

CalSPEED: California Mobile Broadband - An Assessment - Spring 2015

Quality vs Throughput

Ken Biba¹

Managing Director and CTO

Novarum, Inc.

This is an analysis of the Round 7 Spring 2015 dataset for CalSPEED. It shows substantial changes happening in the wireless broadband networks in California for all carriers.

- Major carriers are showing decreasing mean throughput. This slowdown in performance has been apparent for the last two measurement rounds and appears to be a trend. Detailed analysis suggests this is coming from throttling high performance results.
- Broadband coverage is decreasing for all carriers. For the 25 mbps downstream/3 Mbps upstream threshold, Verizon decreased from 16% coverage in fall of 2014 to 4% coverage in the spring of 2015.
- Many measures of quality (latency, jitter, TCP failures, mean packet loss) are improving for most carriers.
- Substantial improvement in estimated OTT MOS is apparent for digital voice service.
- Penetration of LTE in both urban and rural geographic categories appears to have peaked with a floor on 1/2G legacy replacement and on a cap on LTE deployment.
- The performance difference between East and West servers is decreasing suggesting material improvements in backhaul
- In recent rounds, USB modems were ~30% slower than the phones, in Round 7 the tablets (that replaced the USB modems) were about ~20% faster than the phones - reflecting dramatic performance differences between the USB modems and the tablets.
- Program change from using USB modems to tablets complicates the analysis. This paper is based on using the phone measurements alone. Tablets strongly outperform the legacy USB modems.

These trends, considered in aggregate, suggest wireless networks that are being managed more carefully to deliver higher quality of service, particularly for VoIP, at the price of decreased throughput. The stall in LTE deployment in both urban and rural geographic categories can only be speculated on.

¹ Ken Biba is a consultant to the CPUC's CalSPEED project. He was tasked with preparing this analysis of the Spring 2015 CalSPEED field test results. The content of, and conclusions reached in, this Report are Mr. Biba's own, and do not necessarily reflect those of the State of California, the CPUC, its Commissioners or its employees.

1. Calibrating the Mobile Internet Experience

Each of us relies on the Internet to research school papers, find a job, find and buy new products, 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 much 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, funded originally via a grant from the National Telecommunications and Information Administration. CalSPEED is now funded through the California Advanced Services Fund. CalSPEED uses a methodology pioneered by Novarum. The software measurement system is created and maintained by a team at California State University Monterey Bay, led by Professors Sathya Narayanan and YoungJoon Byun. CalSPEED mapping and measurement field operations are managed by the Geographic Information Center at California State University at Chico. Statisticians at CSU Monterey Bay assist the team with detailed geographic and statistical analysis of the dataset.

CalSPEED has now been in use in California for three and a half years with seven rounds of measurement over the entire state collecting over 15,000,000 measurements across California of the four major mobile broadband carriers: AT&T Mobility, Sprint, T-Mobile and Verizon Wireless. This paper does a deep analysis of the first seven rounds of measurement. Previous papers have analyzed the first six rounds of measurement^{2 3}. The methodology has been rigorously analyzed with respect to other available mobile measurement tools⁴.

This paper examines the incremental changes from the previous report extending thru the spring of 2015.

Let's examine what CalSPEED tells us about the state of mobile broadband in California after over three years of measurements in the same locations.

² Ken Biba, "Assessment of California Mobile Broadband Spring 2014", Novarum, September 2014.

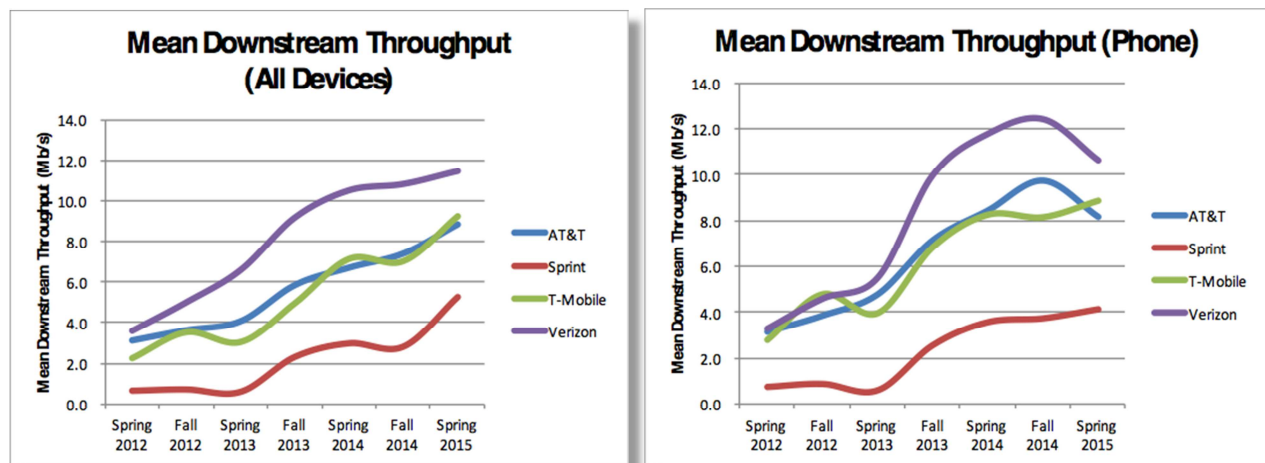
³ Ken Biba, "Assessment of California Mobile Broadband Fall 2014", Novarum, June 2015.

⁴ Ken Biba, "Comparison of CalSPEED, FCC and Ookla", Novarum, Inc., September 2014.

2. Change in Devices

The CalSPEED program for spring 2015 replaced the USB modems that had been used along with Android smartphones with Android tablets. We replaced the USB modems because their performance had been decreased relative to the Android smartphones. We attributed this relative performance decrease to declining market share of USB modems and hence declining commercial support by the carriers to update their technology.

While most major overall trends are consistent, the tablets are clearly outperforming the USB modems and that changes the results for the combined device averages. It is thus imperative to do the comparative trend baseline analysis using only the smartphone data. Further, it is imperative for the integrity of the study going forward that each round use at least one device equipped with the latest cellular radio technology. The next version of LTE to be deployed in the U.S. is LTE Advanced, and we can anticipate the next measurement round to show material increases in throughput based on the improvement in the underlying wireless technology.



The graph on the left shows the mean (average) downstream throughput for both USB modems/tablet and smartphones devices. Prior to the spring of 2015 this was the average of a USB modem (and its attached Windows laptop) and an Android smartphone. From the spring of 2015, it is the average of an Android smartphone and an Android tablet. There is a clear trend towards ever increasing throughput.


However, just looking at the Android smartphone data (the graph on the right) we see a very different story. While there is a clear trend towards increasing throughput across all carriers through spring 2014, there are clear trends that throughput performance peaked for all carriers subsequently and for the highest performing carriers – AT&T and Verizon – downstream throughput has actually decreased subsequently – and rather dramatically. Further, we can see that the lower throughput of the USB modems, brought the combined average down, particularly during the 2014 measurements.

Other analysis (see Section 8) will look more closely at the difference in performance among difference devices – smartphones, USB modems and tablets.

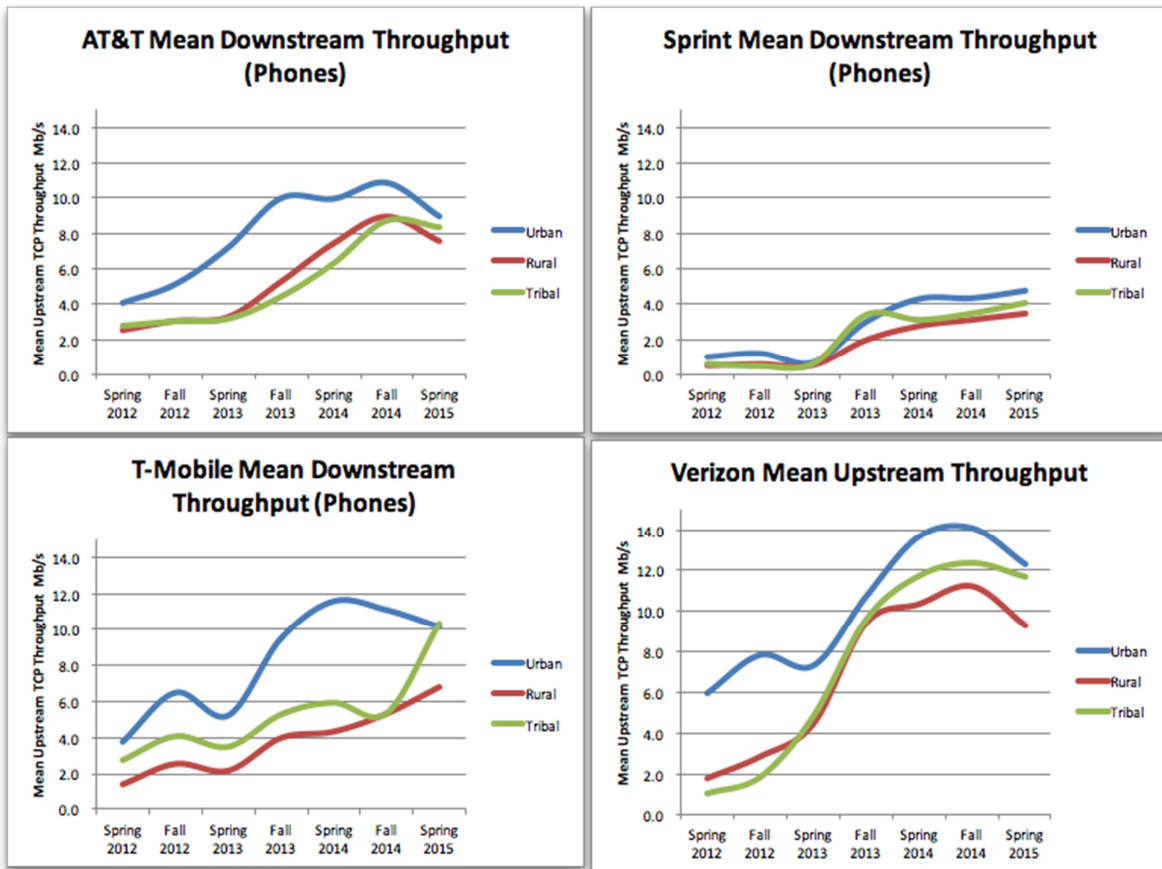
For the rest of this analysis we have chosen to look only at the smartphone data. We think it is more representative of the underlying trends.

3. Throughput Decreasing

AT&T and Verizon are showing clear decreases of ~15% in mean throughput for phones.





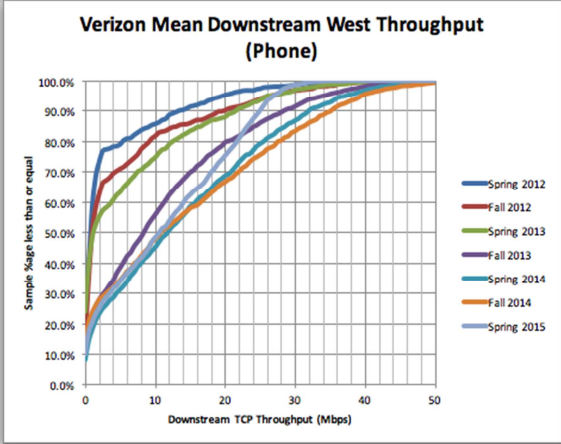
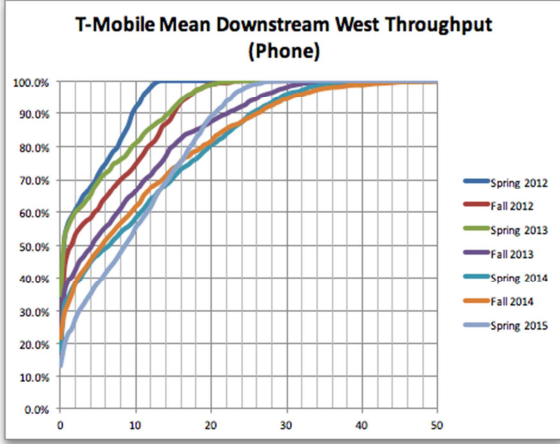
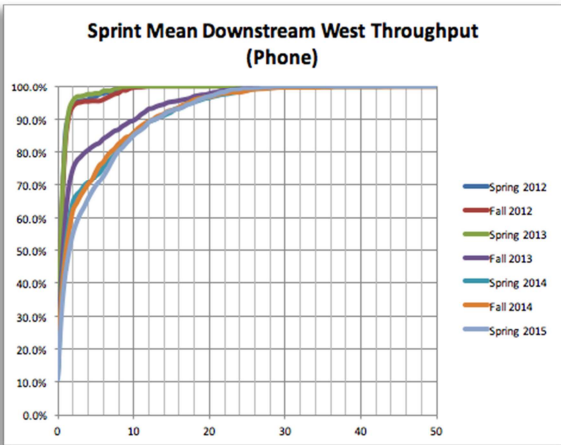
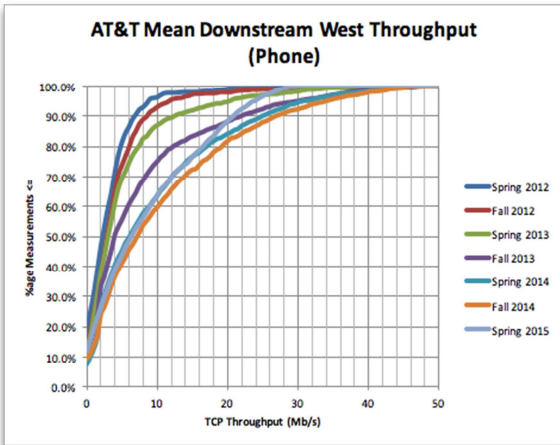
This trend of decreasing throughput appears not to be isolated to just one geographic category.. As the chart below demonstrates, for AT&T and Verizon, downstream throughput is decreasing in all geographic categories, stalled for Sprint in all geographic categories and dramatically decreasing in urban areas for T-Mobile.

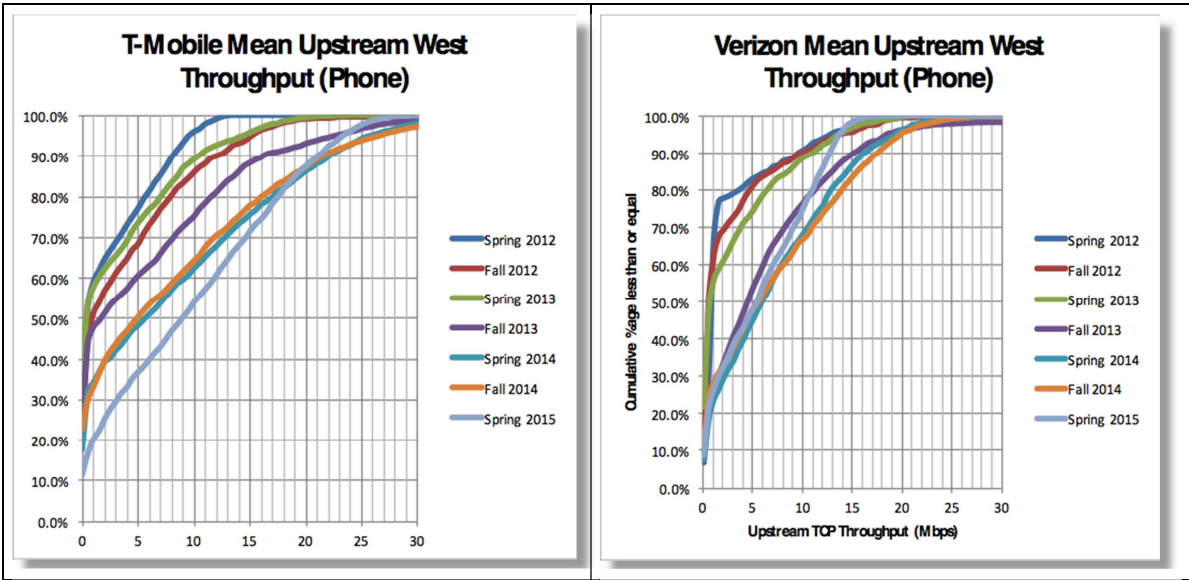


Below are the cumulative histograms of downstream throughput for phones. We can see a clear change in the distribution for spring 2015. In prior measurement rounds we can see the distribution, for every carrier, move further to the right with time – indicating more measurement locations with higher performance. For spring 2015 we can see a clear difference – across all carriers – of the distribution curve moving to the left (both downstream and upstream), documenting a change in the performance distribution towards lower speeds.

This suggests intentional filtering of high performance throughput results.

This decrease is reflected in the full histogram of samples for both downstream and upstream throughput - throttling high throughput both downstream and up.






4. Broadband Coverage Decreasing

With decreasing mean throughput, statewide coverage at any given standard of throughput performance decreases.

Let's first look at the distribution of downstream performance, in particular the percentage of measurements that have downstream throughput greater than or equal to 25 Mb/s and upstream throughput greater than 3 Mb/s.. There is a dramatic decrease in the number of locations with high throughput – for all carriers.


If we then look at the change in coverage - that is the change in the number of locations meeting the 25/3 threshold between fall 2014 and spring 2015 - we see dramatic decrease for all carriers and all demographics.



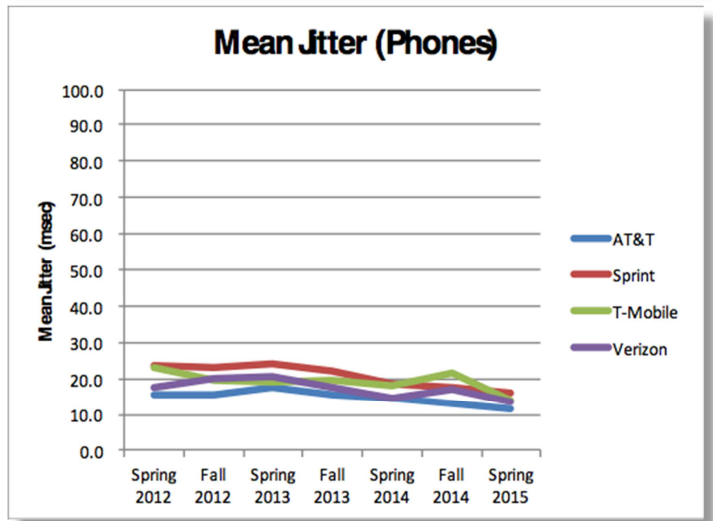
5. Many Measures of Service Quality Improve

Even while both mean throughput and the throughput distribution is decreasing (a suggestion of throttling of high throughput users), many measures of quality improved.

Mean latency improved for all carriers.

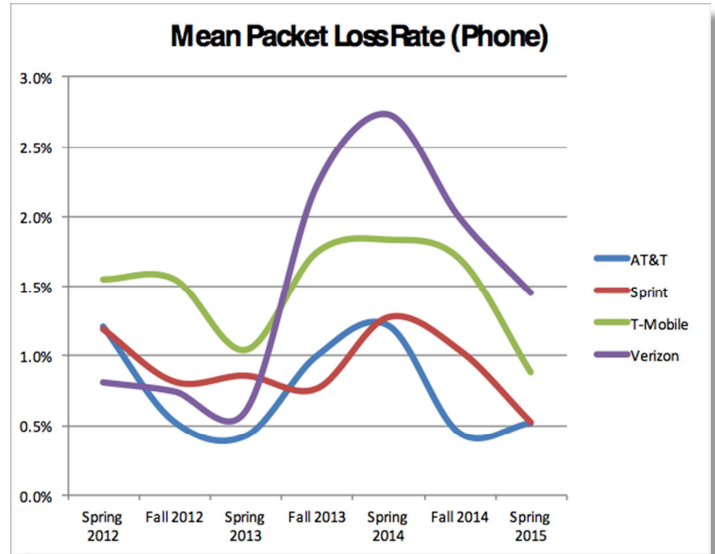


Mean jitter⁵ improved for all carriers.




⁵ Jitter is the measure of variation in latency. It is important component of performance for real time streaming of audio and video.

Mean packet loss⁶ rates improved for all carriers.



TCP quality⁷, as measured by failed TCP connection attempts⁸, slightly improved but there remain profound differences between urban and rural areas with rural users having an almost 400% higher TCP connection failure rate over urban users.





TCP quality, as measured by TCP throughput variation⁹ did not show much improvement in the median variation. This is a measure of the standard deviation of throughput variation during TCP sessions.

⁶ Mean packet loss is the average percentage of packets that are lost during transmission. Small increases in packet loss are particularly degrading for streaming media services.

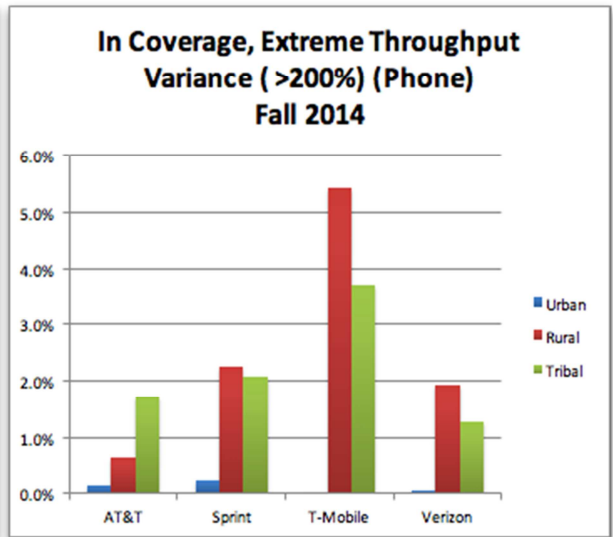
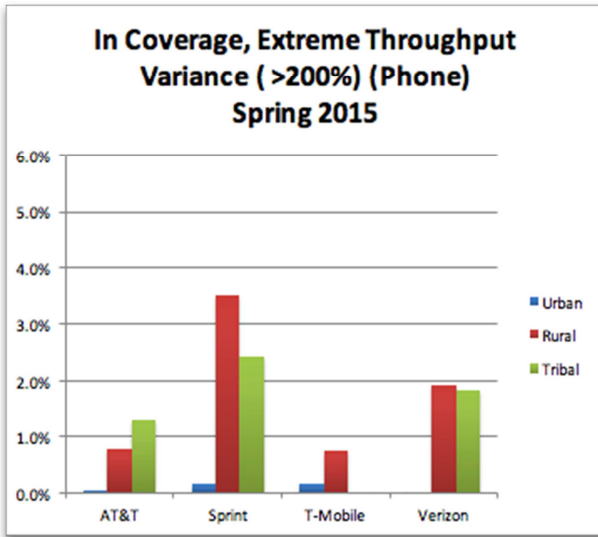
⁷ The fundamental reliable connection service for the Internet is TCP - Transmission Control Protocol. It provides reliable delivery of an ordered stream of bytes and is the foundation service for web browsing, most streaming media services, email, IM and most other user Internet services. CalSPEED measures TCP quality in several ways: the failure rate of making a connection, and the consistency of the throughput during the connection - throughput variation.

⁸ TCP connection failure is a measure of how often TCP attempts to make a connection between source and destination and succeeds or fails to make the connection. It is the Internet equivalent to how often, in making a phone call, the call fails to connect to the destination.

⁹ TCP throughput variation measures how consistent throughput is during a connection - does it vary and by how much. A more consistent connection will deliver higher quality streaming media service.




However, there is modest improvement on the very high end of variation. For connections with extreme throughput variations - where the standard deviation of throughput during a session is > 200% of the mean throughput - there are selective decreases in variability. A decrease in variability translates to a more consistent service particularly for streaming media. The biggest improvement in this metric came from T-Mobile. Modest improvements for AT&T rural and Verizon Tribal. Sprint showed modest increases in high variability connections in rural and tribal.






6. Estimated OTT MOS Improves

The improvement in underlying Internet quality metrics, particularly in latency, jitter and packet loss, means a substantial increase in the OTT MOS quality for all four carriers.







This improvement occurred in all geographic categories.



7. LTE Deployment Coverage Appears Capped

The deployment of LTE as well as the replacement of legacy 1G and 2G wireless access networks seems to be stabilizing - particularly for AT&T and Verizon. Sprint and T-Mobile appear to be still catching up with carriers that deployed LTE earlier and have greater rural footprints. There seems to be substantial stability with the percentage of measurements that still use 1G and 2G technologies – both in urban and rural areas.



As the charts above illustrate, legacy 1/2G technologies still represent as much as 20%+ of the measurements in rural areas for Sprint and T-Mobile and appear to have stabilized at ~7% for Verizon and ~2% for AT&T. The use of 1/2G technologies is only a third as frequent in urban areas decreasing to ~7% for spring and T-Mobile and ~2% for Verizon. These legacy technologies have substantial impact on user performance, with dramatically longer latencies and much slower throughput. VoIP services are impossible on these legacy networks and streaming video service unlikely to deliver acceptable quality.



It also appears that the deployment of LTE has peaked in both rural and urban areas for AT&T, T-Mobile and Verizon. AT&T and Verizon both have peaked at about 85% in urban areas, and below 70% in rural. T-Mobile appears to have peaked at 70% in urban and below 50% in rural. Sprint continues to catch up.

8. Dramatic Performance Difference Among Devices

CalSPEED has always included measurements from multiple devices, recognizing that user performance is a composite between the capabilities of the network and the capabilities of each individual user device. And while the network can enforce stringent standards on network service, there is a wide variation between different brands and models of user devices - even in the same year.

The most dramatic difference was between the USB modems (used through fall 2014) and the Android tablets that were introduced in spring 2015 (shown together as "Other" in the adjacent graphs). Thru spring 2013, USB modems and Android smartphones delivered very similar downstream speeds. Beginning in the fall 2013 survey, USB modem performance increased modestly, but were far outdistanced by the improvement in Android smartphones. But the fall 2014 survey, Android smartphones were over 50% faster than the USB modems.

With the upgrade to the Android tablet in spring 2015 - this difference in downstream throughput was erased.




There remains a difference in latency - even with a common operating system between smartphones and tablets. In the spring 2015 survey, while the downstream throughput difference was small, tablets remained with a 30% latency deficit to smartphones.

Note that all devices for these surveys represent the latest wireless technology for each carrier. Older models of user devices, with older wireless technology, will show much more dramatic differences.

9. Performance Difference between East and West Servers Decreases


We measure performance not just to local servers in CalSPEED, we also measure performance to "distant" servers to get measure of how each carrier chooses to connect to the Internet. Since users will be accessing Internet resources located not just locally, but distributed around the U.S. and the world, how each carrier integrates the full Internet as well as local access is a key component of the wireless broadband experience. For these measurements - we have two test servers, one in the San Francisco Bay Area ("West") and one in northern Virginia ("East").

In the best possible time, the physics of data transmissions¹⁰ adds about 80 msec of additional latency to get from one Coast to another - in addition to any local wireless access latency. Additional latency differences over 80 sec suggests carrier Internet routing choices for traffic between East and West. In the case where the latency difference between servers is zero, we speculate that traffic for both servers is peered through a geographically central location, such as a Kansas, where the Internet distance to either the East coast server or the West coast server is essentially the same.



It is interesting to note that the latency difference between East and West servers has decreased over time - converging on 80 milliseconds. We speculate that each carrier is continuing to optimize its networks to deliver consistent performance throughout the Internet - regardless of the location of the content server.

TCP throughput is related to latency ... the longer the latency, the smaller the throughput¹¹. Historically, we have seen that downstream throughput from the East server to California clients is 10-50% less than throughput from the West server. The chart to the right illustrates this trend.



Note the recent (fall 2014 thru spring 2015) decrease in the difference and the convergence of all the carriers on just 10-20% difference between East and West Internet locations.

We speculate, consistent with the other metrics, that we are seeing a more mature, controlled Internet in which consistent performance is more valued than just pure throughput.

¹⁰ Including the speed of light.

¹¹ A consequence of TCP's data reliability and congestion control mechanisms.

10. Conclusions

This is an analysis of the Round 7, spring 2015 dataset for CalSPEED. It shows substantial changes happening in the wireless broadband networks in California for all carriers.

- Major carriers are showing decreasing mean throughput. This slowdown in performance has been apparent for the last two measurement rounds and appears to be a trend. Detailed analysis suggests this is coming from throttling high performance results.
- Broadband coverage is decreasing for all carriers. For the 25 mbps downstream/3 Mbps upstream threshold, Verizon decreased from 16% coverage in fall of 2014 to 4% coverage in the spring of 2015.
- Many measures of quality (latency, jitter, TCP failures, mean packet loss) are improving for most carriers.
- Substantial improvement in estimated OTT MOS is apparent for digital voice service.
- Penetration of LTE in both urban and rural geographic categories appears to have peaked with a floor on 1/2G legacy replacement and on a cap on LTE deployment.
- The performance difference between East and West servers is decreasing suggesting material improvements in backhaul
- In recent rounds, USB modems were ~30% slower than the phones, in Round 7 the tablets (replacing the USB modems) were about ~20% faster than the phones - reflecting dramatic performance differences between the USB modems and the tablets.
- Program change from using USB modems to tablets complicates the analysis. This analysis is based on using the phone measurements alone. Tablets strongly outperform USB modems.

These trends, considered in aggregate, suggest wireless networks that are being managed more carefully to deliver higher quality of service, particularly for VoIP, at the price of decreased throughput. The stall in LTE deployment suggests maturing networks.

Appendix A: CalSPEED: Capturing the End to End User Experience

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 US Internet. One hosted in the Amazon AWS near San Jose, CA and for many California 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 Android smartphone and current Android tablet for laptops. Both are upgraded for each measurement round to match the latest wireless technology deployed by each carrier. Thru Round 6 (fall 2014), Windows laptops with USB modems were used in place of the tablets.

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 - iPerf¹². 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 foundation assumption of CalSPEED, uniquely 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

¹² <http://en.wikipedia.org/wiki/Iperf>

complete network path, from the client device, through the local access network, through the Internet backbone, to several ultimate 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 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 - 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 than 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.

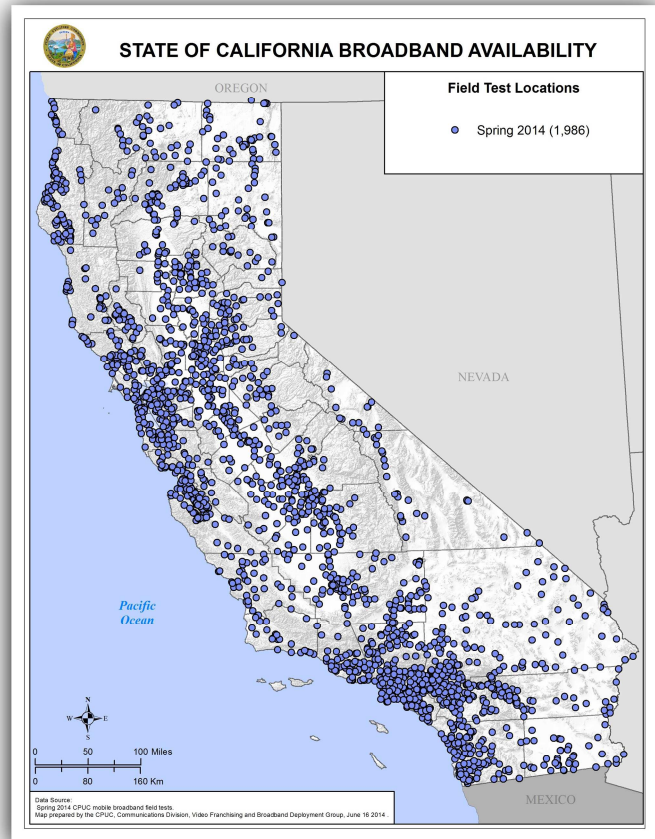
Just the Facts. CalSPEED does not filter any of its results - throughput, coverage, latency or other network metric - rather uses the results of all tests performed and recorded. We believe that just like the user experience with sometimes failing web page loading, all results are valid representing the user experience. Other testing systems filter results in a way that biases results to give a more optimistic expectation of network performance than 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 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 1986¹³ measurement 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.

¹³ Originally 1200, but later increased to improve predictive precision of the interpolation models.

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 crowdsourced 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 (particularly legacy mobile devices) and the performance those user devices are seeing. 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 kriging¹⁴ 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.

CalSPEED has now had six rounds of sampling California (spring 2012, fall 2012, spring 2013, fall 2013, spring 2014, and fall of 2014) and is shortly to finish a seventh round (spring 2015). 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.

¹⁴ <http://en.wikipedia.org/wiki/Kriging>

Appendix B: Terms

Term	Definition
Downstream	The Internet direction from a server to a client.
East Server	Test server located on the East Coast in Northern Virginia
Jitter	The variation in end to end packet latency between user and server.
Kriging	A geostatistical technique for interpolating data from a sample set.
Latency	The end to end round trip delay for a single packet to traverse the Internet from user to server and back.
MOS	Mean Opinion Score. A measurement of VoIP quality
Packet Loss	The rate of loss of packet delivery end to end.
TCP	Transmission Control Protocol. The essential end to end protocol for the Internet that creates a reliable, sequentially delivered byte stream via a sequence of individual IP datagrams.
TCP Connection Failure	Each TCP connection requires a bidirectional packet handshake to initialize data flow. If the handshake cannot occur within a timeout period, the connection fails. The rate of failure is one measurement of the quality of the Internet connection.
Throughput	The number of bytes per second of user data communicated end to end.
Upstream	The Internet direction from a client to a server.
VoIP	Voice over Internet Protocol. Generic name for a family of IP based protocols to replace legacy circuit switched voice with packet based voice.
West Server	Test server located on the West Coast in the San Francisco Bay Area