modest improvements while both Sprint and T-Mobile have much higher latency that is slowly improving.

3.3 Jitter

Jitter, with the exception of AT&T, has remained essentially static for all carriers over the measurement period. The big decrease in mid-2012, as we will see in Section 7, is largely from dramatic improvements in jitter in AT&T's rural network.

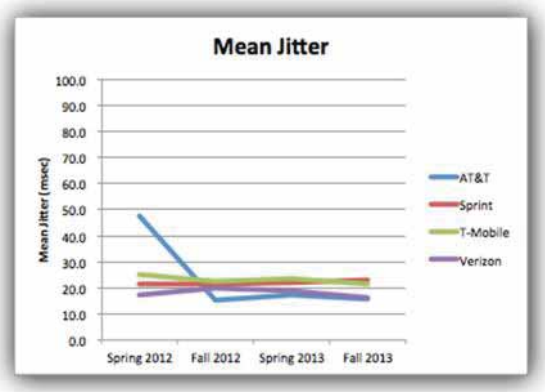

[4] There is a difference between user devices, but it appears to be unique to each device - not structural by technology or carrier. Not all user devices perform equally well.

4. Variation in Mobile Broadband Service

While we see dramatic improvements in average throughput over time, averages do not tell the complete story. The following graphs show a histogram of the measured TCP throughput across the sampled locations to both of the two geographic measurement servers. Much of the growth in average throughput has occurred by dramatic increases in the high performance tail of the distribution. A minority of locations get much better throughput, while the majority of locations have much more modest improvements in throughput. The wide variation in delivered throughput across the entire sample set is apparent. For example, it is possible (though uncommon) to get a downstream throughput for Verizon to a local (West) server that is 50 Mbps even though the Verizon state-wide mean is 17.5 and the median is 13.8 Mbps.

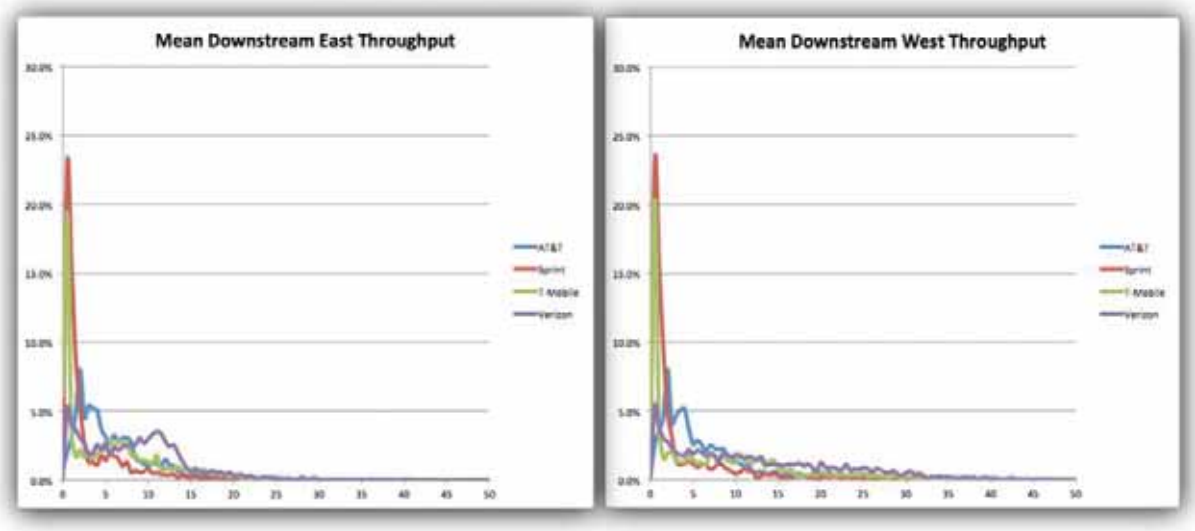

This variation in performance, echoed in other network metrics of upstream throughput, latency and jitter - is a composite of other, more fundamental variations. In order of importance these include:

- Location of user within California
- Choice of carrier
- Location of used server on the Internet
- Local variations
- Time of day.

4.1 Location of the User

The most important variation is location within the state of the user. The following interpolated kriging maps for downstream throughput for the four carriers illustrate this variation. Depending on the carrier there is almost a 25:1 variation between mean downstream performances based on where the user is in the state at the time of Internet access. Similar variations exist for upstream throughput, latency and jitter.

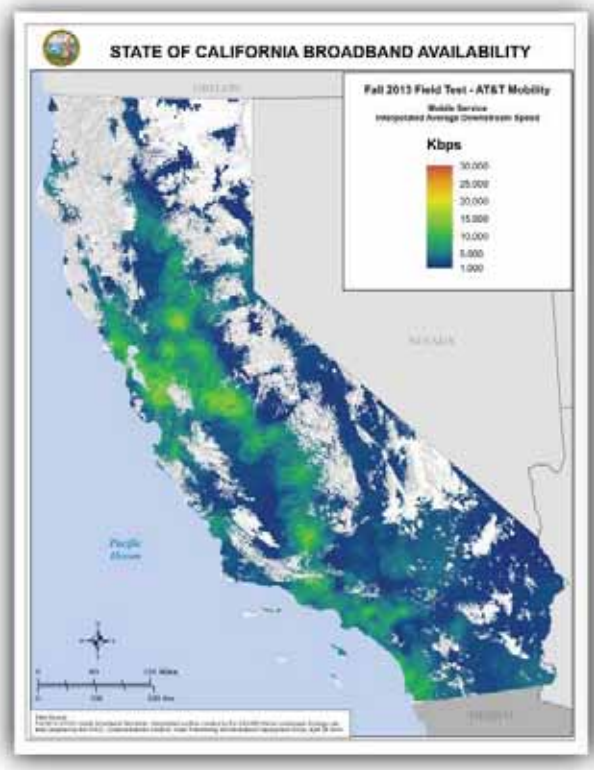
STATE OF CALIFORNIA BROADBAND AVAILABILITY
Fall 2013 Field Test - AT&T Mobility
Mobile Service
Interpolated Average Downstream Speed
Kbps
30,000
25,000
20,000
15,000
10,000
5,000
1,000
Pacific Ocean

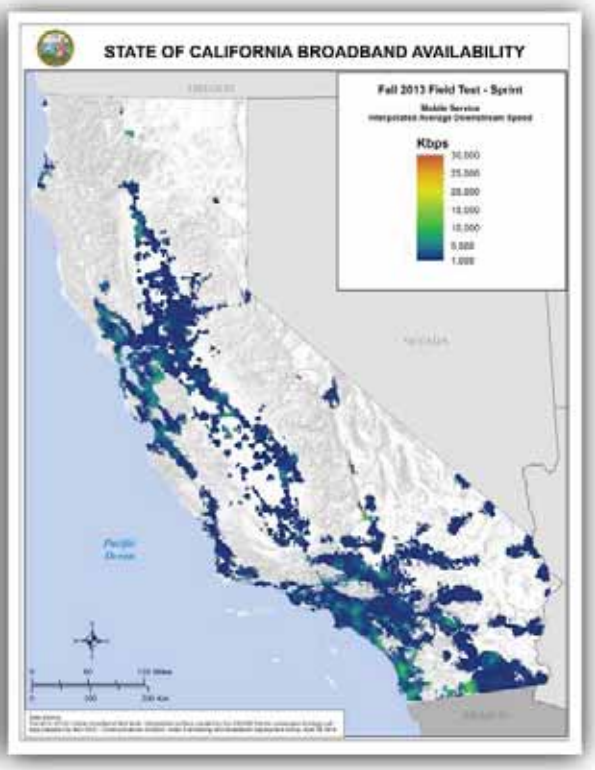
STATE OF CALIFORNIA BROADBAND AVAILABILITY
Fall 2013 Field Test - Sprint
Mobile Service
Interpolated Average Downstream Speed
Kbps
30,000
25,000
20,000
15,000
10,000
5,000
1,000
Pacific Ocean

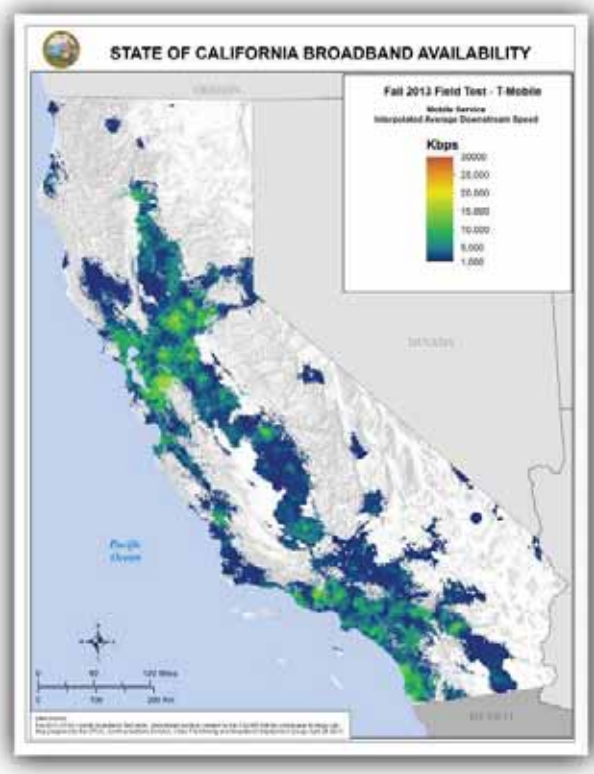
STATE OF CALIFORNIA BROADBAND AVAILABILITY
Fall 2013 Field Test - T-Mobile
Mobile Service
Interpolated Average Downstream Speed
Kbps
30000
25,000
20,000
15,000
10,000
5,000
1,000
Pacific Ocean

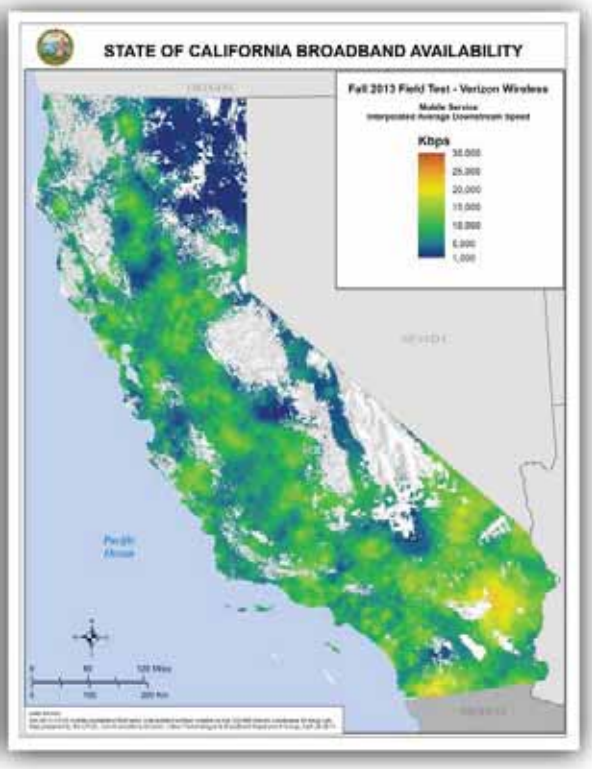
STATE OF CALIFORNIA BROADBAND AVAILABILITY
Fall 2013 Field Test - Verizon Wireless
Mobile Service
Interpolated Average Downstream Speed
Kbps
30,000
25,000
20,000
15,000
10,000
5,000
1,000
Pacific Ocean

4.2 Choice of Carrier

There is a wide variation between the service delivered by each carrier. We will examine this comparison between carriers in greater detail in Section 5. This variation is illustrated in the graphs below charting the overall mean downstream throughput for each carrier across the entire state. We can see a range of about 5:1 between the fastest and the slowest carriers in both upstream and downstream throughput.

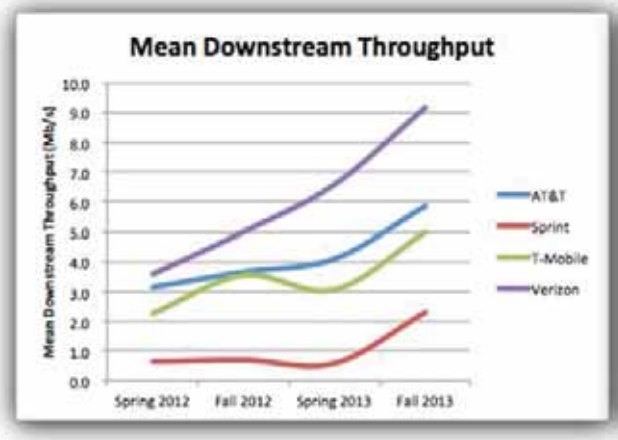

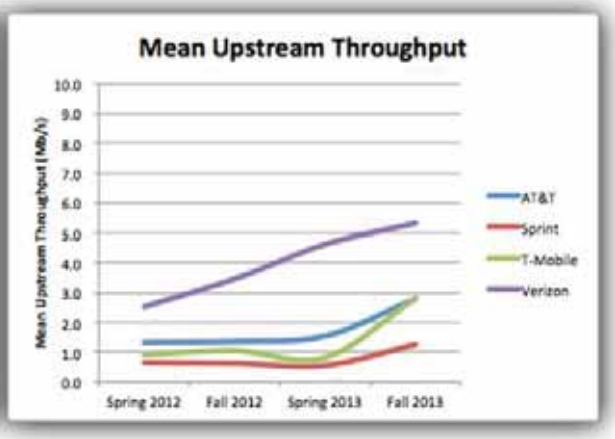

4.3 Location of the Server

The Internet backbone, not just the local wireless access network, has significant impact on user performance. We will examine this impact in more detail in Section 6. The graphs below illustrate the mean downstream TCP throughput to the West and East servers. The difference in mean throughput between the East and West servers is solely due to the impact of the Internet backbone connection strategy chosen by the carrier.

In Verizon's case in Fall 2013, this choice of backbone can result in an almost 50% performance difference between a California user accessing a server on the East Coast vs a server in California.

The data suggest that the effects of location get more pronounced as network performance increases.

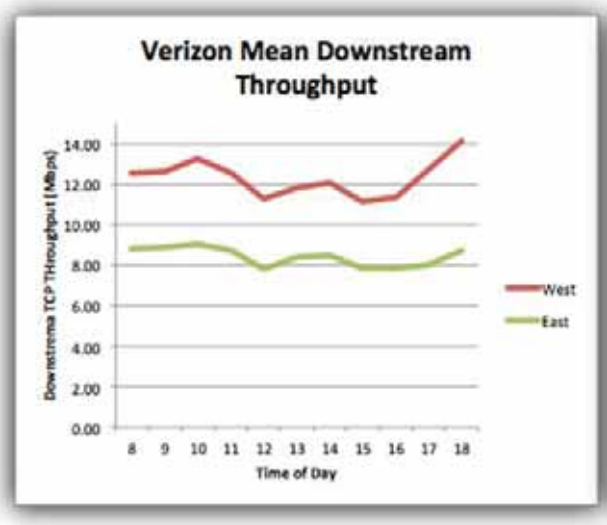

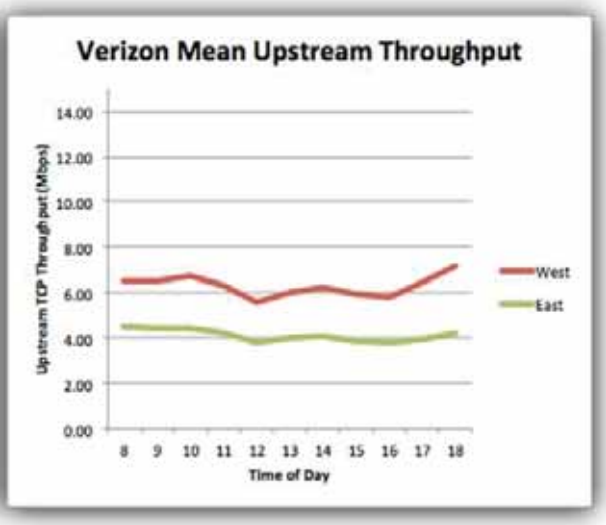

4.4 Local Variations

CalSPEED measures 40 separate TCP throughput measurements in both the upstream and downstream directions for each sample location, for each carrier over a period of about 30 minutes.

The CalSPEED analysis computes a standard deviation for the variation among these measurements for each test location - giving a metric for the local variation in throughput during the duration of the measurement. This local variation depends on carrier and location - as can be seen in the above chart. The horizontal axis plots the normalized standard deviation of downstream throughput during the CalSPEED measurements. The vertical axis plots the cumulative percentage (over the entire CalSPEED dataset for the Fall of 2013) showing that this much variation (or more) was measured. Curves that peak towards the upper left are representative of a more consistent and less variable service than curves that peak to the lower right. The mean normalized standard variation can vary from as little as 25% for an urban Verizon user to as much as over 100% for a rural T-Mobile user. Rural users can expect to see much more variability in their service than urban users.

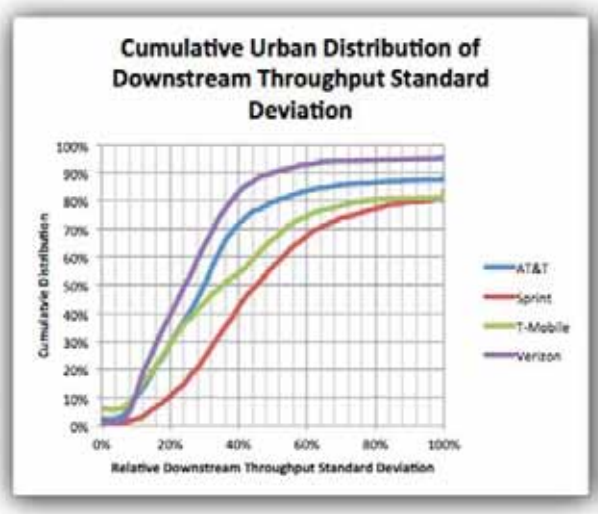

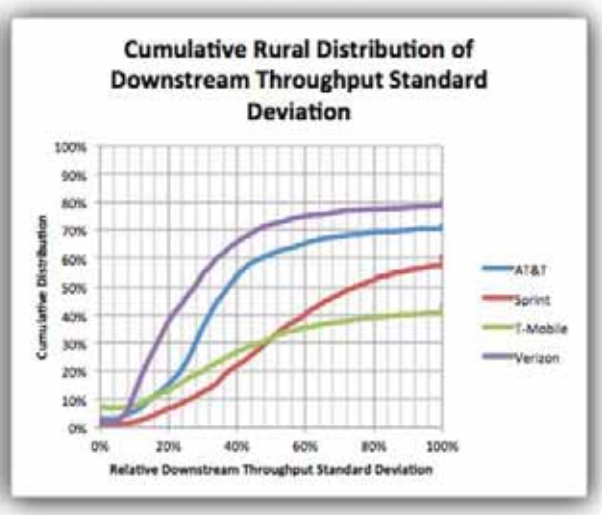

4.5 Time of Day

The least important variation is by time of day. Each CalSPEED measurement records the time of day of the measurement. As the chart below demonstrates for Verizon, the mean downstream throughput shows some variation with time of day, but the variation is on the order of 10% - much smaller than the other sources of variation.

All the carriers show a similar pattern of largely constant average throughput during the day, with a decrease near 4PM each day followed by an increase in mean throughput towards evening.

Our measurements are limited by our choice to only collect data during daylight hours in consideration of the safety of our field teams.

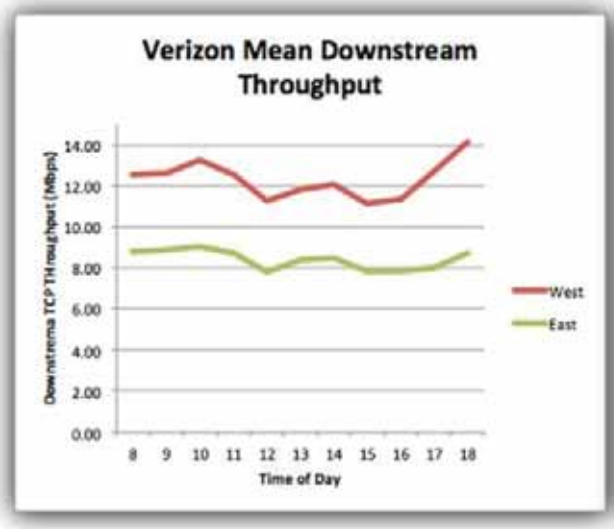

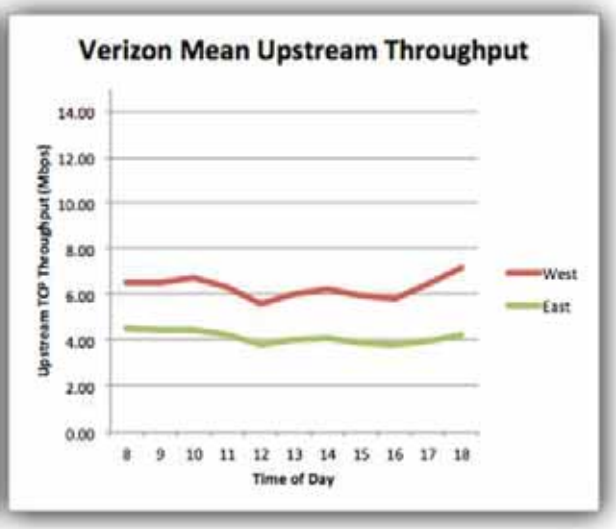

5. All Carriers are Not Equal

There is a dramatic and material difference between the mobile broadband service delivered by each of the four carriers. Let's look at the distinguishing issues of coverage, throughput, latency, jitter, demographic coverage between urban, rural and tribal and finish with differences in deployed wireless access technology.

5.1 Coverage

CalSPEED assesses coverage in several ways. First, let's look at the physical coverage of the state with some measure of service. As the maps below indicate, there is a wide variation between how much of the physical area of the state is offered any service by each of the carriers. Verizon and AT&T offer coverage over substantial parts of the state, while T-Mobile offers coverage largely in urban areas and major highways.

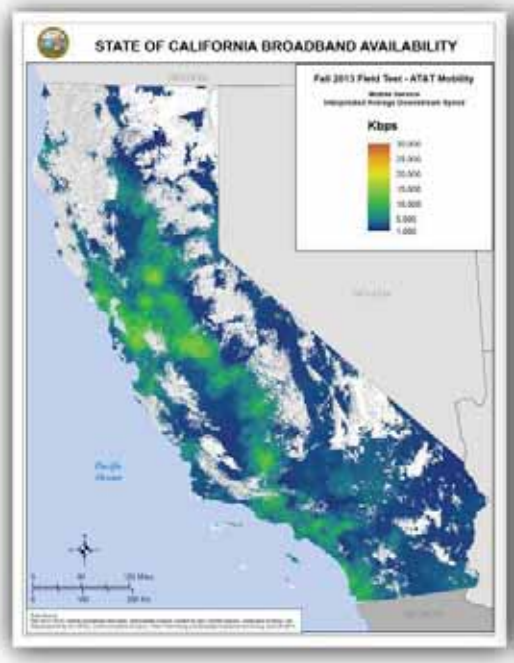
STATE OF CALIFORNIA BROADBAND AVAILABILITY
Fall 2013 Field Test - AT&T Mobility
Kbps

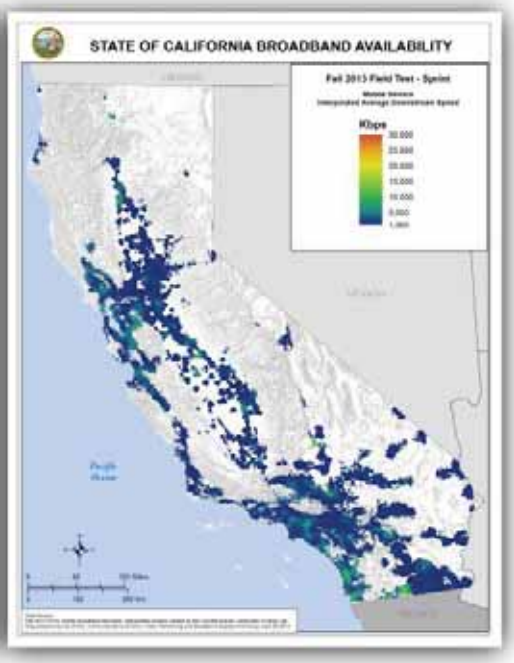
STATE OF CALIFORNIA BROADBAND AVAILABILITY
Fall 2013 Field Test - Sprint
Kbps

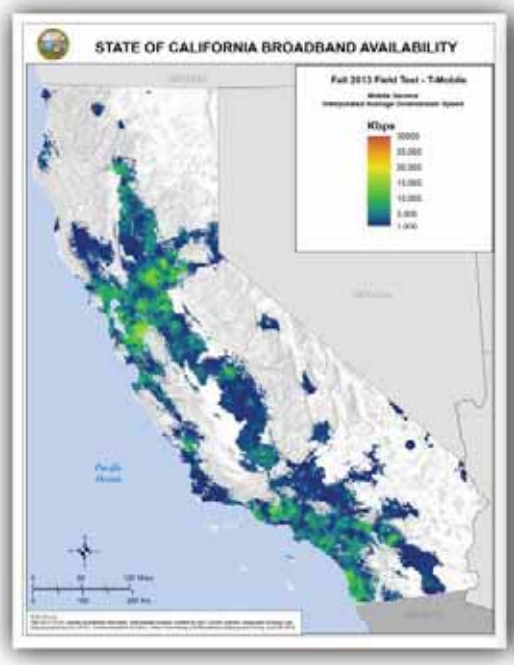
STATE OF CALIFORNIA BROADBAND AVAILABILITY
Fall 2013 Field Test - T-Mobile
Kbps

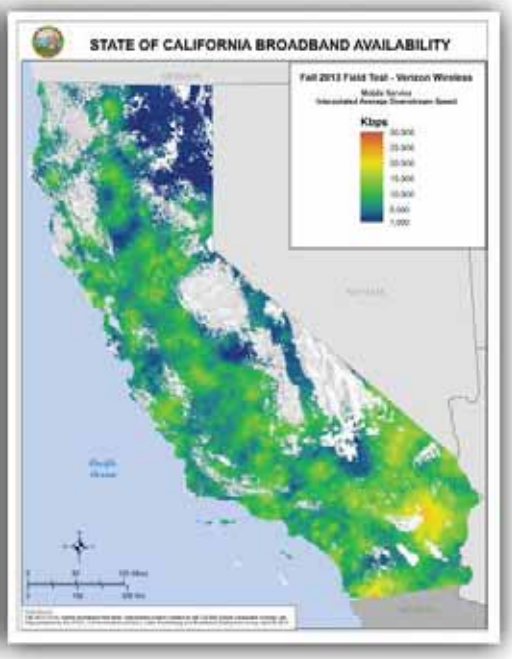
STATE OF CALIFORNIA BROADBAND AVAILABILITY
Fall 2013 Field Test - Verizon Wireless
Kbps

A second measure of comparative coverage is the quality of service offered within the announced service footprint of each carrier. The chart to the right documents the percentage of sample locations, within the announced service footprint of each carrier, that meets or exceeds the current California standard for sufficient broadband service or "served" broadband status requires speeds of at least 6 Mbps downstream AND 1.5 Mbps upstream. Areas that do not meet this standard are eligible for broadband infrastructure subsidies.

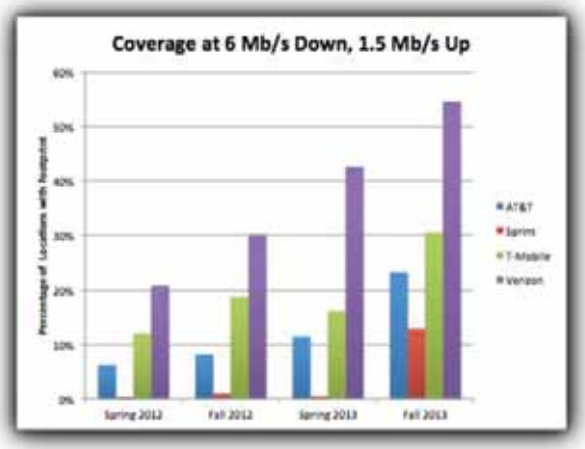

All carriers show consistent improvement over time, but Verizon clearly has the lead with not only an extensive footprint of coverage within California, but a quality of service within that footprint that delivers broadband service at California's "served" speed benchmark to over 50% of its California footprint.

5.2 Throughput

We can compute an overall average TCP throughput across all samples for each measurement round. The charts below graph this overall average by carrier over time for both downstream and upstream throughput.

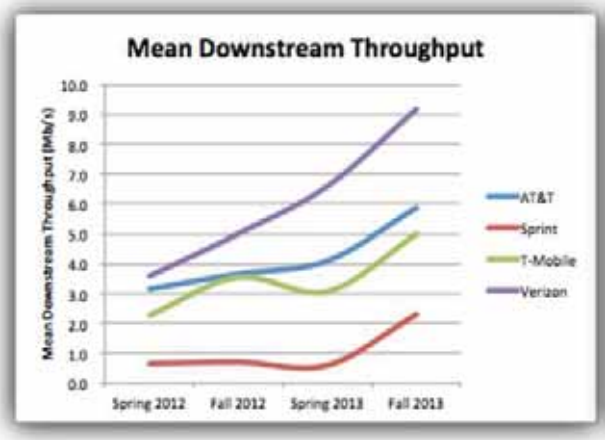

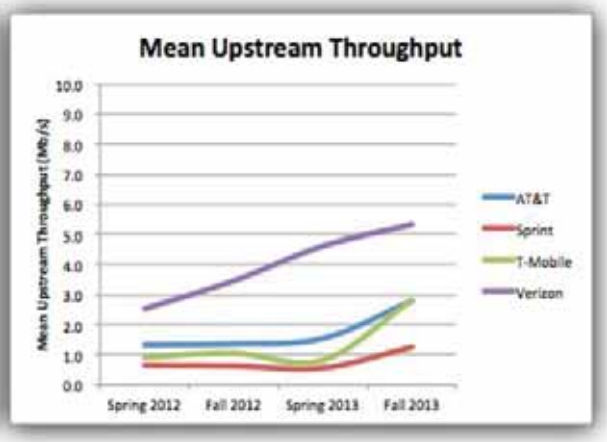

Downstream Throughput Statistics

	East		West	
	Mean	Median	Mean	Median
AT&T	8.7	7.2	10.7	8.2
Sprint	1.9	0.6	2.5	0.6
T-Mobile	7.5	8.0	9.7	8.0
Verizon	11.8	11.4	17.5	13.8

Upstream Throughput Statistics

	East		West	
	Mean	Median	Mean	Median
AT&T	3.4	4.7	4.5	4.7
Sprint	1.0	1.0	1.1	1.0
T-Mobile	2.6	1.8	4.0	1.8
Verizon	4.6	5.5	6.5	6.0

Looking at the Fall of 2013, we can see material differences between the carriers with Verizon having a commanding (and growing) lead in both downstream and upstream average throughput. This is almost 50% greater than the closest competitor, AT&T, in downstream throughput and over 400% greater than the lowest performing competitor, Sprint.

5.2 Latency

All mobile networks have shown a trend towards lower end-to-end latency despite the wide gap between the best and worst performing networks. Verizon improved its latency a modest 5% but AT&T improved 26% during the measurement period. With much greater average latency, both Sprint and T-Mobile improved about 20%.

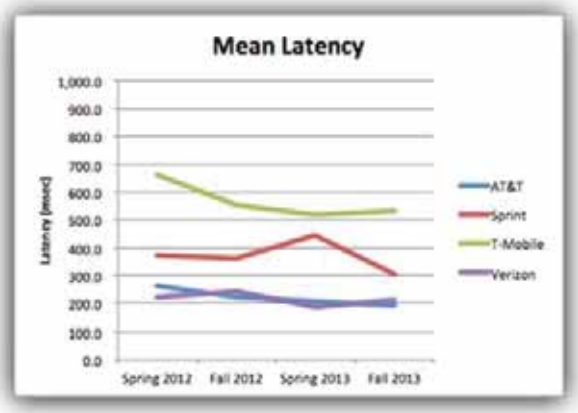

As with most other metrics, AT&T and Verizon are in the lead showing the lowest overall absolute average latency. As of Fall 2013, AT&T and Verizon both show average latencies less than half the worst performing network T-Mobile.

5.3 Jitter

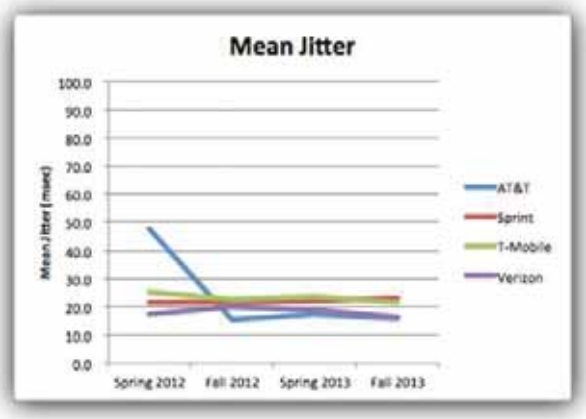

Overall, carriers have been trending towards lower average jitter over time. This was particularly noticeable for AT&T during 2012, and as we shall see later, that was primarily due to reducing jitter in the rural portion of their network.

AT&T, Verizon and T-Mobile's mobile networks trend towards lower end-to-end jitter. Sprint has actually modestly increased jitter over the measurement period by about 6%.

5.4 Service Across Urban, Rural and Tribal Areas

As we will see in Section 7, there is a wide variation between the major carriers on the quality of service between urban and rural and tribal demographics.

5.5 Wireless Access Technology

Each CalSPEED measurement records the type of local wireless data access network being used for this particular measurement, on this carrier, at the location at the time of the measurement. For the purposes of this analysis, we have grouped EVDO, HSPA and HSPA+ as 3G technologies. For this analysis, GPRS technology, delivering essentially < 10 Kbps dial-up network performance, is classified as a 1G network technology. EDGE and 1xRTT are classified as 2G, and only LTE is classified at 4G technology.

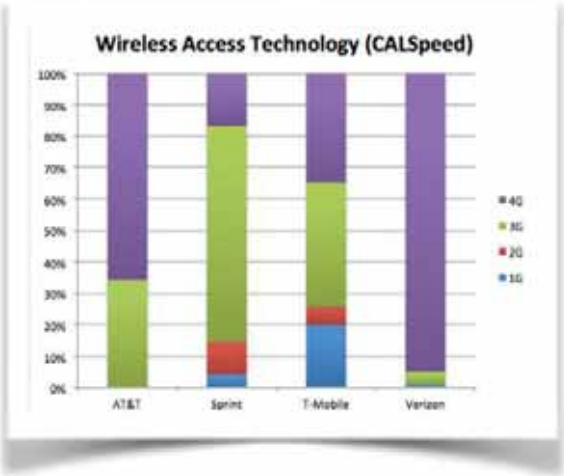

The figure to the right illustrates the dramatically different state of wireless access network technologies in use in the state by measuring the percentage of coverage of each carrier's wireless technologies. Verizon has a 95% 4G LTE footprint in contrast to Sprint, in particular, that has less than 20% 4G footprint. Its footprint is instead dominated by 3G CDMA technology. T-Mobile, in this sample, even has a significant legacy footprint of 1G GPRS service.

6. Mobile Broadband Service is not Just Local - The Internet Backbone Matters

Mobile network performance, as delivered to the end user, depends on a combination of the user's choice of handset, the radio access network of the wireless carrier and the wired backbone network of the chosen carrier that connects the carrier's radio network to the overall Internet - and thus ultimately to the chosen application service.

The recent Netflix issues with peering have pointed out that the choices of Internet peering have substantial impact on the end-to-end user performance experience.

CalSPEED data illustrates that backbone choice by each mobile carrier strongly influences the user experience.

6.1 Latency

We begin with latency. Latency is the measure of end to end round trip packet delay and is the foundation of network performance. Latency drives performance - both in raw throughput, but also in quality of service. In mobile broadband, latency is a composite between local wireless access network latency, the latency of the backhaul each carrier uses to deliver packets from the local access networks within their internal backbone network and the peering strategy each carrier uses to connect to the overall Internet and ultimately to the destination servers. We see the influence of all of these components in the measured latency.

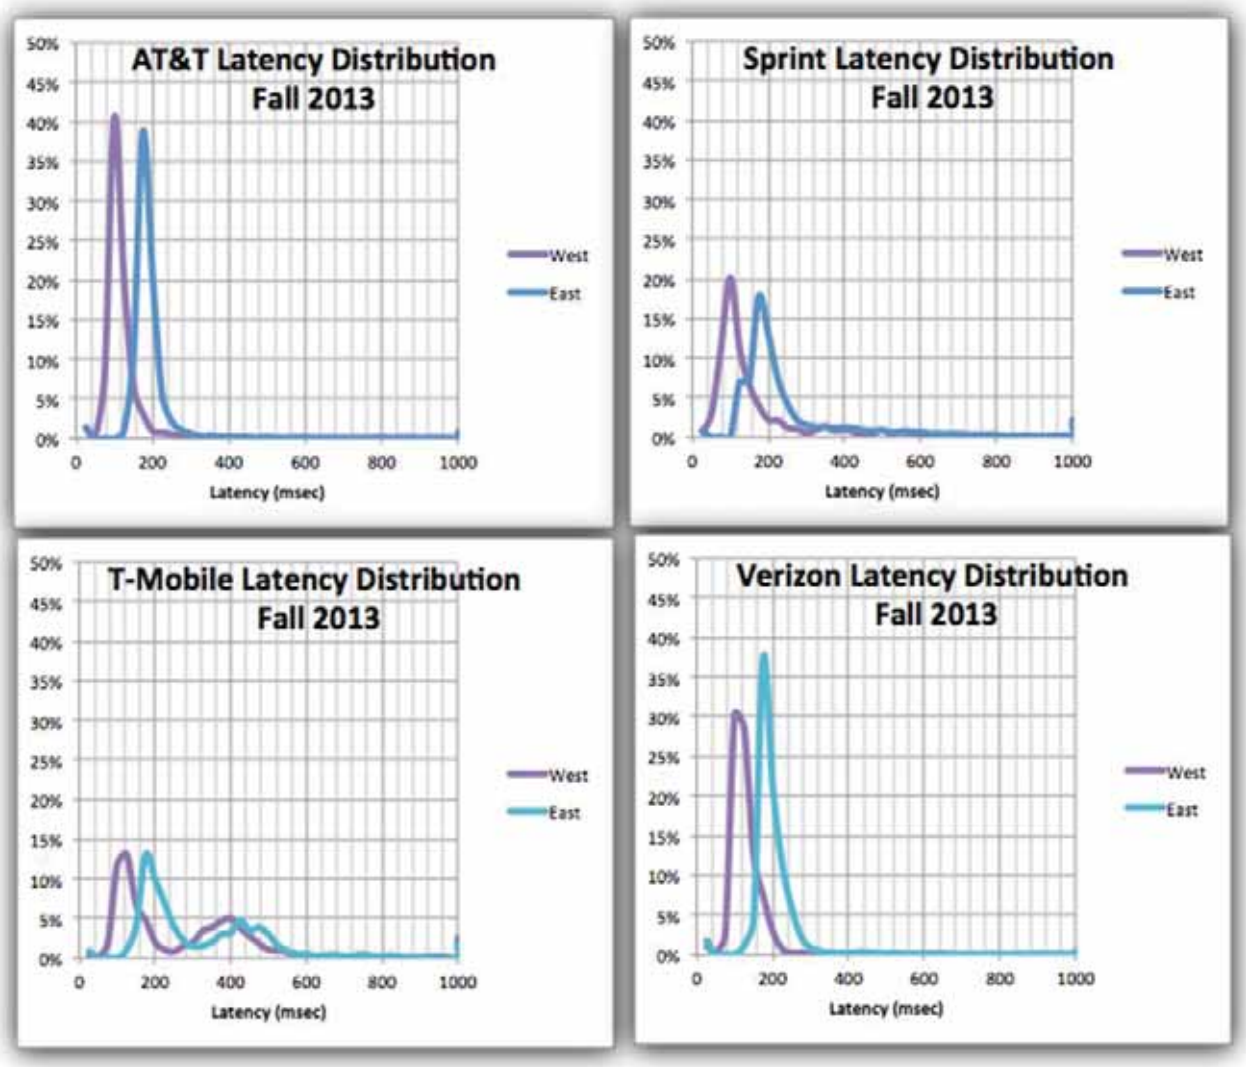

The above charts illustrate that CalSPEED measures an almost 2:1 difference in latency between East and West servers for all carriers. Typical latencies for the West server for all carriers were about 100 msec. Typical latencies to the East server were about 180-200 msec. This difference is the backhaul latency from servers in different locations in the Internet.

T-Mobile had a significant secondary population with much longer latency - of about 250 msec longer in latency to the same server. Since this secondary latency is present for both geographic servers, we hypothesize that it represents locations with higher local radio latency or older radio access technology for a subset of the measurement locations rather than an artifact of backhaul. And we can see that the tabulation of wireless access technology by carrier in Section 5.5 supports this hypothesis.

Sprint also has a significant secondary population with about 90 msec longer latency for the same server. Like T-Mobile, we believe this is an artifact of differing radio access subnetwork latency in a subset of the locations in our sample.

For modern 4G wireless access networks, user measured latency is about double when accessing services on the “far side” of the Internet.

6.2 Jitter

Jitter is the measure of variance in latency. It is critically important for real-time streaming media applications such as voice and video. CalSPEED measures jitter on the UDP measurement data streams. Jitter is a key metric of the ability of a network to sustain high quality streaming media. The graph below shows the jitter distribution for each carrier to the East and West servers in Fall 2013.

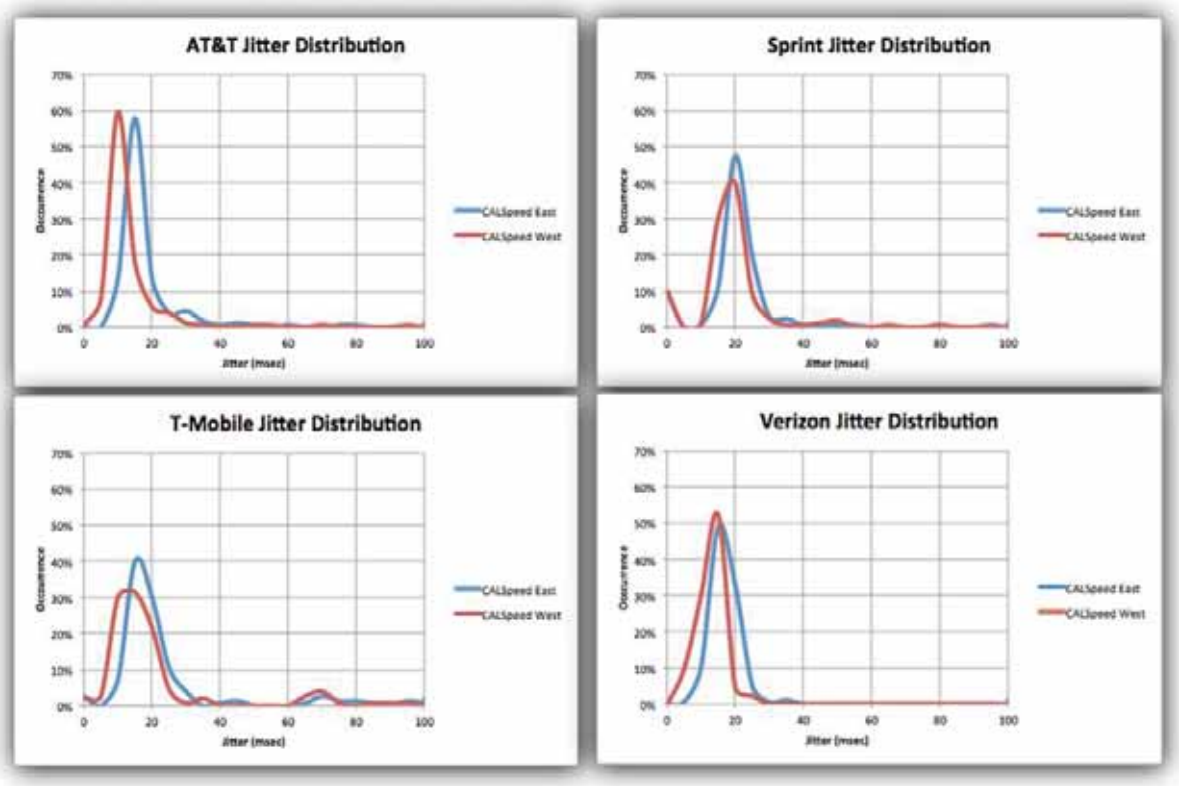

Each carrier shows a unique jitter distribution that reflects unique backhaul, routing and network

management choices. AT&T shows distinctly different jitter distributions for East and West servers with clear concentrations at 10 msec for the West server and 16 msec for the East server. Sprint shows little difference between servers - and shows the highest jitter of all the carriers at 20 msec. The jitter distribution for each carrier is distinct and clearly differs with location.

6.3 Packet Loss

Packet loss is the measurement of how often packets are lost and not delivered to their destination regardless of latency. CalSPEED measures packet loss to two locations in the Internet and both the absolute packet loss rates and the differences between the servers is of interest.

AT&T and Verizon have very low absolute packet loss rates in Fall 2013, with the loss to the East server marginally higher than to the West server. Both carriers had packet loss rates much less than 2% over the entire dataset.

Sprint and T-Mobile have much higher packet loss rates to both East and West servers. Packet loss rates to the closer West server are slightly smaller than to the East server.

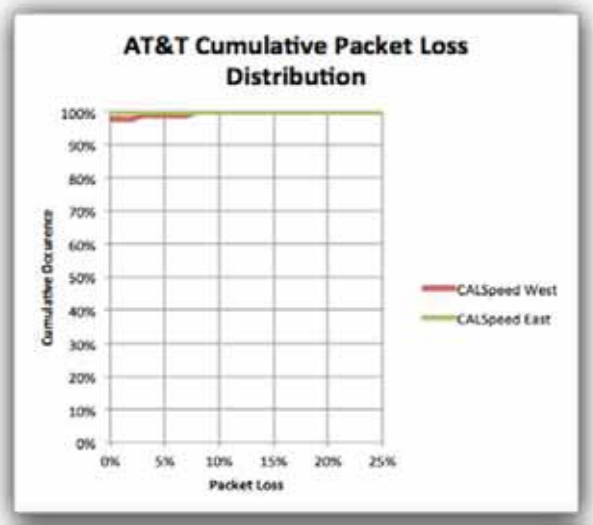

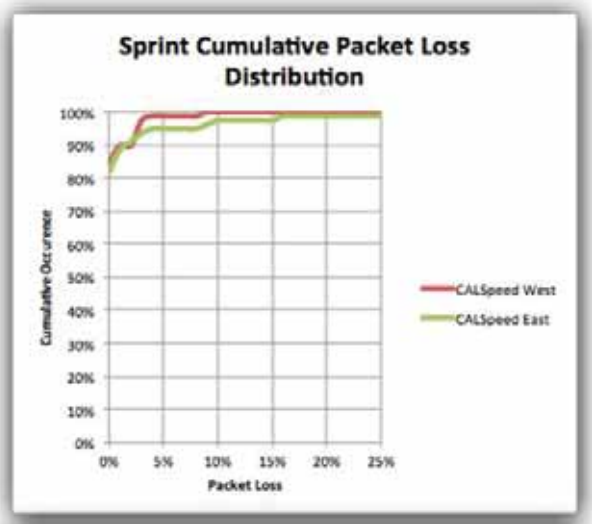

6.4 Downstream Throughput

TCP is the foundational data delivery protocol of the Internet and is used to reliably deliver data streams. CalSPEED focuses on measuring TCP throughput both downstream (from the server to the user device) and upstream (from the user device to the server). TCP throughput is strongly affected by the latency of the network and longer latency times for acknowledgements will slow a TCP connection down.

The following graph illustrates that for each carrier, higher downstream throughput is more commonly delivered from the closer West sever than from the East server.

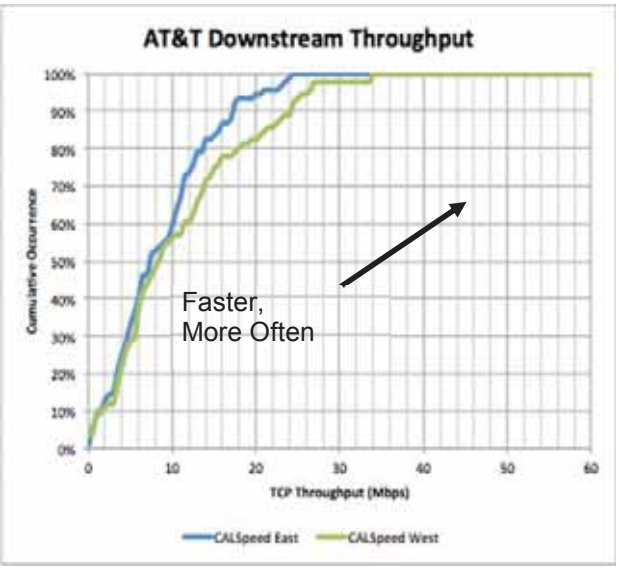

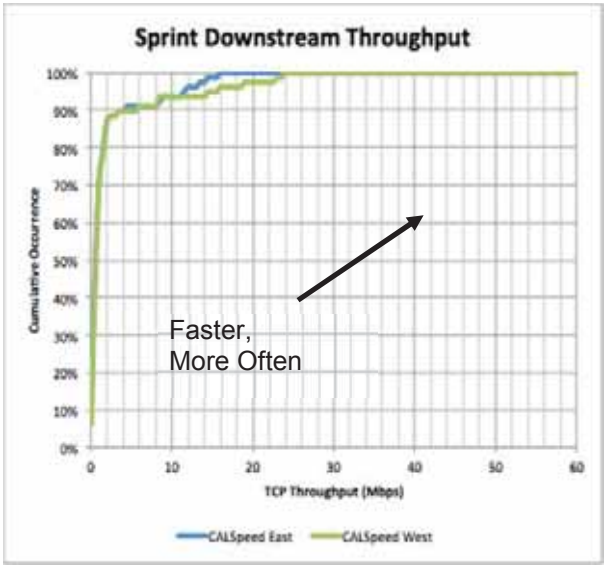

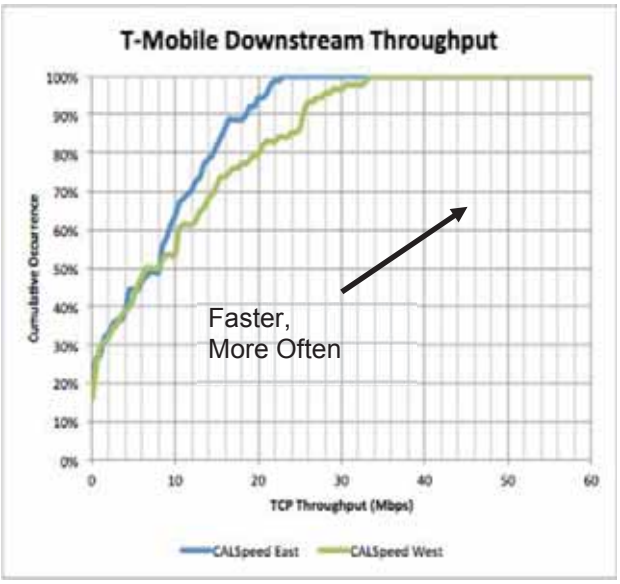

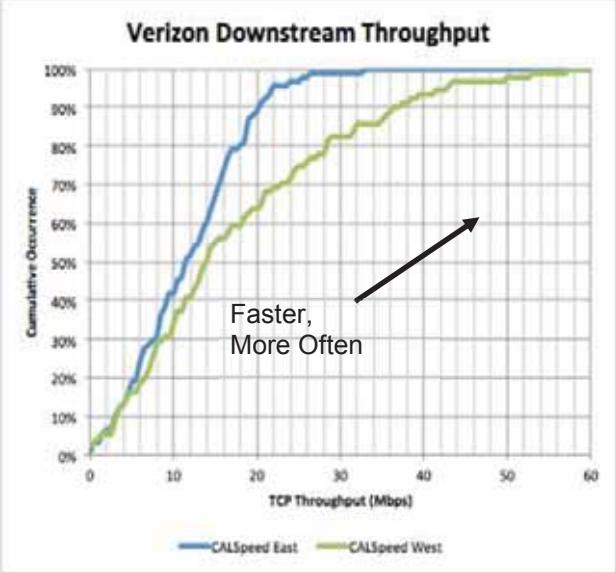

Further, it appears that as the capacity and performance of the network increases the difference due to location of the server - look at the sequence of Sprint, AT&T, T-Mobile and Verizon. For all these carriers, throughputs exceeding 20 Mbps to the East server are quite rare.

While Sprint is least affected by distance (and is also the lowest performing carrier), AT&T, T-Mobile and Verizon commonly see improvement in downstream throughput between 20-50%, paralleling decreases in latency of 50% or more, as traffic moves from the more distant East server to the closer West server. TCP's performance degrades with increasing latency but improves with decreased latency.

6.5 Upstream Throughput

We see the same effect of latency (and Internet distance) on upstream throughput. There is a distinct difference between throughput to the East and West servers, again with strong divergence above the median. The following graphs plot the cumulative frequency distribution on upstream TCP throughput.

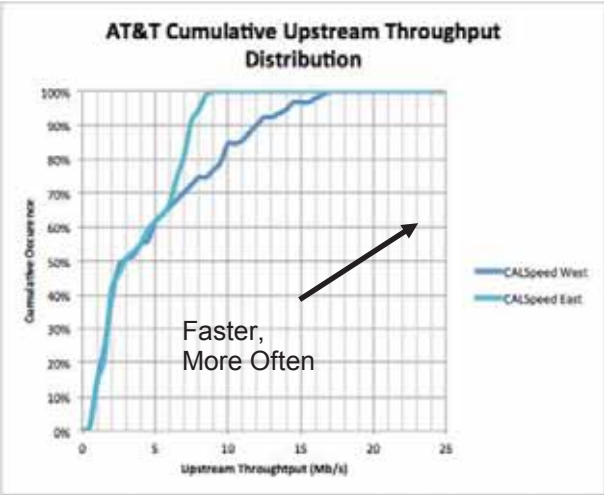

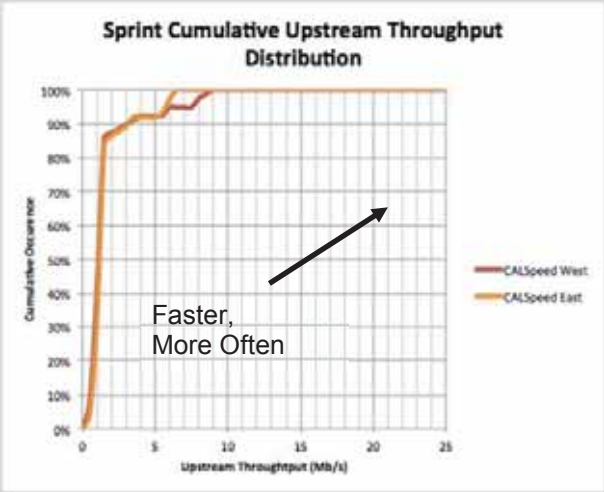

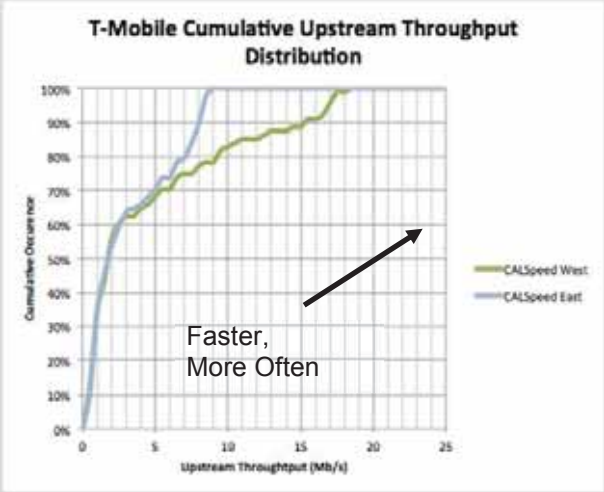

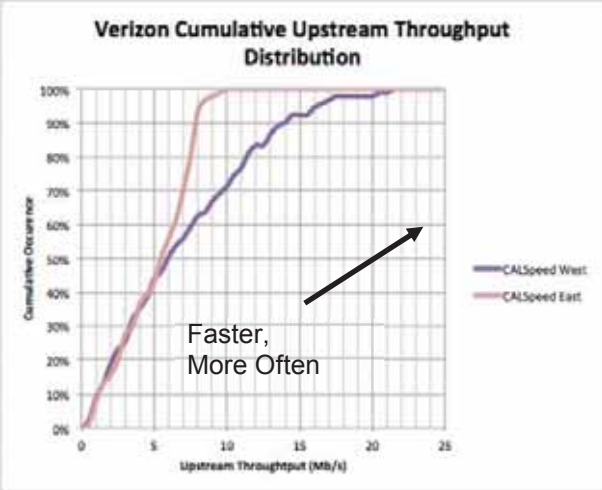

Note the strong low pass filter to the East server for > 9 Mbps upstream for all carriers[5].

Location (and thus differing latency) makes a difference in upstream throughput. As with downstream throughput, Sprint shows the least effect of distance (and also the lowest upstream throughput service) while Verizon, AT&T and T-Mobile show effects ranging from increases of 10-60% in upstream throughput as service is directed to a closer West server than a more distant East server.

7.0 The Growing Mobile Digital Divide

The CalSPEED data suggests a new, and growing, digital divide for mobile broadband - between urban, rural and tribal demographics. Let's look at how the CalSPEED data assesses the key measurements of coverage, throughput, reliability, streaming media capability and wireless technology between urban, rural and tribal areas. In each of these metrics, the CalSPEED data indicates that mobile broadband service for rural and tribal areas in California is materially less robust than mobile service for urban areas.

7.1 Reduced Broadband Coverage

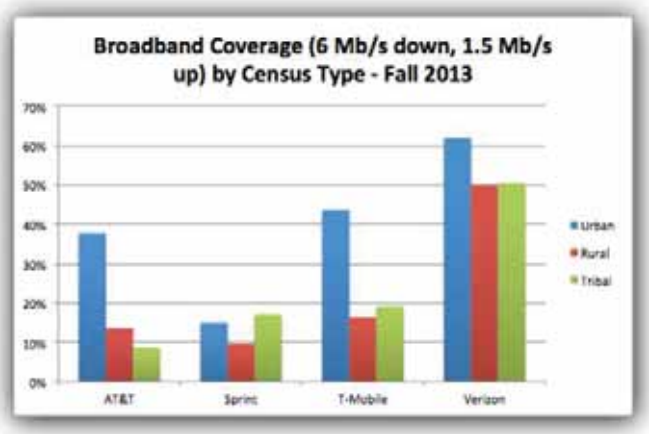

The standard in California for adequate or "served" broadband service status requires downstream speeds of at least 6 Mbps <u>and</u> upstream speeds of at least 1.5 Mbps. CalSPEED data shows significant differences in the scale of coverage to this standard exist among urban, rural and tribal areas and among the various carriers. Of the four carriers, only Verizon offers a similar level of broadband service among all census areas - roughly about 50% of all sampled locations in all demographics offer service at a level California considers as "served." AT&T and T-Mobile have substantial differences in the scale of broadband coverage between urban as compared to rural and tribal.

5 We have examined our methodology and we do see greater than 9 Mbps upstream service and greater than 20 Mbps downstream service to the East server - particularly thru WiFi connections to a wired access network. We therefore believe this is not a test artifact but rather measured performance in the field for mobile broadband.

7.2 Lower Average Throughput

Let's now examine the mean downstream throughput by carrier by urban, rural and tribal areas. As we can see in in the graph below, only Verizon has upgraded its California mobile network to offer similar downstream throughput, on average, in all census areas. Verizon's service with regard to throughput appears to be converging between urban and rural/tribal demographics over time. AT&T and T-Mobile mobile networks have dramatically different (and lower) downstream throughput in rural and tribal areas than in the urban areas - and that difference appears to be increasing in time. Sprint's service is very much lower in all demographics than the other three carriers.

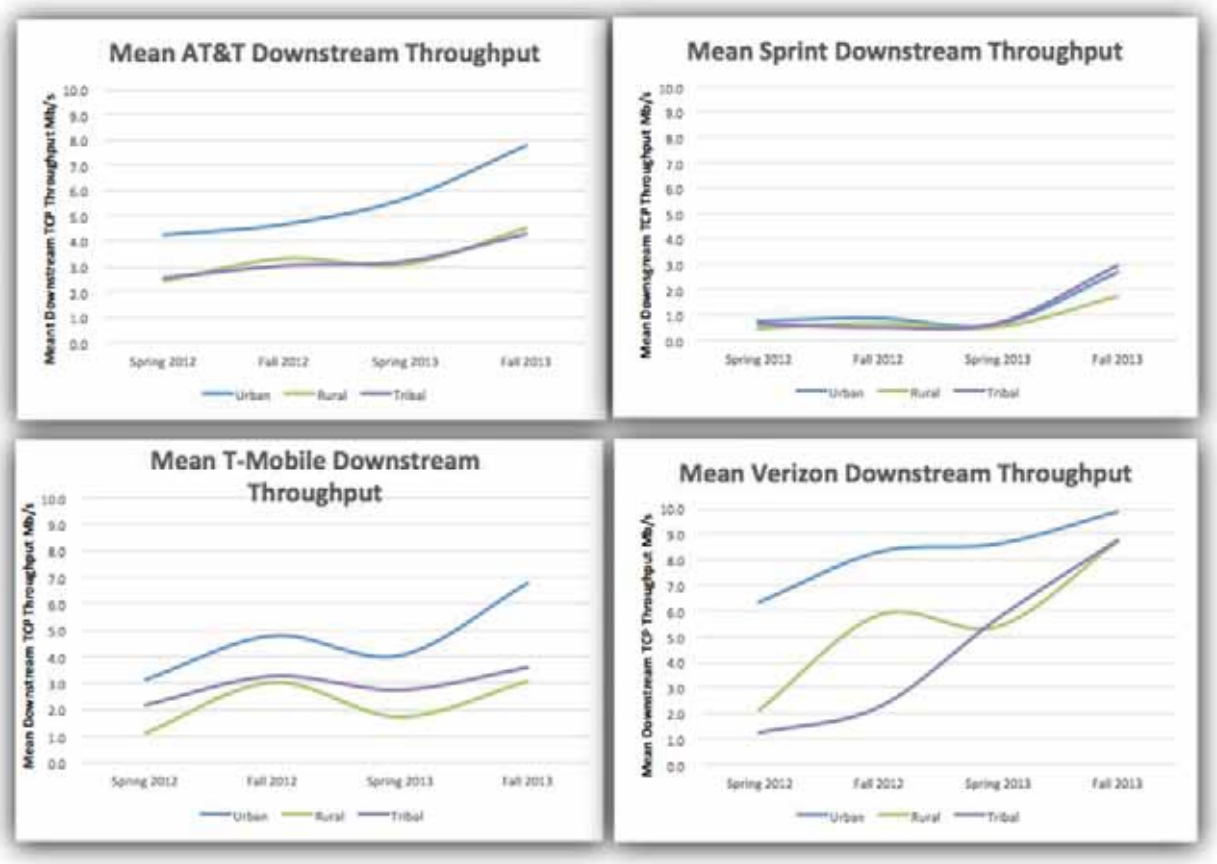

7.3 Lower Quality of Service

Key measurements of quality of service, in addition to raw throughput, are throughput variance, latency and jitter. As we can see from the following graphs, quality of service for rural and tribal users is often materially worse than for urban users.

Rural demographics are delivered not only a lower average throughput, but the service delivered has a materially higher variability - that is, sequential user requests for a TCP connection will vary much more in rural demographics than in urban demographics. We can see that in the following chart for

Fall 2013. The horizontal axis plots the normalized standard deviation of downstream throughput during the CalSPEED measurements. The vertical axis plots the cumulative percentage (over the entire CalSPEED dataset for the Fall of 2013) showing that this much variation (or more) was measured. Curves that peak towards the upper left are representative of a more consistent and less variable service than curves that peak to the lower right.

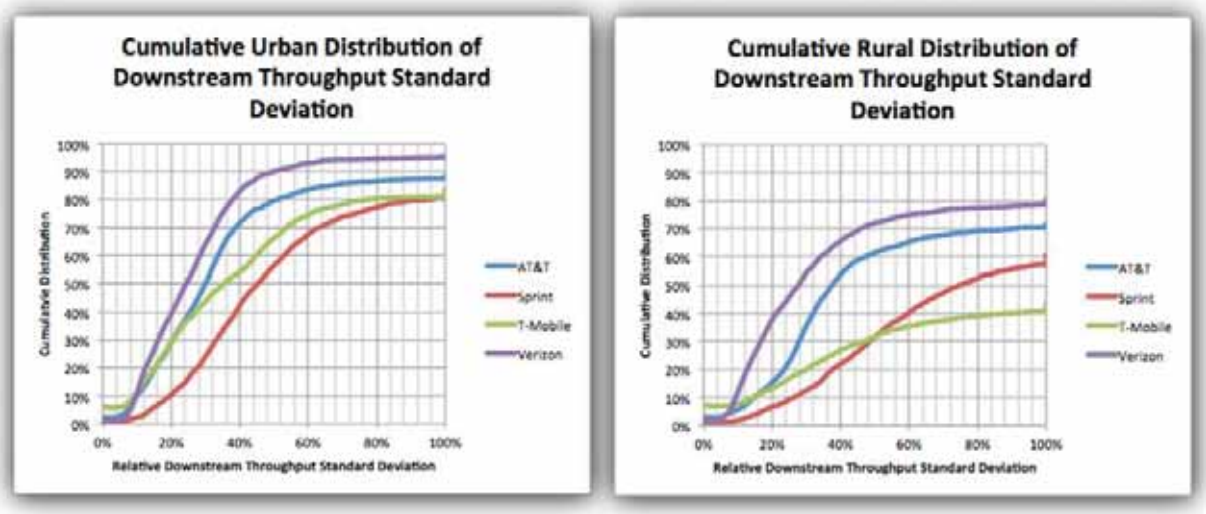

The graph on the left is the cumulative distribution for urban demographics of the standard deviation of local TCP throughput. Sprint provides the worst service with a median downstream throughput relative standard deviation of about 45%. In the rural demographic, this variation increases to 75%.

Latency shows a pattern of urban demographics having lower (better) latency than both tribal and rural demographics. The following charts graph latency for each carrier over time in these demographics.

Note that while Verizon has shown a convergence of throughput among the three demographics, the CalSPEED data indicates that latency for tribal users still remains substantially higher than rural and urban users.

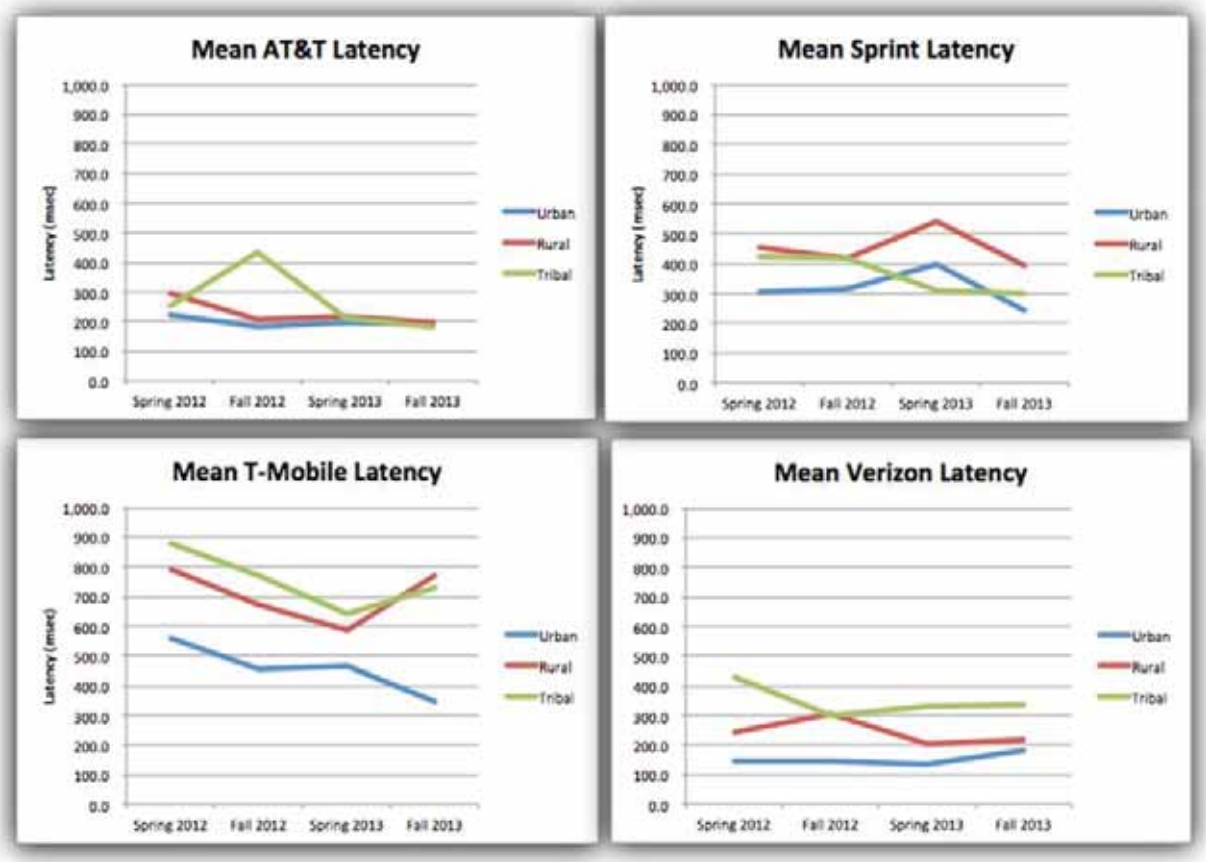

While average latency for all carriers is generally decreasing, in all demographics, urban latency tends to be less than rural or tribal latency.

Jitter is less impacted by demographics but in general urban demographics have lower (better) jitter. Jitter is key to good real-time streaming (voice and video) performance.

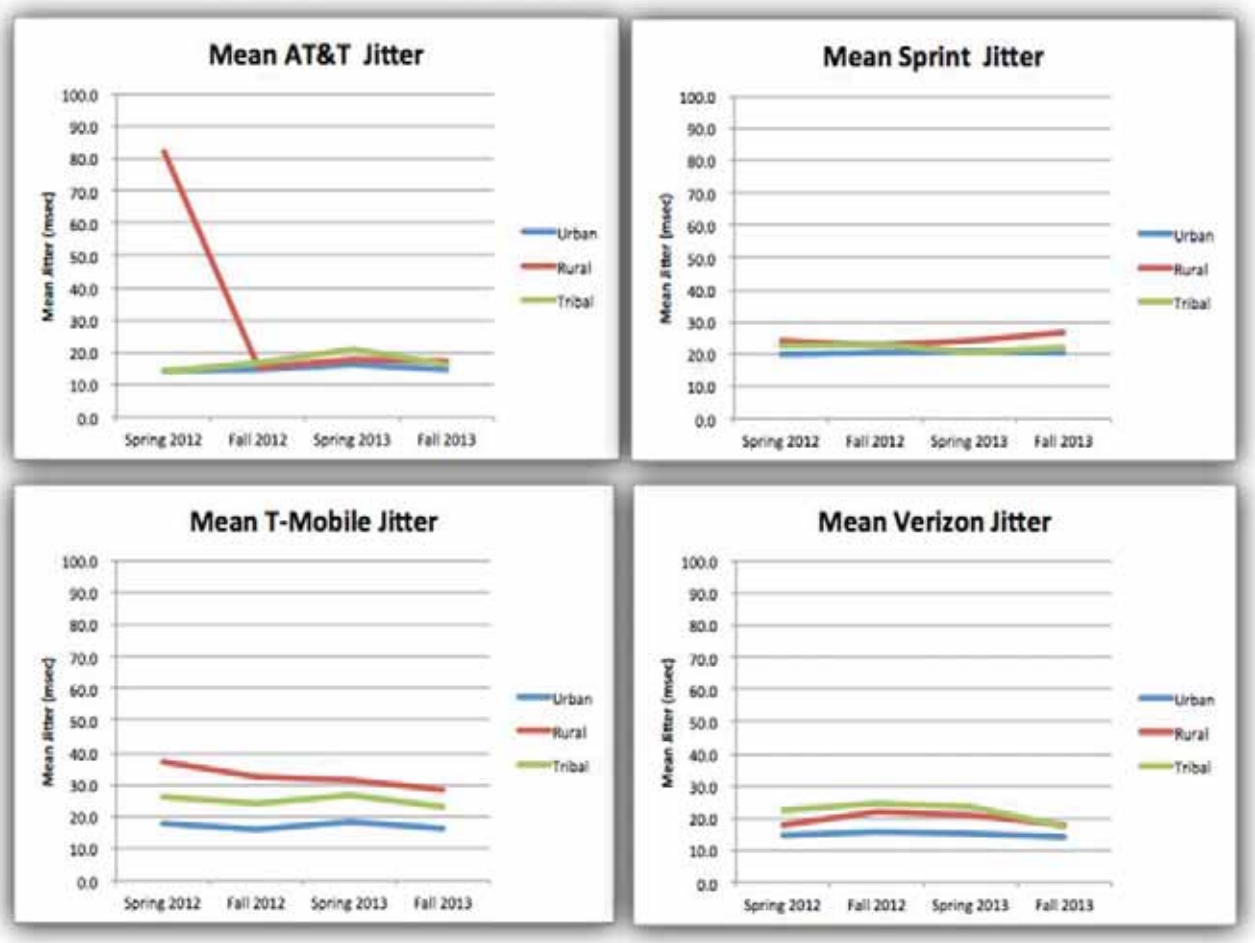

7.4 Lower Service Reliability

The CalSPEED measurement makes a number of TCP connections between client and the several geographic measurement servers during the course of the measurement. Some of these connections fail for a number of reasons, such as lack of coverage, variable wireless connectivity, TCP errors, packet loss -. This failure is not uncommon in TCP networks - each Internet user has had the experience of clicking on a web link that stalls and has to be clicked again to restart the TCP connection to get to the destination web page. The prevalence of connection failures is a measurement of the reliability of the network in delivering quality service. We can examine the percentage of attempted TCP connections that fail as a measurement of reliability.

In the above charts two themes can be seen. First, there is a variation between carriers of service reliability. Second, there is an almost 2:1 variation in service reliability between each demographic. For AT&T and Verizon, there is a substantial increase in connection failure rate from the roughly 7% (for AT&T and Verizon) in urban to about 15%+ in rural and tribal demographics.

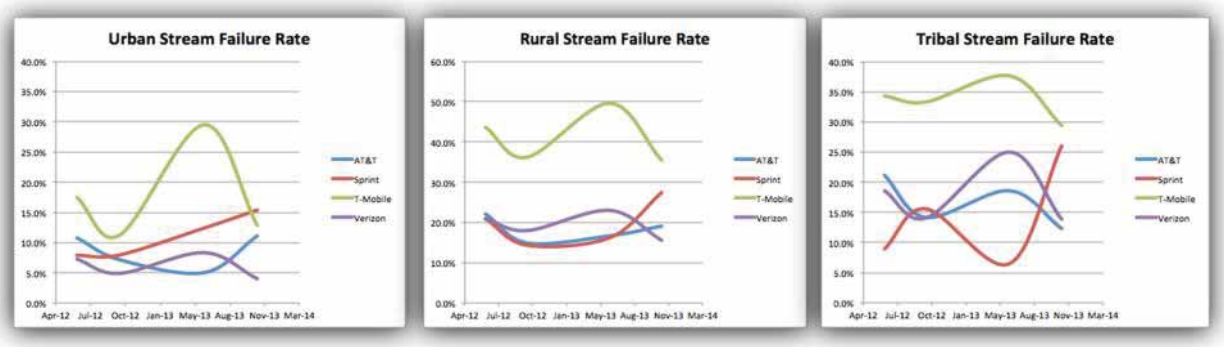

7.5 Less Streaming Media Capability

The CalSPEED analysis computes a synthetic Mean Opinion Score (MOS) as a metric for quality of streaming media, particularly VoIP. For all carriers, even AT&T and Verizon, which have the largest LTE footprints in California, a significant percentage of each carrier's entire network is not VoIP capable[6].

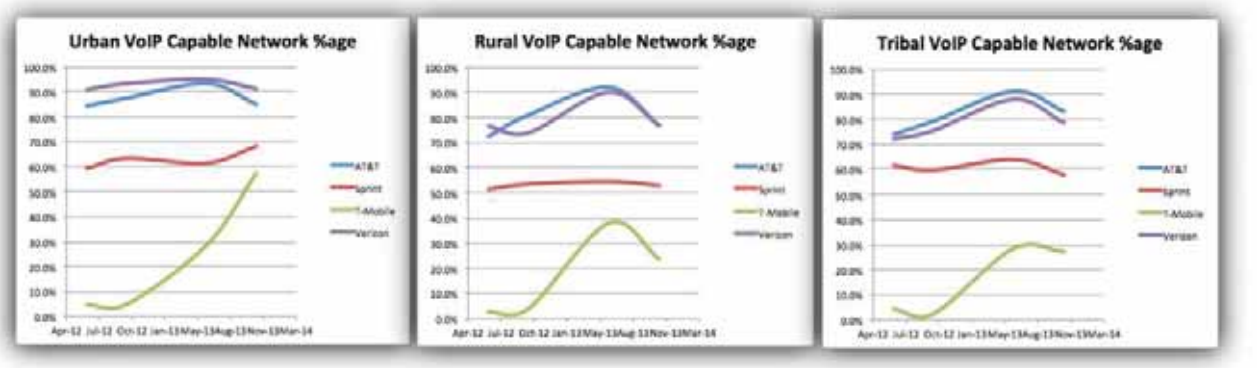

As the charts above graph, both Verizon and AT&T decrease from roughly 90% VoIP capable in urban demographics to 80% VoIP capable in rural and tribal demographics.

Both Sprint and T-Mobile show much lower VoIP capability in all demographics.

7.6 Obsolete Wireless Technology

One of the key drivers of mobile performance and service is the radio access technology deployed from each tower, in addition to the backbone connection of that tower by each carrier peered into the overall Internet. Over the last 5 years, there has been astonishing upgrade in radio access technology from a largely 2G network in 2009, to a dramatically increasing footprint of 4G LTE. One way of assessing the mobile capabilities is the distribution of generations of radio access technology over time, census type and among carriers.

Let's look over time, first at the proportion of sampling locations with 1G[7] and 2G radio connections during the measurement, and then the proportion of sampling locations with an LTE connection. As the above chart illustrates, there is a general trend over time, and across carriers, for fewer connections using legacy 1G and 2G radio access networks and more service using LTE radio access networks. However, there is a dramatic difference between carriers - with Sprint and T-Mobile still offering 10-15% of sampled locations in the Fall of 2013 with 1G/2G radio access. It is only within the last year that LTE was substantially deployed by carriers other than Verizon.

---

[6] We consider that a MOS score of greater than or equal to 4 is required to be VoIP capable.
[7] We have found not only substantial legacy prevalence of 2G service such as CDMA and Edge, but also some remaining 1G GPRS service!

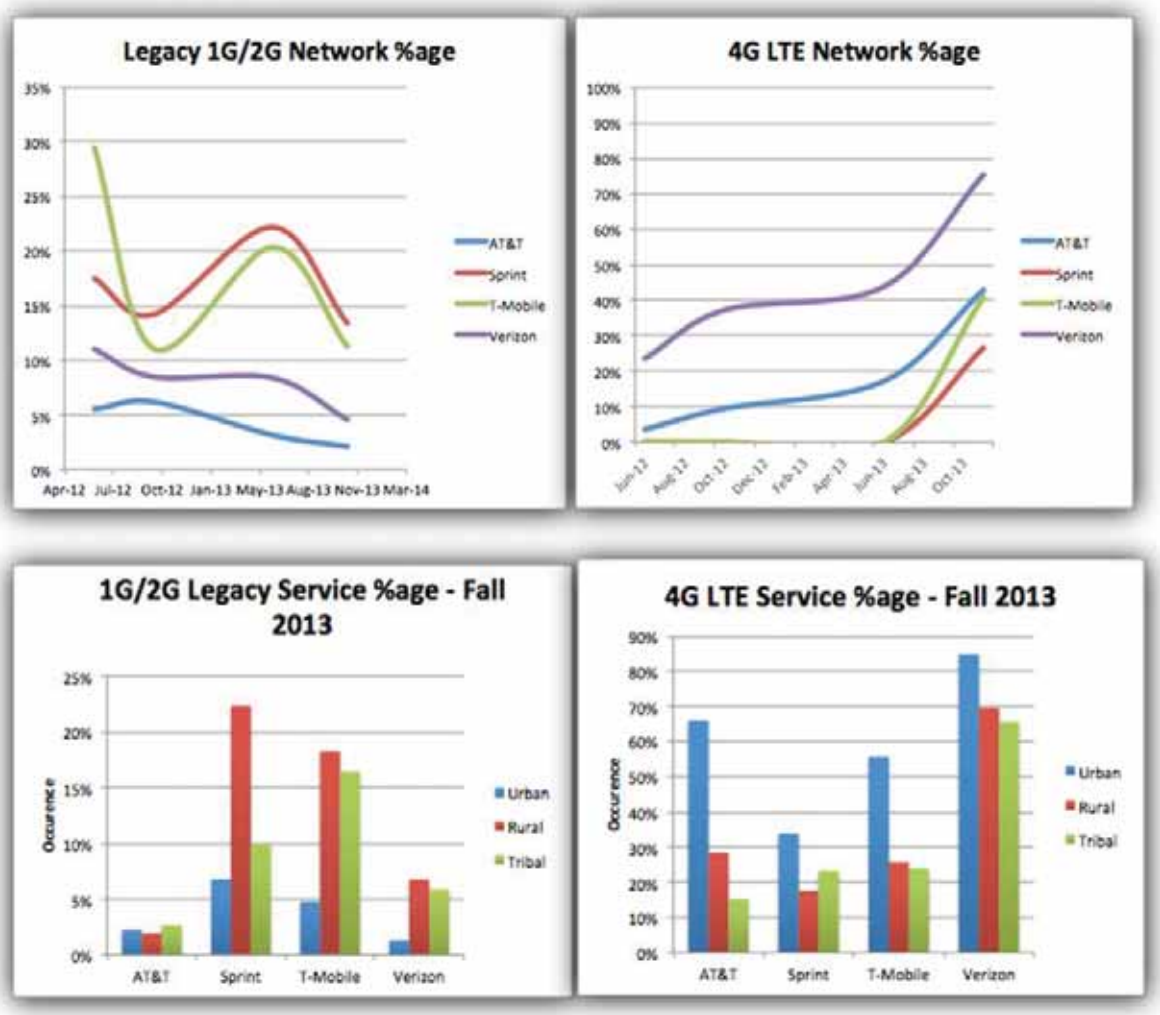

This persistence of legacy wireless access technology is biased towards rural and tribal demographics. As the above charts illustrate, carriers (particularly Sprint and T-Mobile), appear to have comparatively larger usage of legacy networks in rural and tribal demographics and a correspondingly low usage of 4G LTE wireless technology outside of urban areas.

8. VoIP Over Mobile Not Universally Available

Streaming data service - voice and video - are some of the newest and most important services to be carried by the Internet. CalSPEED estimates the quality of streaming audio, the foundation of VoIP, by computing latency, jitter and packet loss for a simulated UDP data stream at a typical encoded VoIP streaming rate of 88 kb/s. For this calculation we average measurements between East and West servers since the location of the voice destination could be anywhere on the Internet.

MOS - Mean Opinion Score - is a synthetic metric that integrates packet latency, packet loss and jitter as a measurement of voice quality. A value of equal to or greater than 4 indicates an acceptable level of voice quality. The following chart documents the percentage of each carrier's footprint that is MOS ready as of Fall 2013.

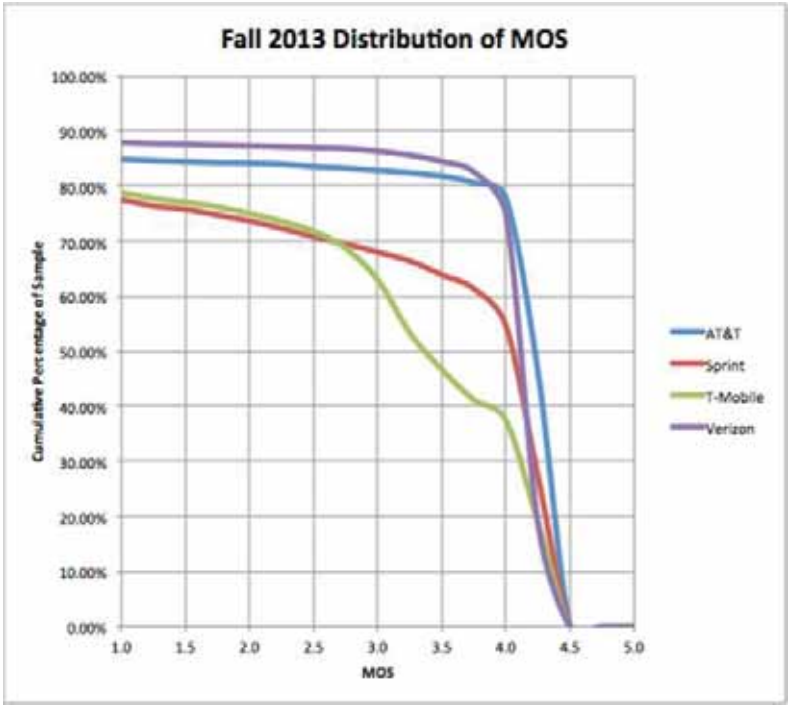

This difference among carriers in MOS capability is exacerbated by the dramatic differences in coverage footprint among the carriers. CalSPEED analysis provides an interpolated kriging map of the estimated VoIP capable MOS service for each carrier as of the Fall 2013.

Red color denotes interpolated VoIP capable coverage for each of the carriers.

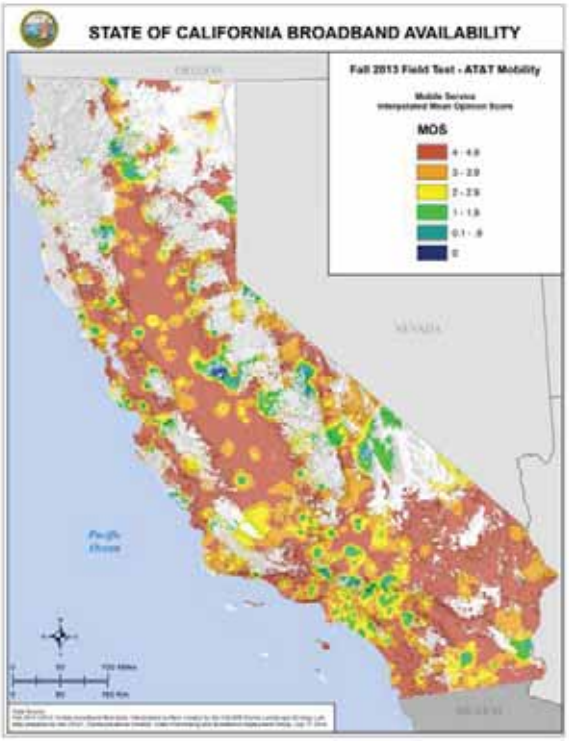
STATE OF CALIFORNIA BROADBAND AVAILABILITY
Fall 2013 Field Test - AT&T Mobility
Mobile Service
Interpolated Mean Opinion Score
MOS
4 - 4.9
3 - 3.9
2 - 2.9
1 - 1.9
0.1 - .9
0
Pacific Ocean

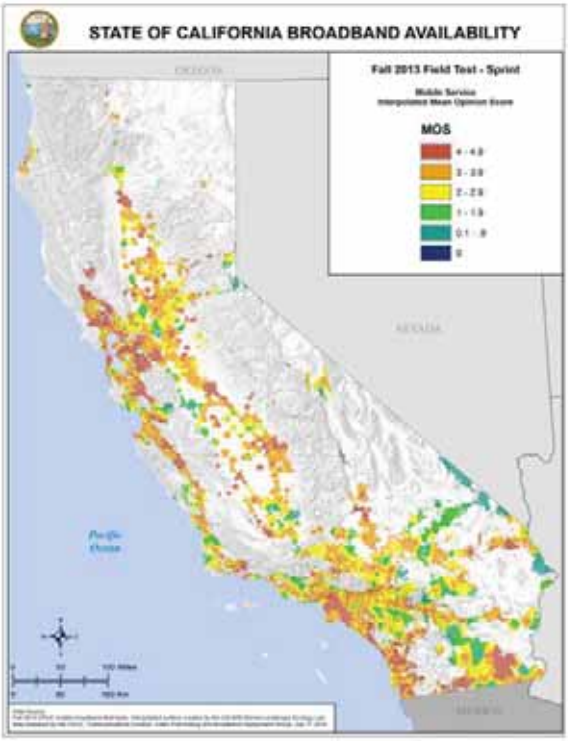
STATE OF CALIFORNIA BROADBAND AVAILABILITY
Fall 2013 Field Test - Sprint
Mobile Service
Interpolated Mean Opinion Score
MOS
4 - 4.9
3 - 3.9
2 - 2.9
1 - 1.9
0.1 - .9
0
Pacific Ocean

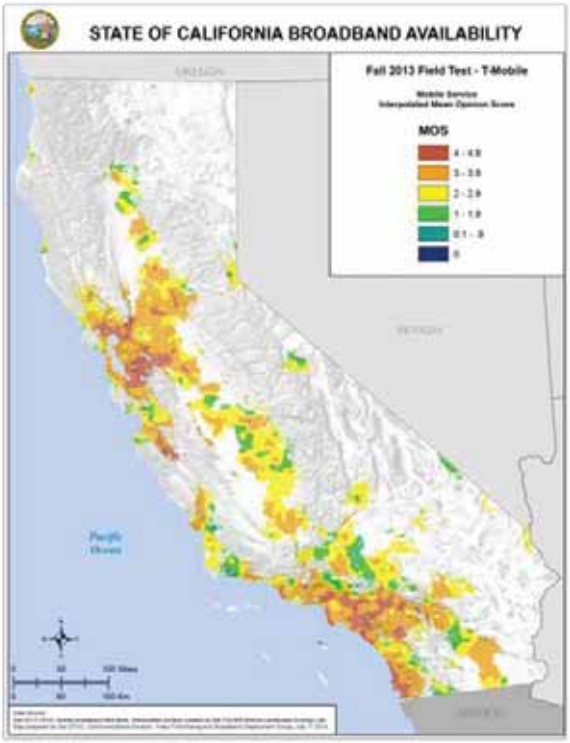
STATE OF CALIFORNIA BROADBAND AVAILABILITY
Fall 2013 Field Test - T-Mobile
Mobile Service
Interpolated Mean Opinion Score
MOS
4 - 4.9
3 - 3.9
2 - 2.9
1 - 1.9
0.1 - .9
0
Pacific Ocean

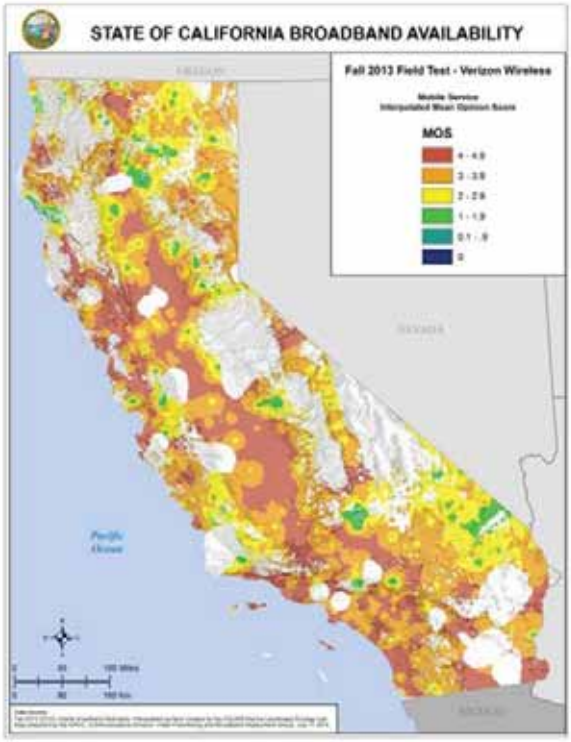
STATE OF CALIFORNIA BROADBAND AVAILABILITY
Fall 2013 Field Test - Verizon Wireless
Mobile Service
Interpolated Mean Opinion Score
MOS
4 - 4.9
3 - 3.9
2 - 2.9
1 - 1.9
0.1 - .9
0
Pacific Ocean

9. Measured Coverage Less Than Advertised Coverage

Each carrier advertises coverage with a specified “up to” downstream TCP performance. Using our kriging maps, CalSPEED can compute the percentage of the population covered at each performance level for each carrier. We can then compare the histograms of coverage, comparing how much of the California population is covered at each performance level - and then compare that to the implied advertised population coverage provided by the carriers.

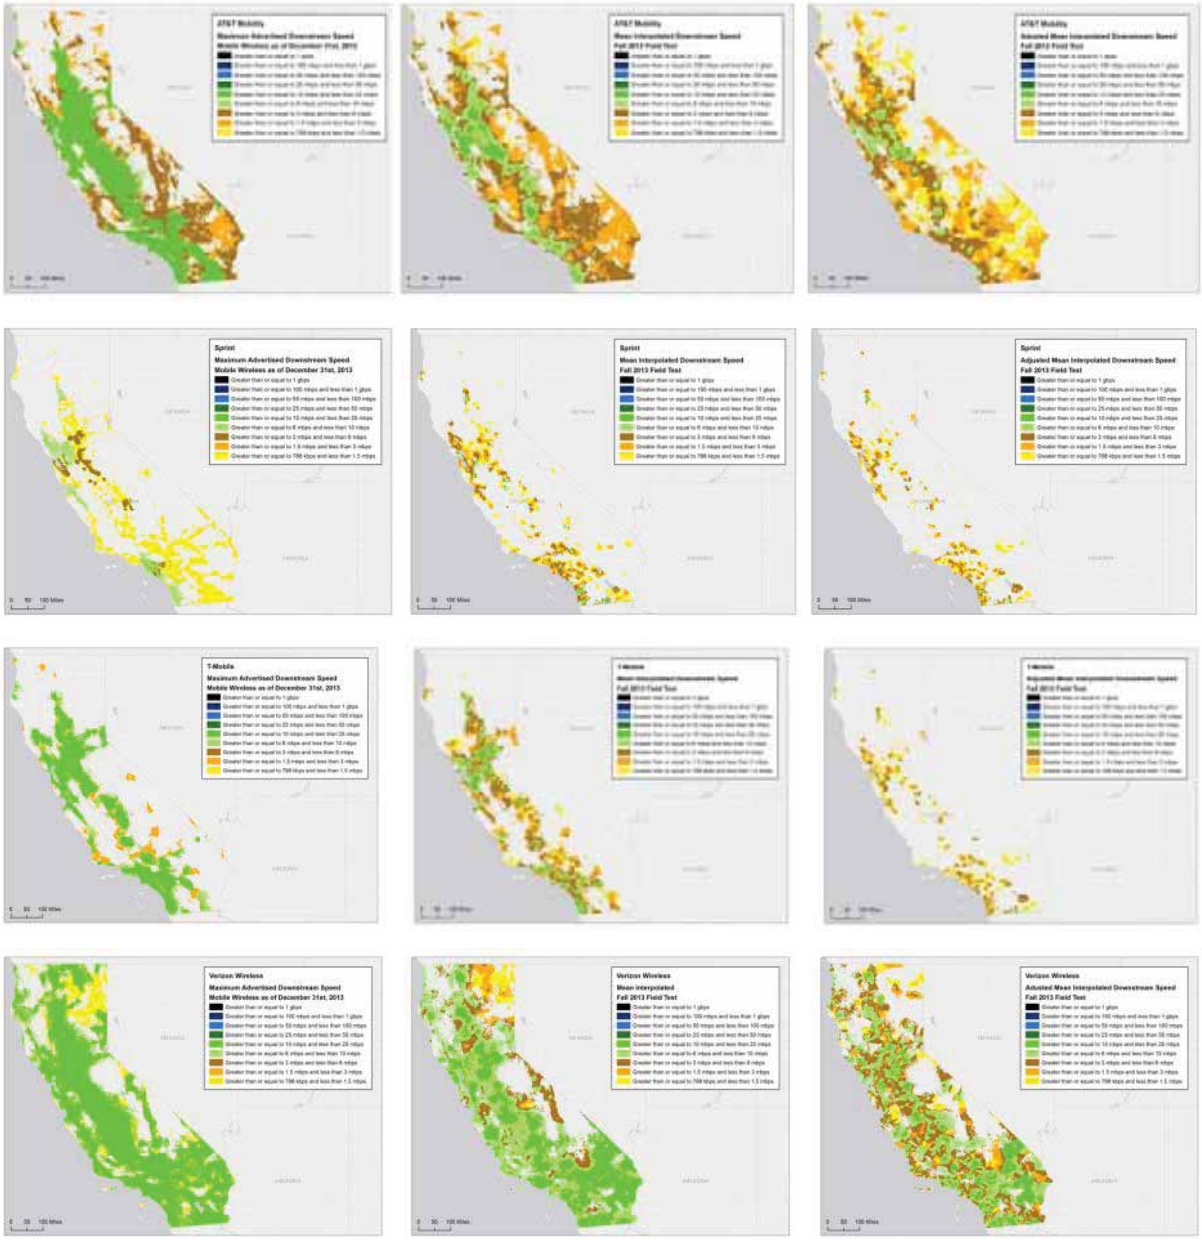

In the graphs above for each carrier, the left map documents the carrier's advertised "up to" performance level, the middle map documents CalSPEED's interpolation of mean downstream throughput and the map on the right documents an interpolation of adjusted measured mean throughputs at test locations of one standard deviation less the measured mean. As we have observed, mobile traffic has high variation and often a high standard deviation of service. Lowering the mean speeds down by one standard deviation is necessary to approximate what most consumers can expect from their mobile service most of the time.

10. What Is Mobile Broadband?

There are various standards for the definition of broadband service in the US - mostly derived from the quality of wired service, rather than wireless. The FCC currently uses 768 kbps down and 200 kbps up as "broadband" even though service is adequate only if service is reaches the benchmark of at least 4 Mbps up and 1 Mbps down. California uses the federal 768 Mbps/ 200 Mbps standard to determine whether a service is "broadband." However, service below 6 Mbps/ 1.5 Mbps is considered "underserved" Unserved (no service or slower than 768 Mbps/ 200 Mbps) and underserved areas are eligible for state and federal infrastructure grants, With the rapid advances in technology, particularly with the rapid deployment of 4G LTE, it is appropriate to consider what a modern mobile broadband standard might be.

All wireless carriers are rapidly deploying 4G LTE (and soon Advanced LTE) so that makes an appropriate baseline to examine. In California, the most evolved LTE network is Verizon.

Recognizing that downstream performance is rapidly increasing, let's begin looking at downstream and upstream TCP performance for Verizon over time.

A quick look at these cumulative histograms shows rapid change. The median mean downstream throughput has changed from 1 Mbps in the Spring of 2012 to about 8 Mbps in Fall 2013. The median up stream throughput has changed from under 1 Mbps in Spring 2012 to just over 4 Mbps in Fall 2013.

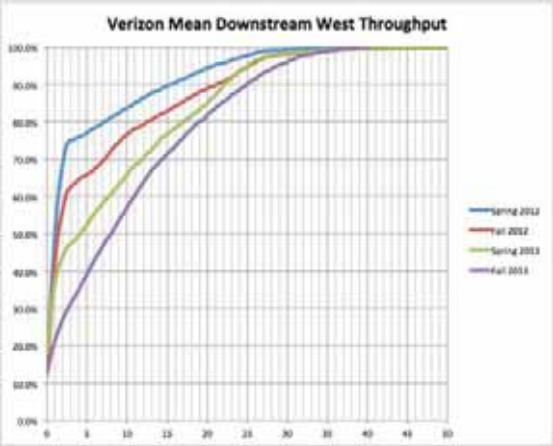

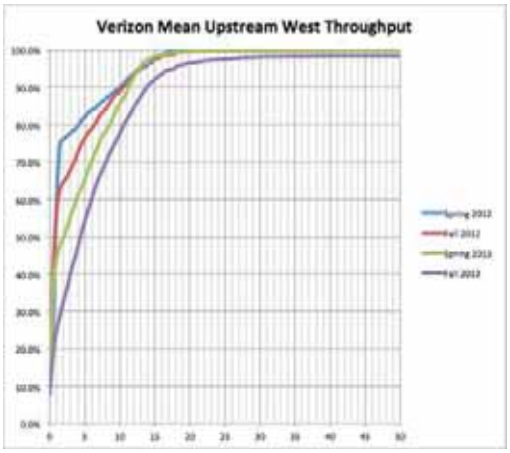

As all carriers will migrate to offer service profiles using 4G LTE similar to or exceeding Verizon, and we can expect that further LTE deployment (as well as the next generation of LTE - LTE Advanced) with further move these histograms to the right, it is appropriate to set a broadband standard that recognizes both the speed of the technology in performance and the speed of its deployment. A 10 Mbps downstream and 4 Mbps upstream standard is well within reach.

And the network will advance with technology. Any standard should be reevaluated periodically based on developing application and usage patterns.

10.1 What are Appropriate Mobile Speed Benchmarks

The FCC is exploring whether its current broadband benchmarks of 4 Mbps down and 1 Mbps up, used to determine CAF grant-eligible areas, should be increased, and what the impact of various benchmarks would be. The chart below shows the percent of California population and land area that would be considered served (i.e., at least as fast as the benchmark,) and what percentage would be considered unserved (i.e., below the benchmark),using various benchmarks.

Impact of Various Speed Benchmarks on Served Status in California
As of December 31, 2013

STATE OF CALIFORNIA		Population and Land Area "Served"			
		Served		Unserved	
Speed Benchmarks		**Total**	**%**	**Total**	**%**
Mean; 4 mbps down, 1 mbps up	Pop.	37,966,622	**99.0%**	373,453	**1.0%**
	Sq. Mi.	99,797	63.62%	57,076	36.38%
Mean - 1 STD; 4 mbps down, 1 mbps up	Pop.	36,429,622	**95.0%**	1,769,274	**4.6%**
	Sq. Mi.	80,032	51.02%	76,841	48.98%
Mean; 6 mbps down, 1.5 mbps up	Pop.	37,397,891	**97.5%**	832,202	**2.2%**
	Sq. Mi.	88,248	56.25%	68,625	43.75%
Mean - 1 STD; 6 mbps down, 1.5 mbps up	Pop.	24,617,299	**64.2%**	13,722,776	**35.8%**
	Sq. Mi.	53,806	34.30%	103,067	65.70%
Mean 10 mbps down, 1 mbps up	Pop.	19,690,366	**51.4%**	18,527,960	**48.3%**
	Sq. Mi.	44,753	28.53%	112,120	71.47%
Mean - 1STD; 10 mbps down, 1 mbps up	Pop.	2,444,043	**6.4%**	35,754,931	**93.3%**
	Sq. Mi.	13,115	8.36%	143,758	91.64%
Mean; 10 mbps down, 4 mbps up	Pop.	19,654,246	**51.3%**	18,568,291	**48.4%**
	Sq. Mi.	44,049	28.08%	112,824	71.92%
Mean - 1 STD; 10 mbps down, 4 mbps up	Pop.	2,436,091	**6.4%**	35,763,248	**93.3%**
	Sq. Mi.	12,845	8.19%	144,028	91.81%

Sources:
Broadband Drive-Test Results, Prepared by the CPUC in collaboration with CSU Chico and CSU Monterey Bay
Household data from the California Department of Finance, January 1, 2014 estimate.
*-1STD indicates we have lowered the each provider's estimated mean throughput by one standard deviation to arrive at service speeds likely to be experienced by most customers. The served and unserved calculations estimate the number of customers with service from any provider most of the time that meets the stated benchmark.

Using a benchmark of 4/1, 99% of the state is served (leaving aside issues of price and usage limits). When that benchmark is moved up to 6/1.5, (the levels used by the CPUC for prior infrastructure grants), the percentage served is still 97.5%. Either benchmark will result in considering virtually the whole state served, with little area left to subsidize. Yet we know there is an unmet need for fast mobile broadband. At the FCC's proposed 10/1 benchmark, only 51.4% of the

population would qualify as served. Interestingly, when the benchmark is raised to 10/4, the percent served only falls to 51.3.

The speed of deployment in mobile broadband service can support a new benchmark standard of 10 Mbps down and 4 Mbps up based on the deployed capabilities of modern LTE networks. When subsidizing mobile deployment, it would certainly be prudent to require deployment at speeds of at least 10 Mbps down/4 Mbps up, properly configured to be able to provide VoLTE and other real-time streaming services.

11. Conclusions

This paper has examined the key findings of the first four measurement rounds of CalSPEED covering 18 months of measurement between Spring 2012 and Fall 2013. There have been rapid changes during that time and the data provide a solid foundation for seven key findings.

- Mobile broadband, at its best, is getting much better VERY quickly.
- Mobile broadband shows wide variation in performance across California.
- Not all carriers are equal - there is a substantial and growing difference in the coverage and quality of mobile broadband service between the best performing and worst performing carrier.
- Mobile broadband service quality is not just local - the user experience is governed not just by local radio access but also by Internet backbone interconnect strategies of the carriers.
- There is a real and growing mobile digital divide between urban and rural and tribal.
- The bulk of the mobile network is not yet VoIP ready.
- Service availability as measured by CalSPEED is less than carriers' "advertised service".

Appendix A: How CalSPEED Measures

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 (4 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 10 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 uses two identical measurement servers on the opposite ends of the Internet. One hosted in the Amazon AWS near San Jose, CA and for many users has performance like a CDN server. The second measurement server is in the Amazon AWS in Northern Virginia.

CalSPEED uses two device measurements - a current smartphone and current USB datastick for laptops. Both are upgraded for each measurement round to match the latest wireless technology deployed by each carrier. For the Fall 2013 measurement survey, these were the devices used.

Carrier	Smartphone	Datastick
AT&T	Samsung Galaxy S4 (LTE)	AT&T Beam
Sprint	Samsung Galaxy S4 (LTE)	Connect Tri-Mode
T-Mobile	Samsung Galaxy S4 (LTE)	Samsung LTE Mobile HotSpot PRO
Verizon	Samsung Galaxy S4 (LTE)	UML295

Appendix B: Maps

CalSPEED's kriging methodology creates maps plotting a number of mobile broadband metrics. The body of the paper included the maps for mean downstream TCP throughput, this appendix includes the maps for mean latency.

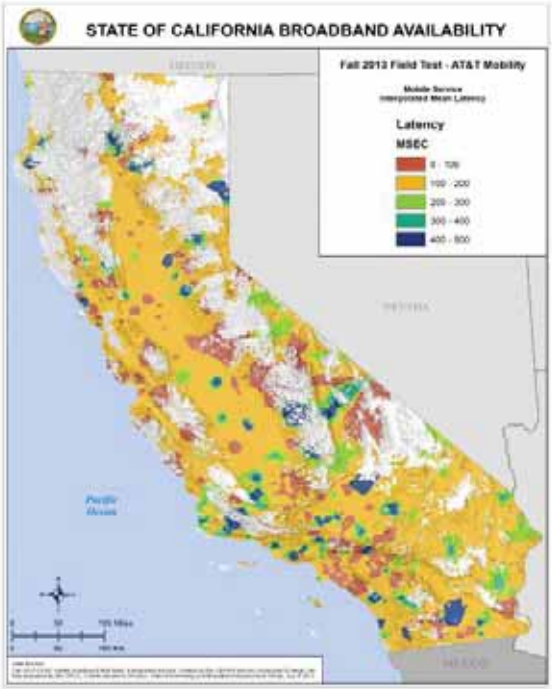

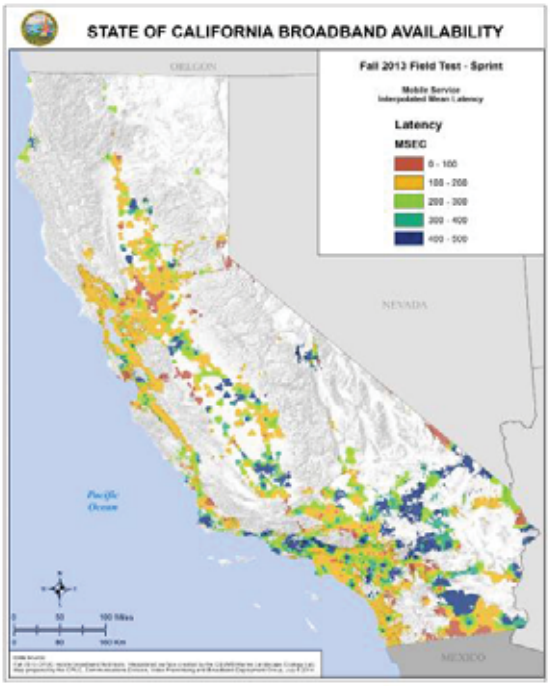

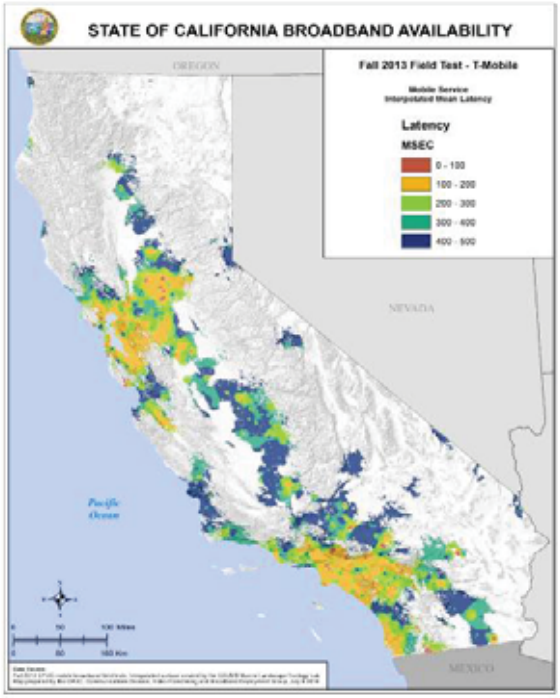

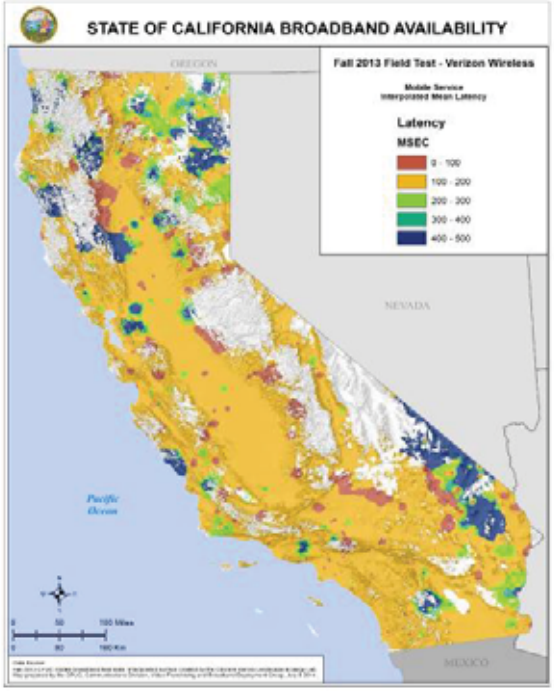

APPENDIX B

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*.

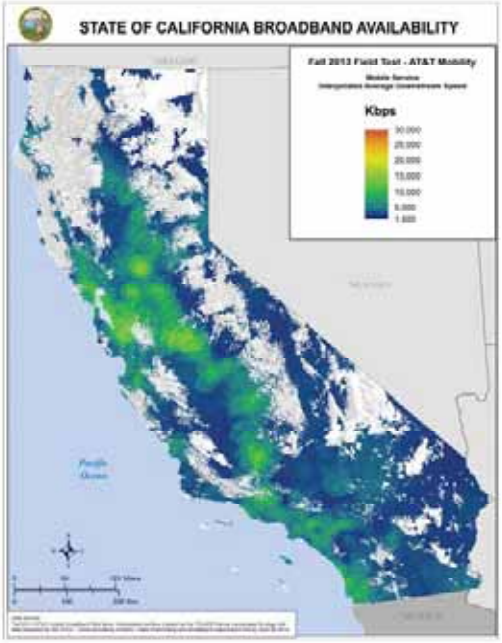

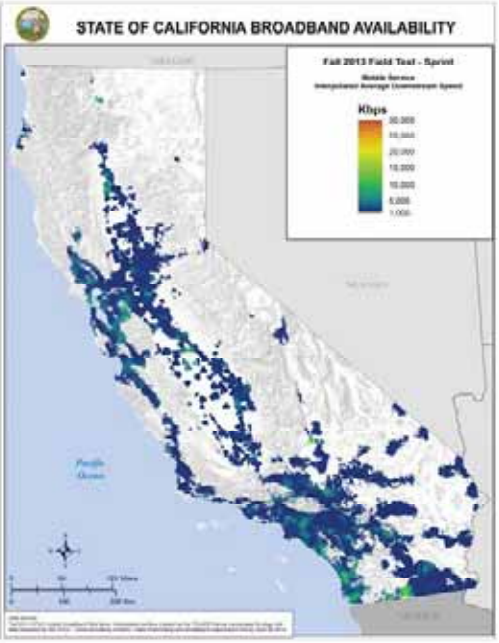

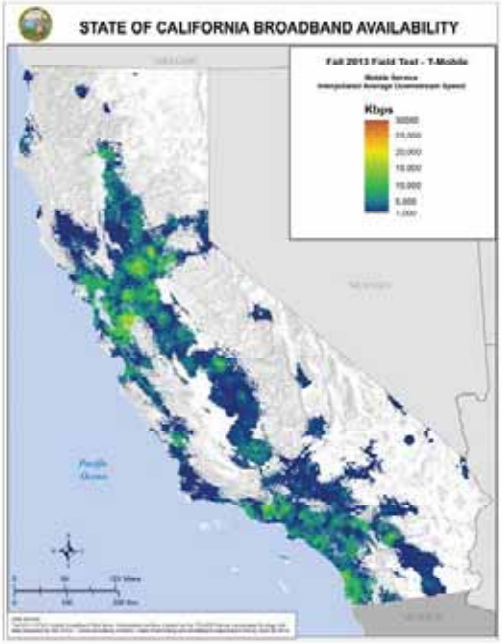

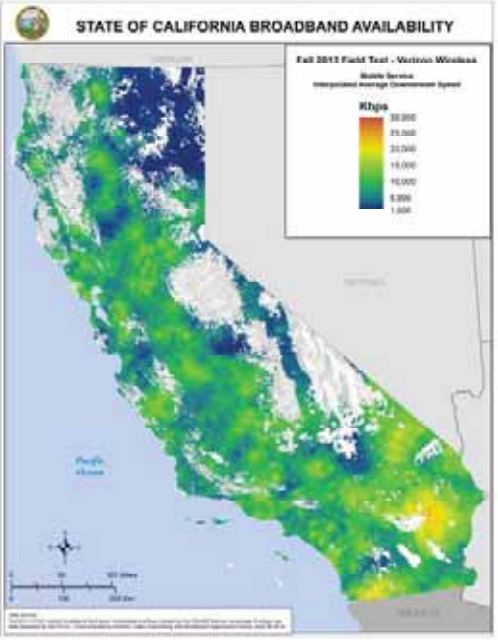

2.2 CalSPEED Methodology

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?

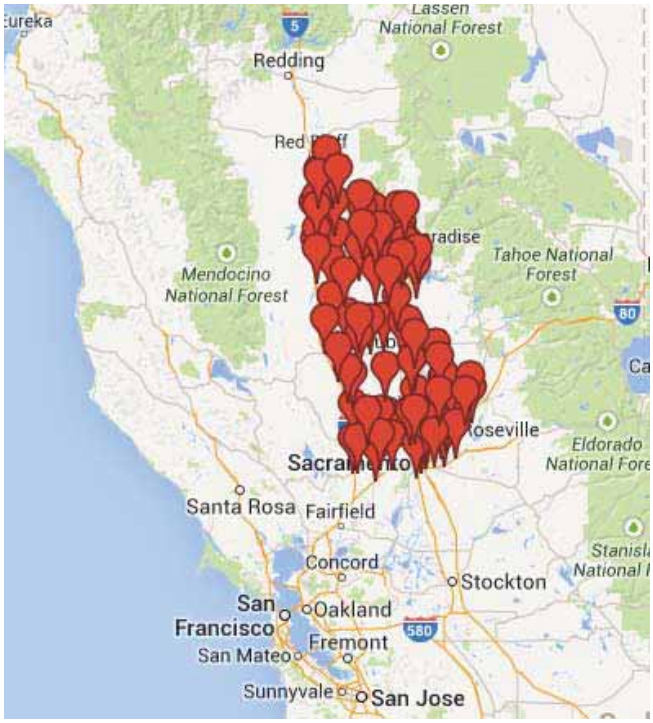

CalSPEED is one of several network performance measurement tools available free of charge to the public. We measured at the same locations, at the same time with both the latest FCC[3] and Ookla[4] test on the same smartphones, in the same locations and at about the same time[5].

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.

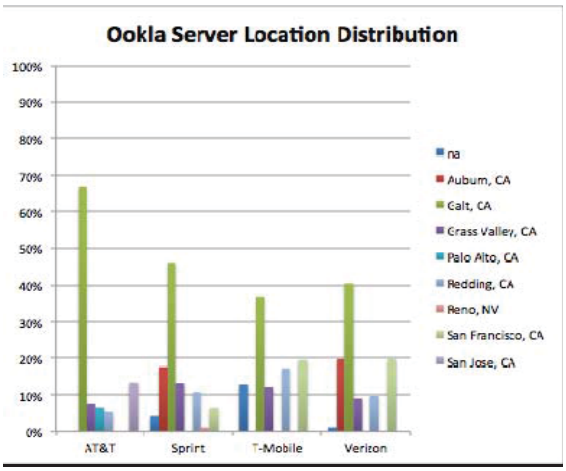

The composition of servers chosen differs between each carrier which makes comparing carriers, in a given location or across the complete set of locations, challenging since we are comparing different measurements. For example, 67% of the tests run for AT&T select the Galt, CA server while only 37% of 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.

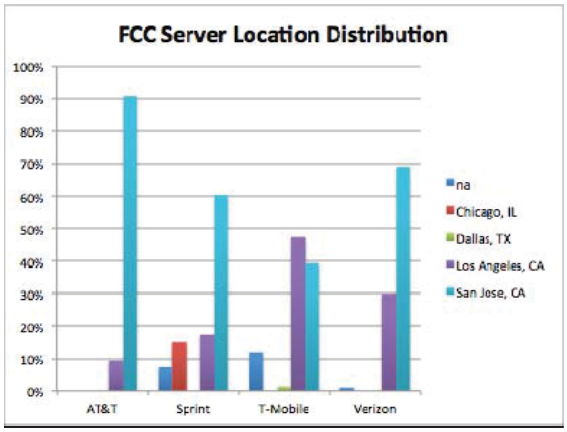

As with the Ookla tests, the distribution of server locations differs between depending on which carrier is tested making direct comparison of test results between carriers either at any given test location or among the complete set of test locations challenging. Each test independently filters to a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

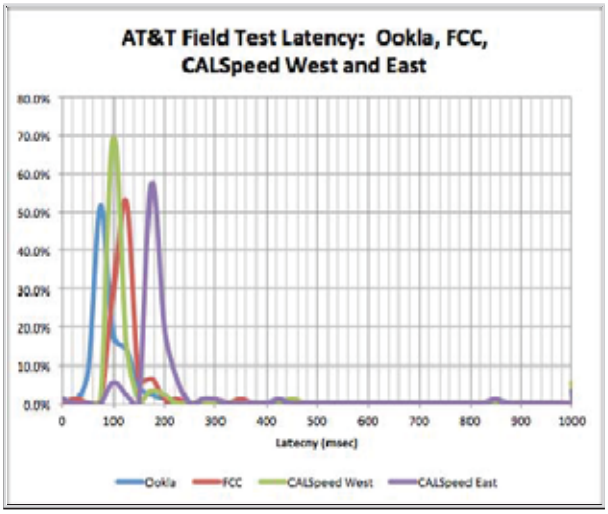

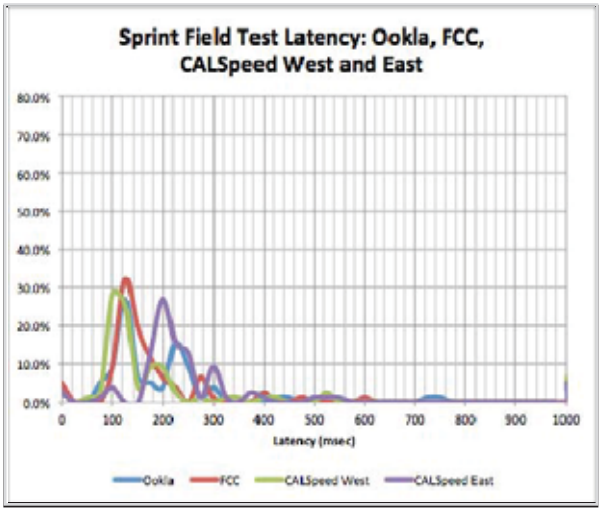

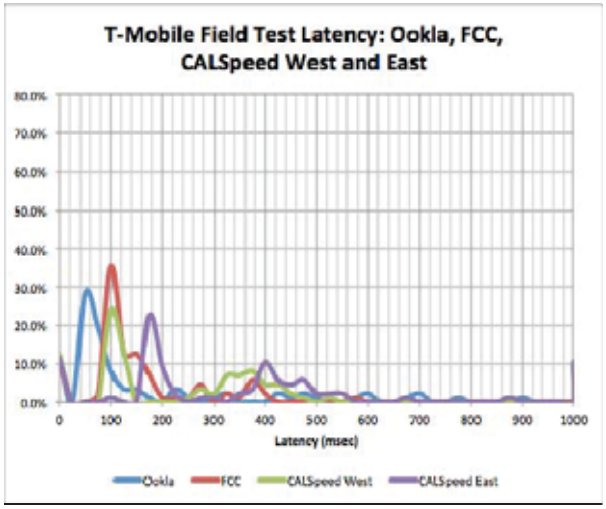

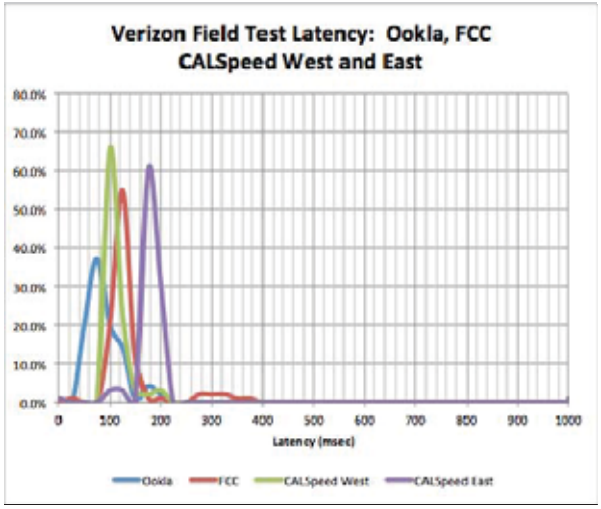

As we might expect given the geographic distribution of server sites, Ookla generally delivers the lowest absolute latency, on the order of 20% less than CalSPEED West and about half of CalSPEED East. This difference reflects the substantial difference in the length of the Internet to be traversed between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

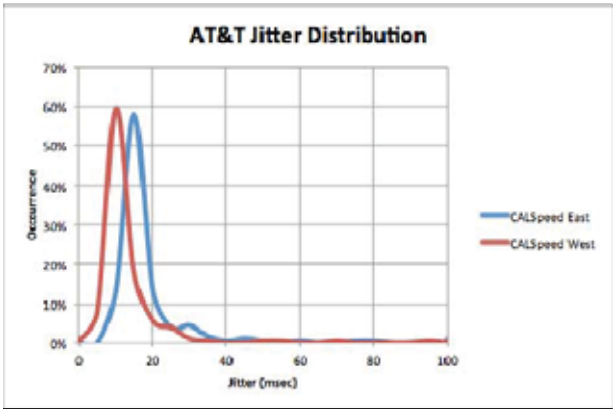

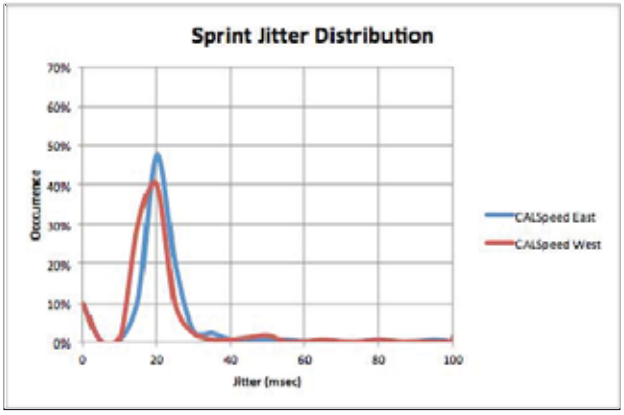

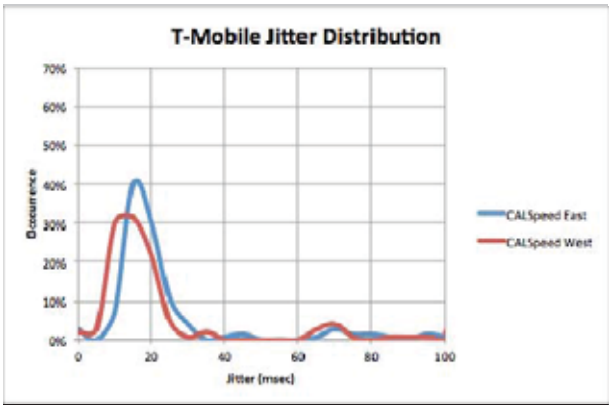

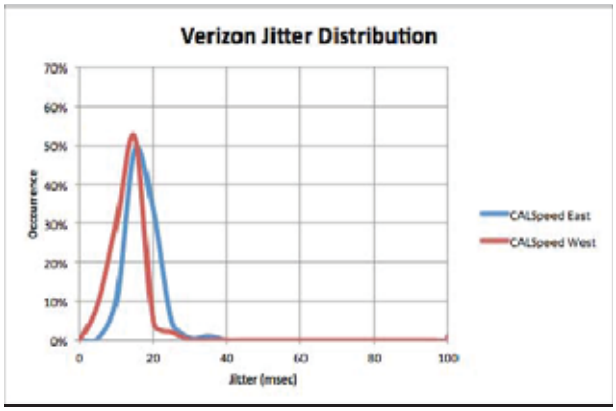

3.7 Packet Loss

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

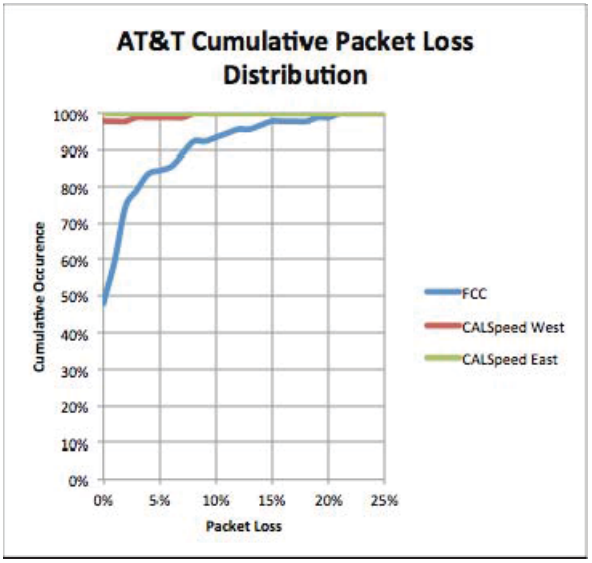

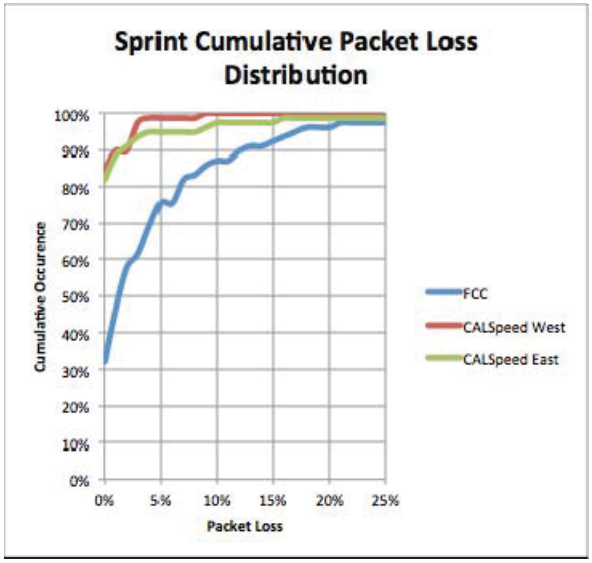

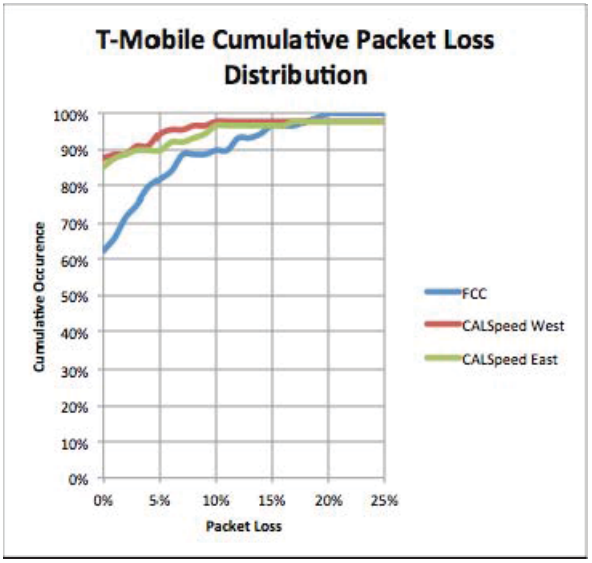

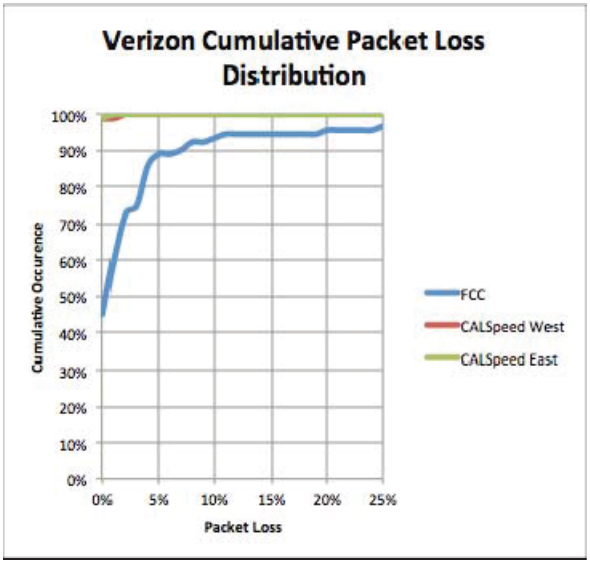

The FCC measurement shows materially higher packet loss numbers than either CalSPEED West or East. We do not have an explanation for these much higher packet loss measurements, but such high numbers would make VoIP or other real-time streaming services very difficult.

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.

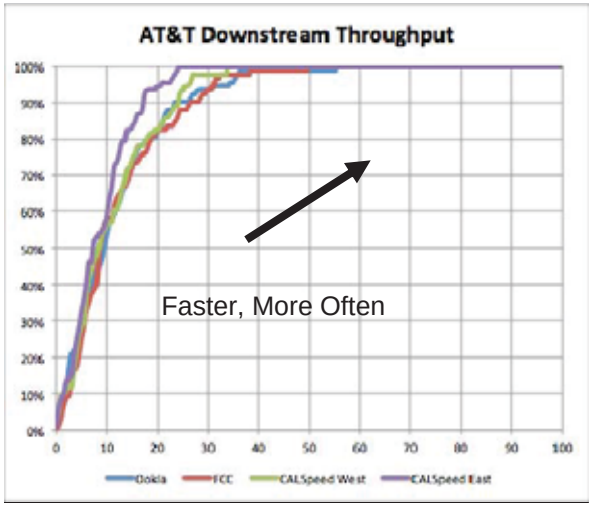

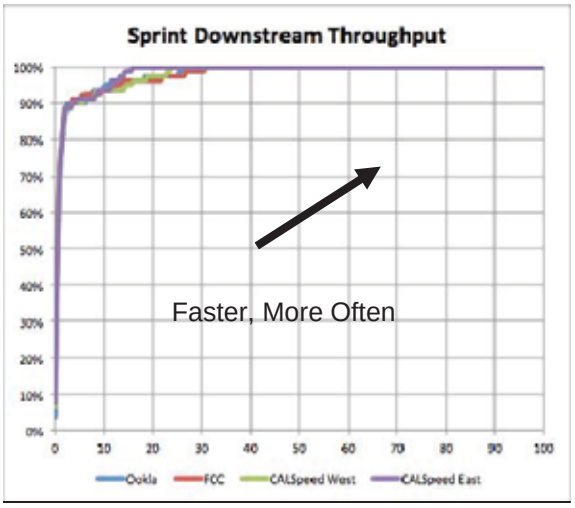

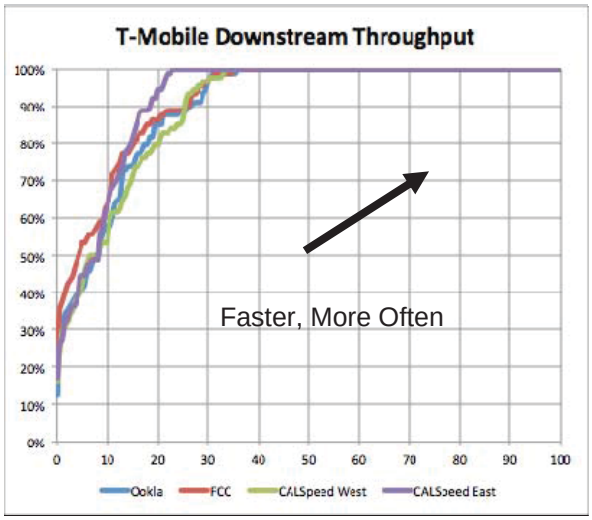

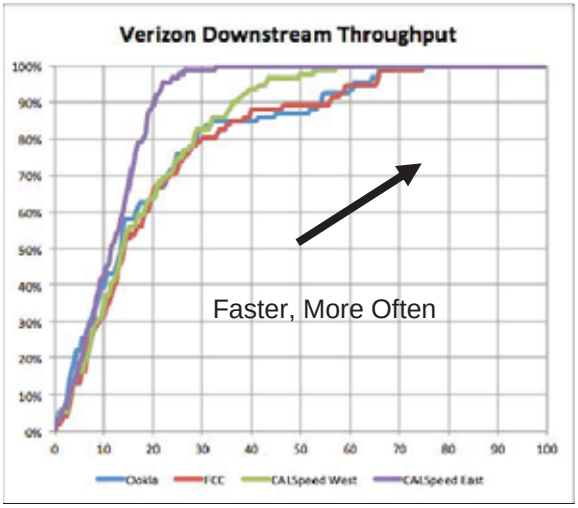

Selecting the server with the lowest latency biases the result to a higher average throughput measurement, particularly at higher performance levels. The differences between CalSPEED and FCC and Ookla are modest at low speed networks (Sprint and T-Mobile) with the advent of LTE we can see that the selection 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.

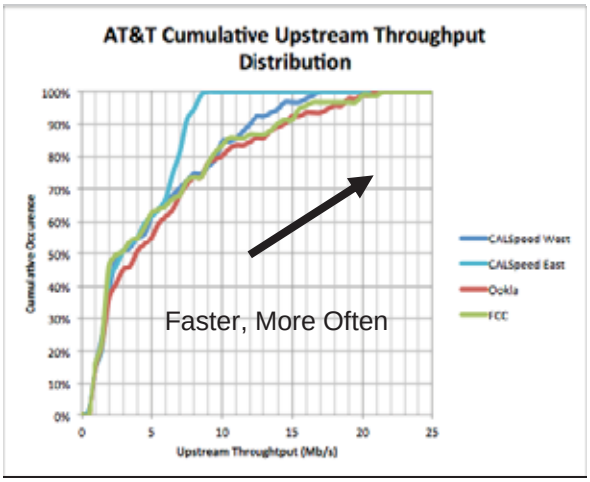

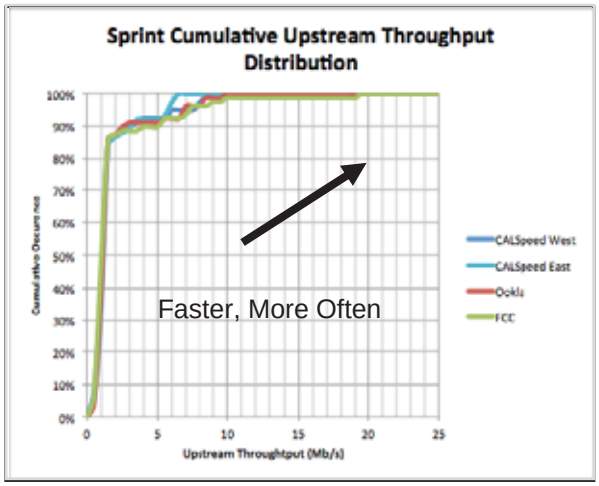

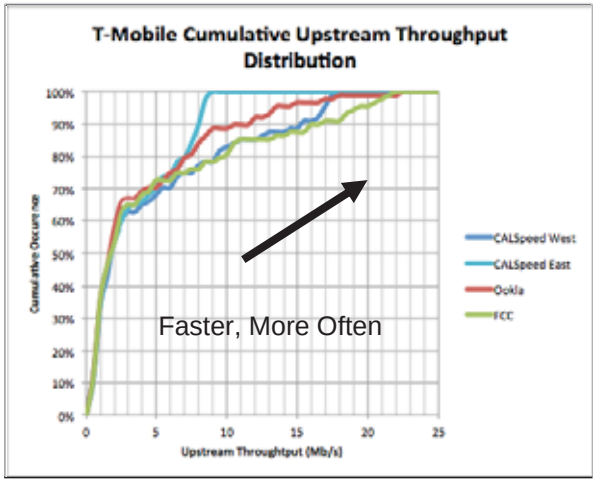

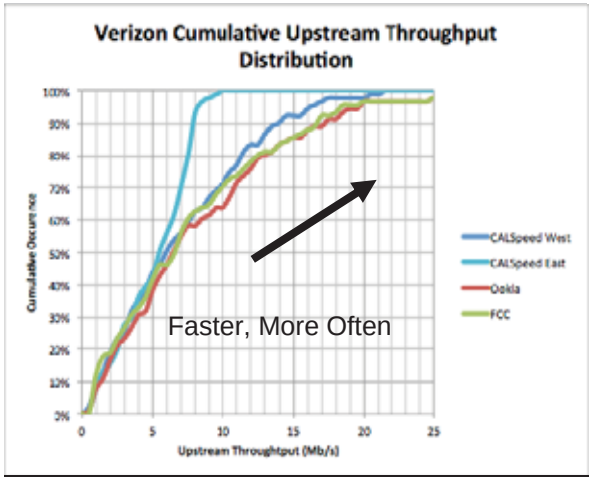

3.10 MOS

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.

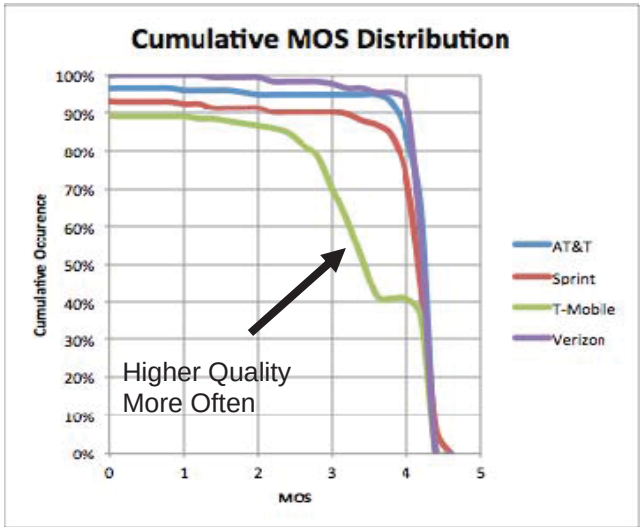

3.11 Wireless Network Technology

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.

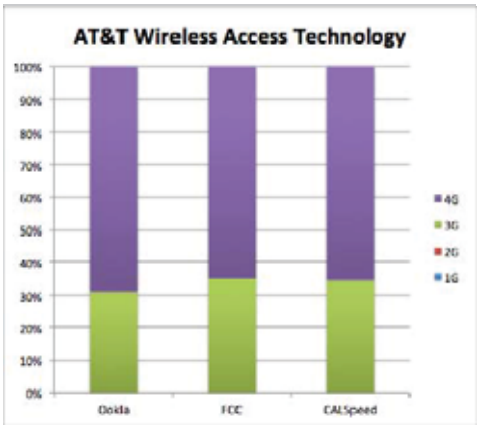

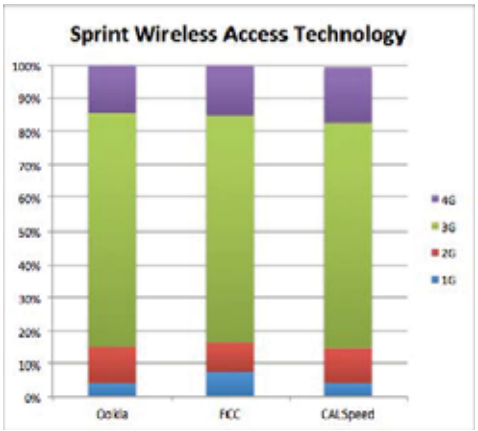

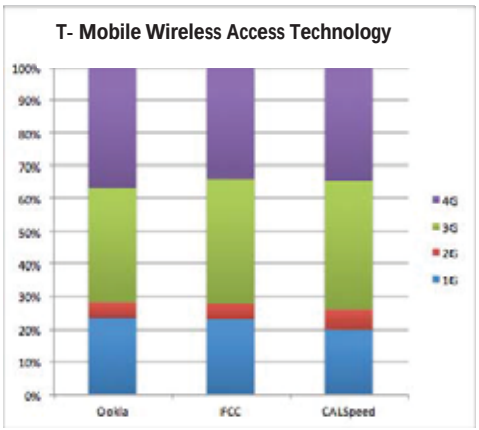

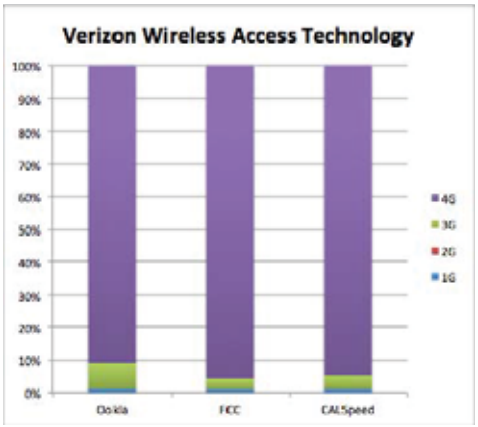

3.12 Interpolation and Mapping

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.

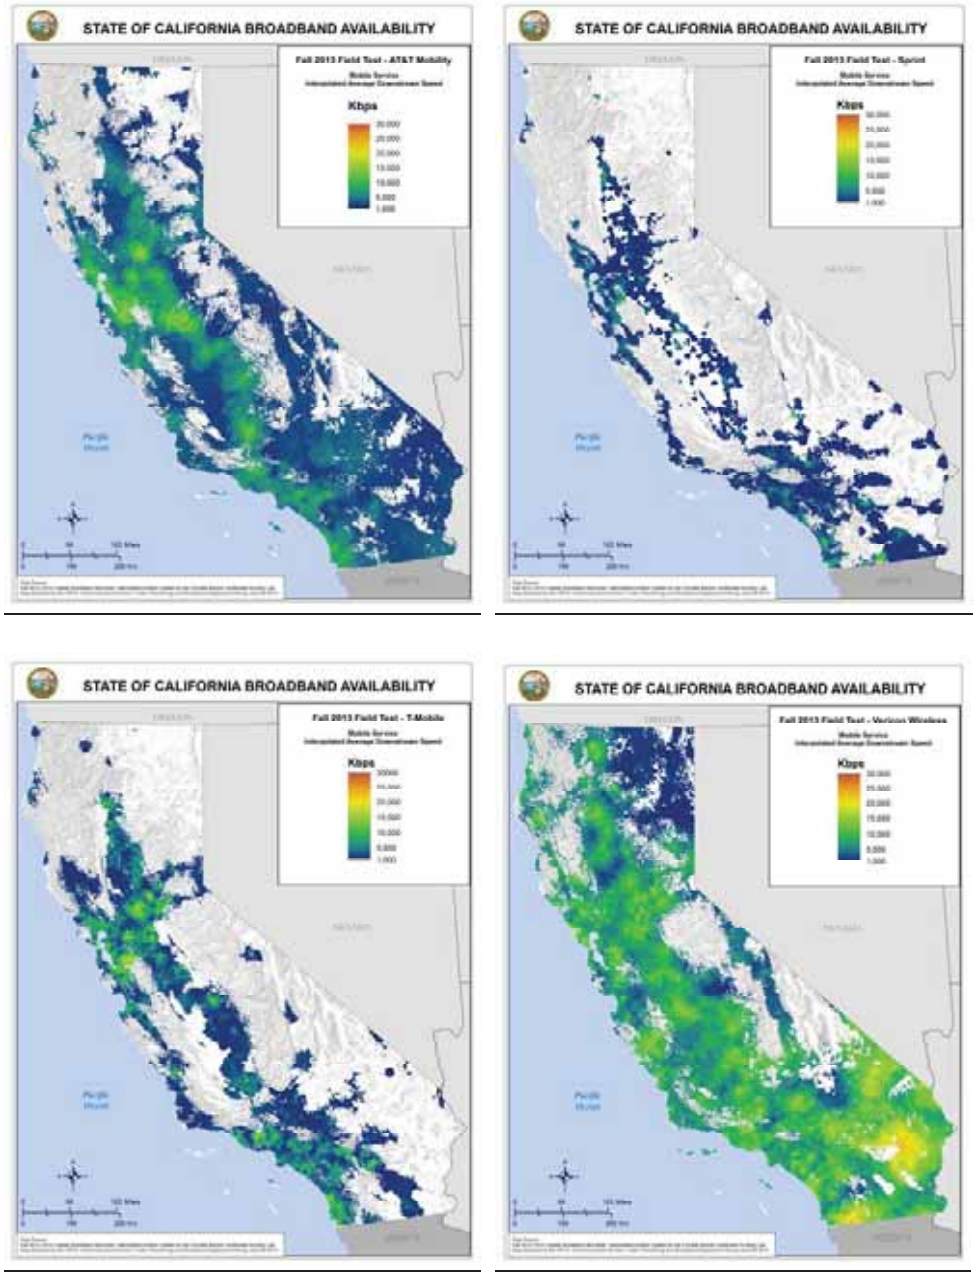

4.0 Conclusions

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.