

CalSPEED: California Mobile Broadband - An Assessment - Spring 2017

Persistent Urban/Rural Digital Divide Persists, Overall Mobile Broadband Returns to Year Ago Quality and Performance and Carrier Performance Throttling Continues

Ken Biba
Managing Director and CTO
Novarum, Inc.

This is an analysis of the measurement Round 11, spring 2017 dataset for CalSPEED. The key results for California mobile broadband include the following:

Persistent mobile digital divide	The rural/urban mobile digital divide appears to be persistent with rural users having material (about 1/3) poorer mobile broadband service over a period of years with little indication of improvement. This is true for throughput, latency, packet loss, MOS, jitter, broadband coverage and deployed technology. Spring 2017 shows additional incremental degradation.
Mean throughput worsens	Mean throughput worsens for all carriers (small increase for Sprint downstream and for Verizon upstream), and returns to levels last seen at least a year ago.
Carrier throughput throttling continues	Throughput throttling for all carriers continues from fall 2016 into spring 2017. The downstream throttling threshold continues from fall 2016 at ~42 Mb/s. The upstream throttling threshold continues from fall 2016 at ~20 Mb/s (15 Mb/s for Sprint).
Mobile broadband coverage decreases	Broadband coverage at the 25 Mb/s down, 3 Mb/s benchmark continues below year ago levels at ~13%. Using the for mobile broadband benchmark of 10 Mb/s down and 1 Mb/s up, improves “coverage” to ~50% for all carriers.
Underlying service quality mixed	Mean rural latency worsens, jitter worsens, packet error rate improves, and TCP connection reliability worsens for all carriers.
Rural and Tribal TCP connection failure rates 3-7x urban	TCP connection attempts fail much more often for rural users than for urban users for all carriers. An urban AT&T or Verizon user can expect 2% of TCP connection attempts to fail (often invisibly retried by the using applications) while ~20% or rural users can expect TCP connection failures.

Over the Top VoIP quality modestly improved, continued relative impairment for rural and tribal	VoIP quality continued improved after substantial degradation in spring 2016. Rural and tribal VoIP quality remains degraded.
LTE deployment coverage has peaked	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. Notably, there was no legacy 1/2G service detected for AT&T in California.
Internet latency stabilizes	The latency difference between east and west servers has stabilized at ~80 msec.
Internet video and conferencing results mixed	Internet OTT streaming video and conferencing availability modestly improved from fall 2016. A “Valley of Video Death” appears as a consistent area of No Service for OTT video in rural Southern California for all carriers.
East/West throughput differences continue at all time high	Ratio of east/west throughput continues pattern of last year with an all-time high of 60% for all carriers but Sprint at 40%.
Little difference between devices	There continues to be no material throughput or latency difference between current smartphones and tablets.
Weak relationship between signal strength and throughput	There is a modest correlation between Signal to Noise Ratio (SNR) and west throughput, but not for Received Signal Strength Indicator (RSSI). Moreover, the wide range in throughput values for any particular SNR or RSSI means both are poor predictors of throughput.

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.

Unlike many speed tests that offer just a horse race between carriers, CalSPEED tries to understand the quality of the mobile user broadband experience.

CalSPEED has now been in use in California for five years with eleven rounds of measurement over the entire state collecting over 30,000,000 measurements, at the same locations, 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 eleven rounds of measurement¹. Previous papers have analyzed the prior rounds of measurement. The methodology has been rigorously analyzed with respect to other available mobile measurement tools². This paper examines the incremental changes from the previous report extending thru the spring of 2017.

¹ Refer to Ken Biba's analysis in the Novarum mobile test reports on the CPUC web site, <http://cpuc.ca.gov/General.aspx?id=1778>

² Ken Biba, "Comparison of CalSPEED, FCC and Ookla", Novarum, Inc., September 2014

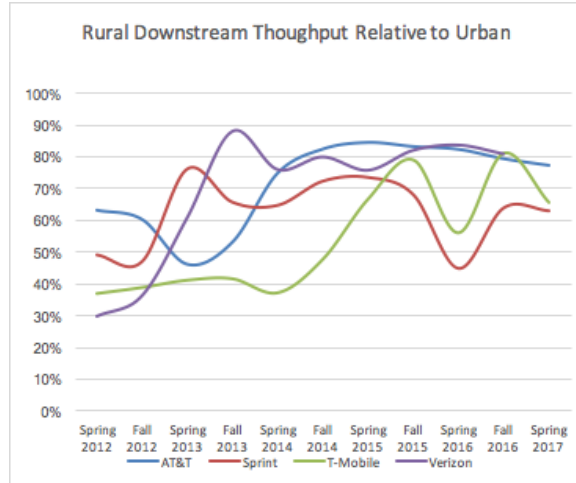
2. Persistent Mobile Digital Divide

A substantial collection of network metrics demonstrates that rural mobile broadband networks consistently (by approximately 1/3) underperform urban networks - both in performance and quality, for all carriers.³ The data strongly suggests that this underperformance has continued for years, and trend lines suggest little improvement.

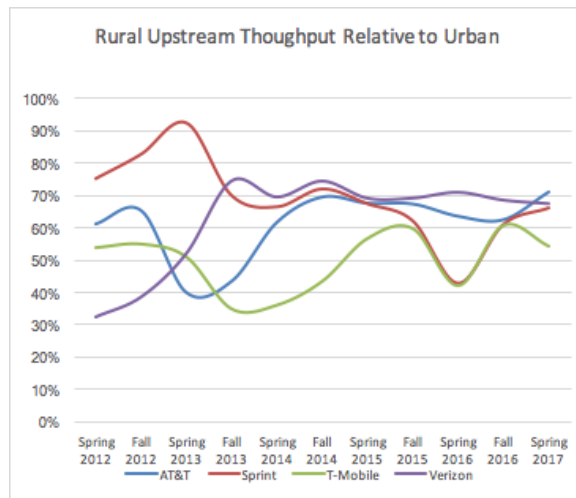
CaISPEED metrics help document this persistent digital divide.

³ This analysis uses the urban vs. rural delineation as assigned in the 2010 Census Blocks geography. The classifications are defined here: <https://www.census.gov/geo/reference/ua/urban-rural-2010.html>. For the 2010 Census, an urban area will comprise a densely settled core of census tracts and/or census blocks that meet minimum population density requirements, along with adjacent territory containing non-residential urban land uses as well as territory with low population density included to link outlying densely settled territory with the densely settled core. To qualify as an urban area, the territory identified according to criteria must encompass at least 2,500 people, at least 1,500 of which reside outside institutional group quarters. The Census Bureau identifies two types of urban areas: a) Urbanized Areas (UAs) of 50,000 or more people; b) Urban Clusters (UCs) of at least 2,500 and less than 50,000 people. "Rural" encompasses all population, housing, and territory not included within an urban area.

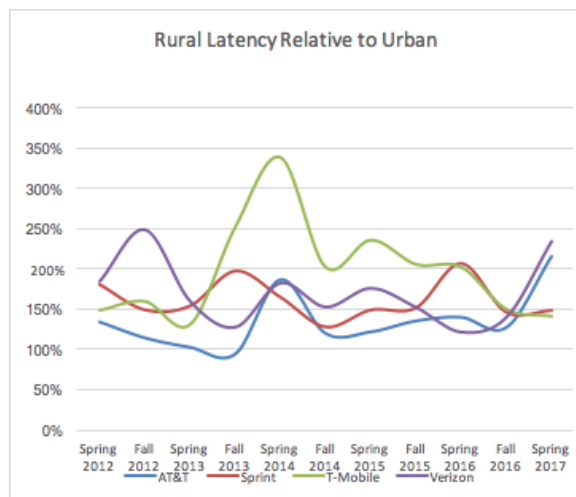
Rural mean downstream TCP throughput is consistently only 80% (and for Sprint and T-Mobile sometimes as little as 50%) of urban.



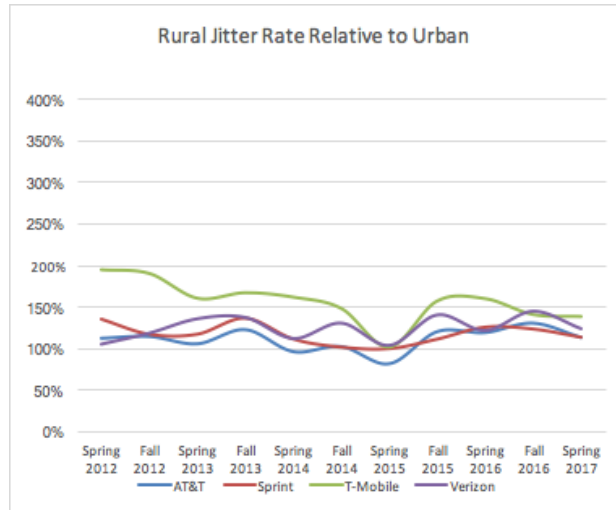
Rural mean upstream TCP throughput is consistently only 60-70% of urban.



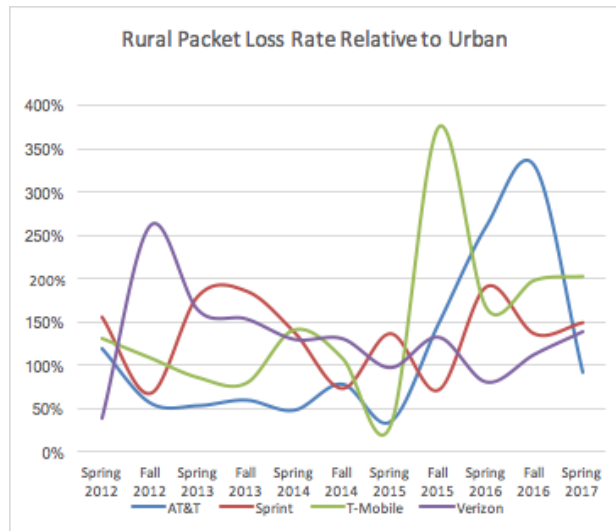
Rural mean latency is consistently at least 50% (T-Mobile and Sprint) to 100% (AT&T and Verizon) worse than urban.



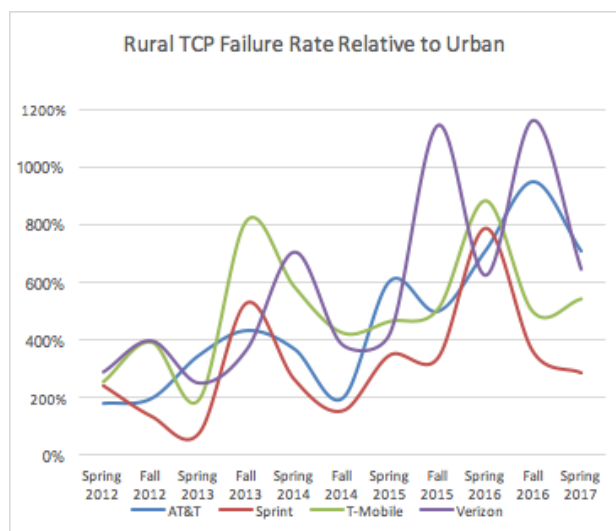
Rural mean jitter is consistently at least 20% worse than urban.



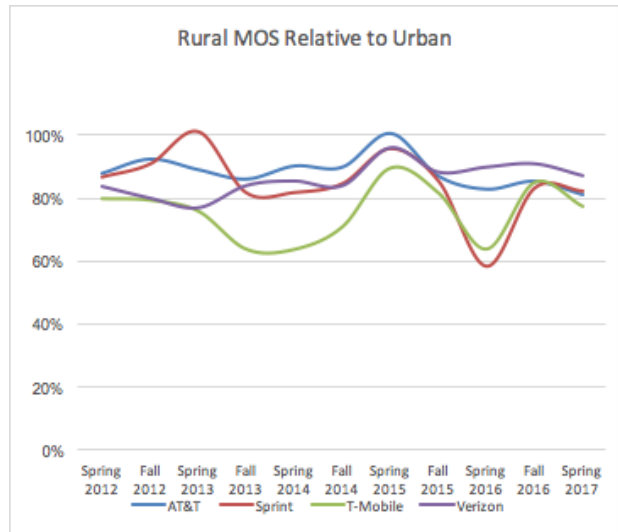
Rural packet loss is consistently 100-200% worse than urban. The poor performance of rural latency, jitter and packet loss strongly implies that real-time streaming voice and video services will perform materially worse in rural than urban areas.



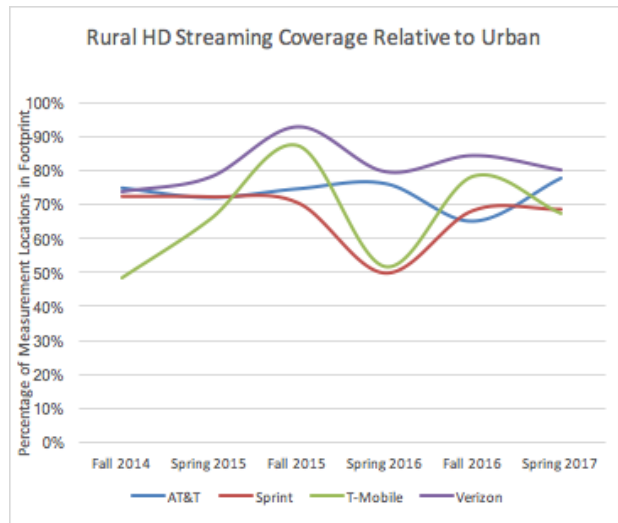
Rural TCP connection failure rate is consistently ~5x worse than urban with indications that this disparity is getting worse. For rural and tribal users of all carriers, about 1 in 5 TCP connection attempts fail.



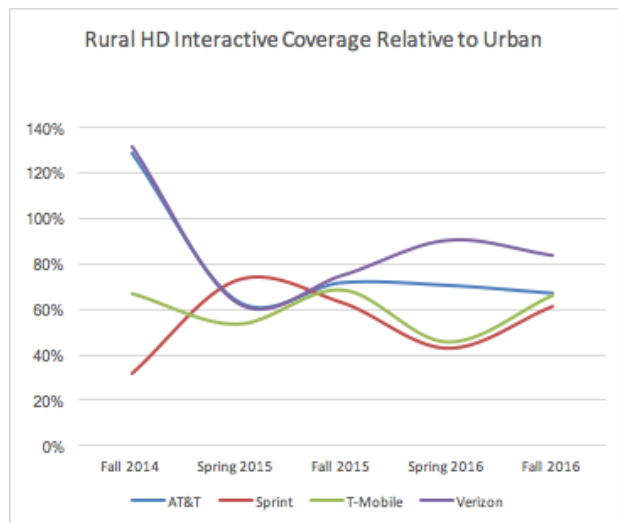
Rural OTT digital voice coverage is consistently only 85% of urban.



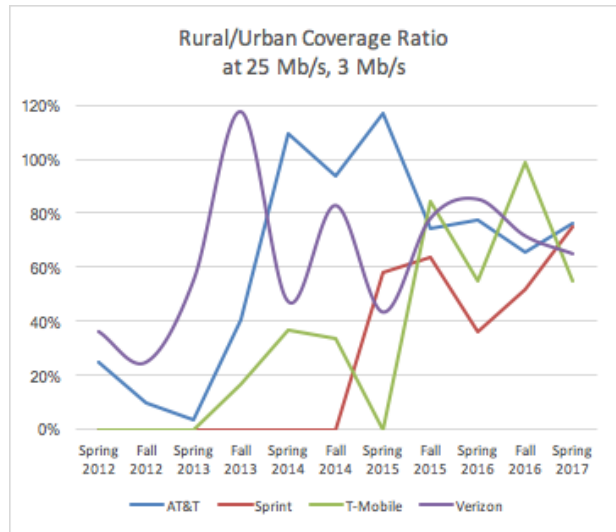
Rural streaming video coverage from the west server is consistently 60-80% of urban. Rural users will encounter no streaming availability 3x more frequently than urban users.



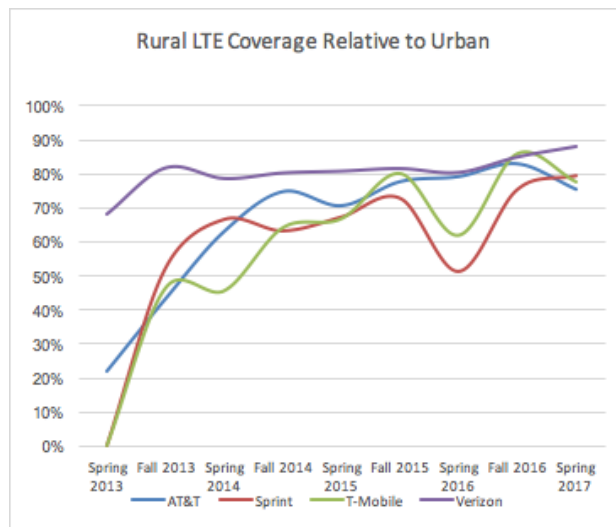
Rural conference video coverage is consistently 60-80% of urban. Rural users will encounter no conference availability 2x more frequently than urban users.



Rural mobile broadband coverage at the 25 Mb/s down and 3 Mb/s up benchmark is consistently 60-80% of urban.

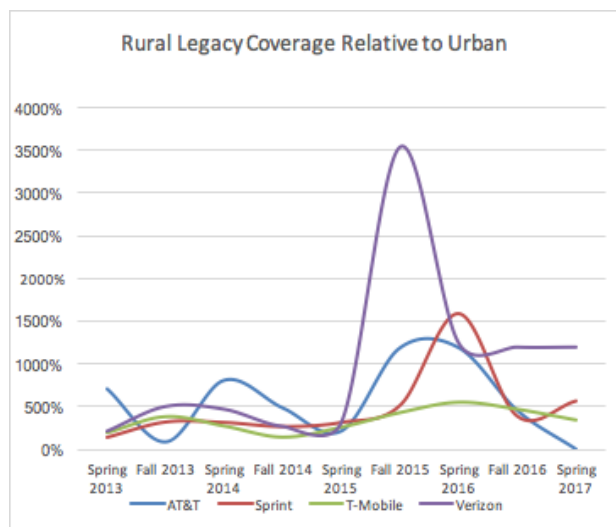


Rural LTE coverage is consistently 80% of urban. The trend suggests rural LTE availability will cap out below 85% of urban. This has substantial implications for the use of mobile broadband as a replacement for wired broadband in rural areas and for the capability of public safety services.



Rural mobile broadband consistently has more than 500% use of legacy 1G and 2G mobile broadband technology than urban. It is much more likely for a rural user to encounter very poor legacy mobile broadband service than an urban user.

A notable change in spring 2017 is the effective elimination of legacy rural coverage for AT&T.

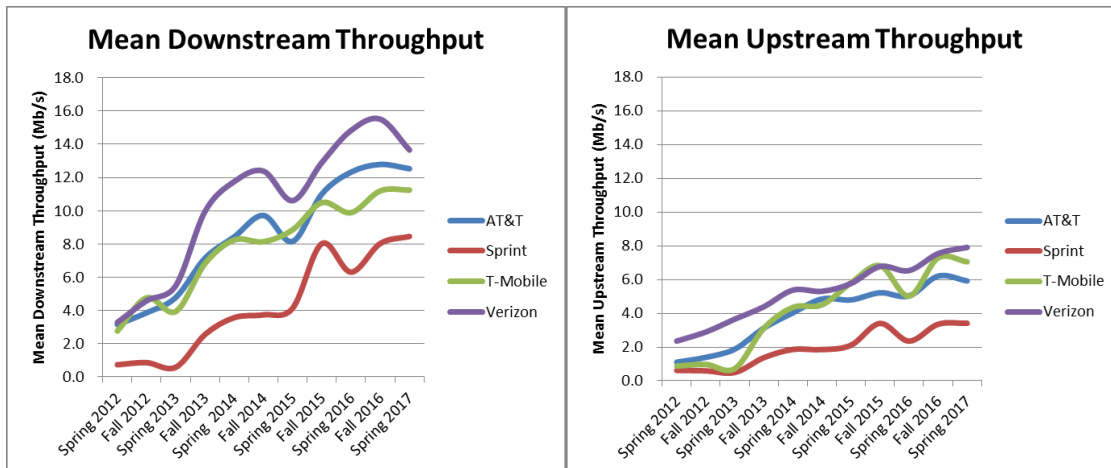


3. Mean Throughput Goes Back in Time

All mobile carriers show variation over time of mean downstream and upstream throughput, and it is expected that there will be variation from measurement period to measurement period. The general trend has been towards higher performance as carriers deploy newer, higher performance infrastructure technology and more cell towers. Variations in performance can be interpreted as correlated to changes (and mismatches) in network capacity and user demand.

For ease of analysis and consistency across all rounds of testing, we examine data from only smartphones⁴. As shown in the graphs below, mean downstream throughput decreased in spring 2017 for (almost) all carriers. Mean downstream throughput levels for AT&T, Sprint and Verizon returned to levels from a year ago- and T-Mobile showed no improvement from fall 2016 and only small improvement since fall 2015. There is a suggestion that historic throughput increases have leveled off. Further measurement rounds will help decide.

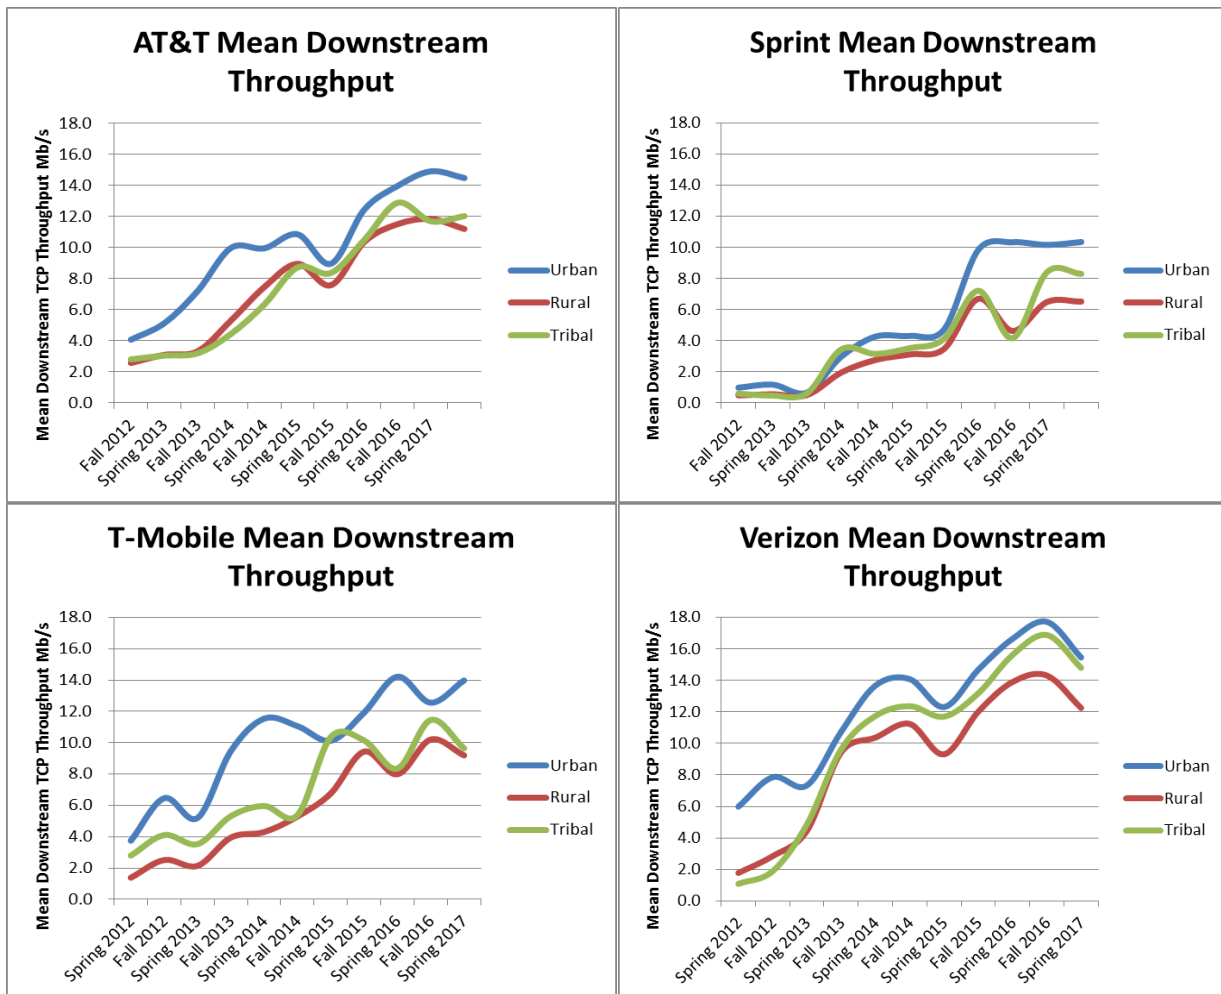
Mean carrier upstream throughput also regressed to levels from at least a year ago for AT&T, Sprint and T-Mobile.



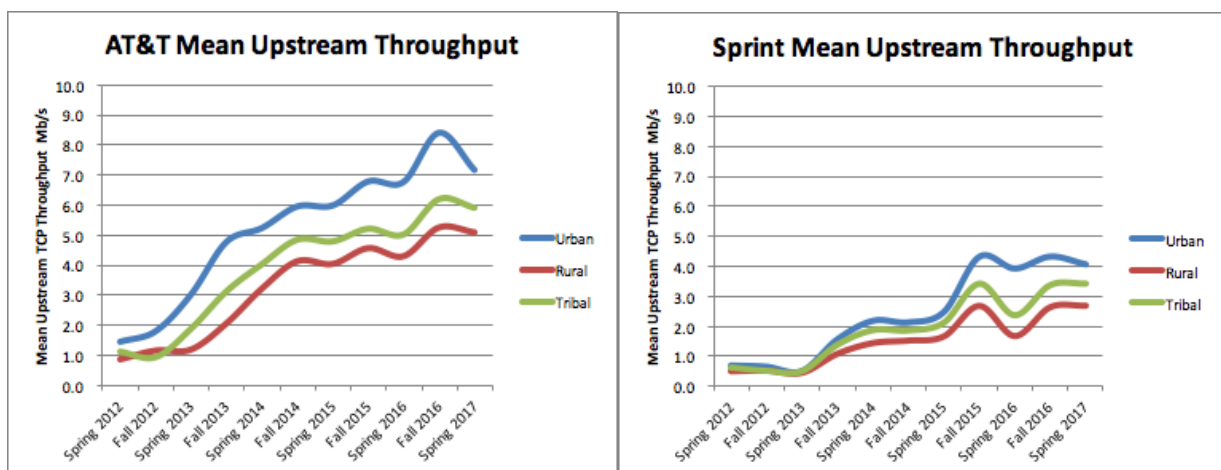
This regression of mean throughput was not limited to one demographic - we can see that all demographics, for most (AT&T, Sprint and T-Mobile) carriers show weakness in mean downstream AND upstream.

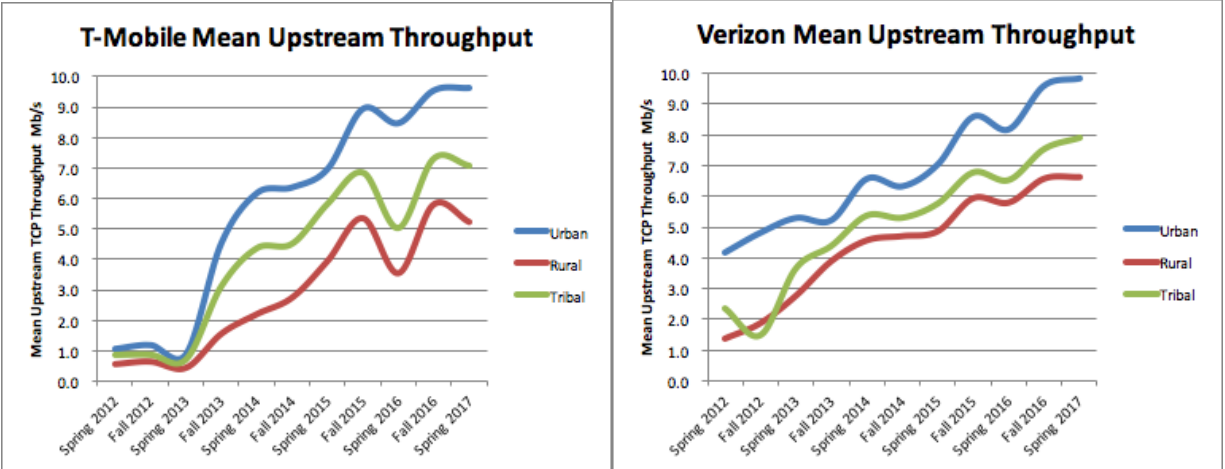
⁴ For some of the earlier rounds of mobile field testing, we tested smartphones and USB modems side by side. Beginning with spring 2015 testing we replaced the USB modems with tablet devices. We see very little difference in performance between the tablets and the smartphones.

Downstream Throughput by Testing Round in Megabits-per-Second (Mb/s)



Upstream Throughput by Testing Round in Megabits-per-Second (Mb/s)



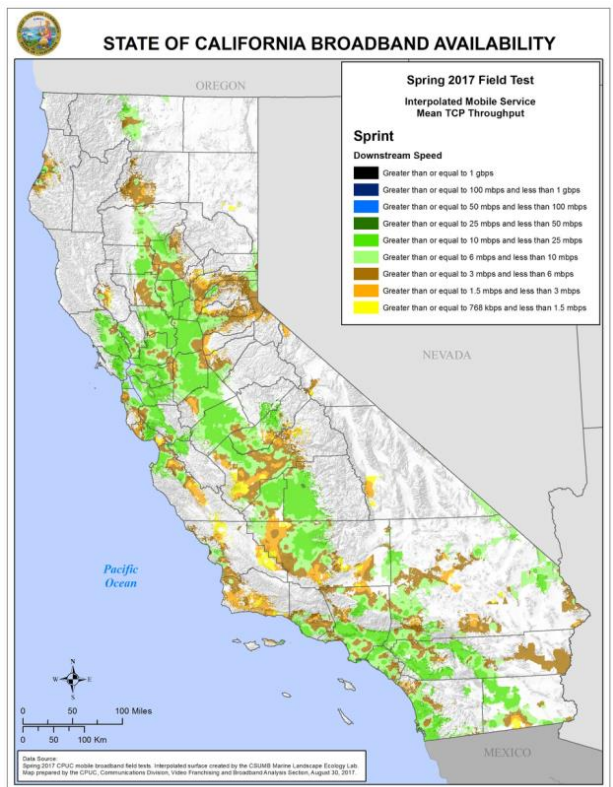
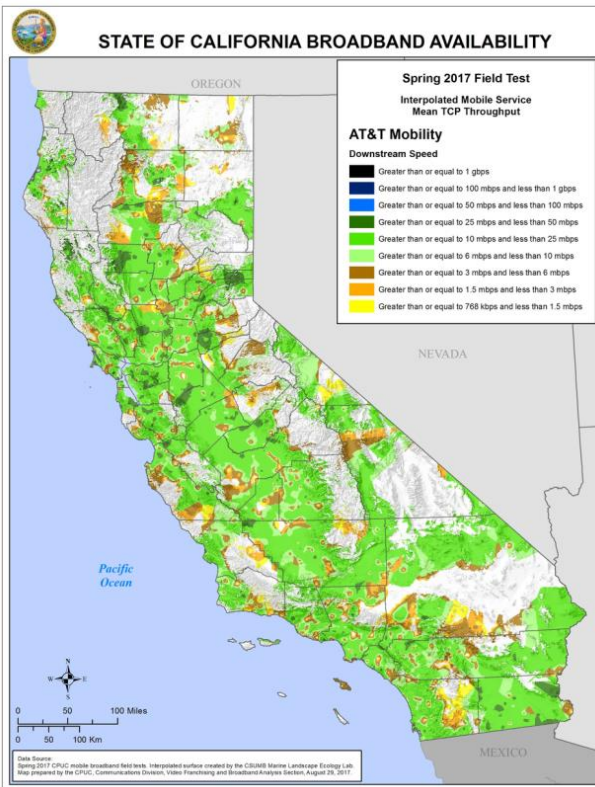


CalSPEED maps of the estimated mean downstream and upstream throughput across the state are clipped to the carrier-supplied coverage footprint.

First, estimated mean downstream throughput. AT&T and Verizon clearly have both a greater coverage footprint in California. Sprint and T-Mobile have a much smaller geographic footprint, particularly in rural California. Their rural coverage is largely confined to major road transportation corridors.

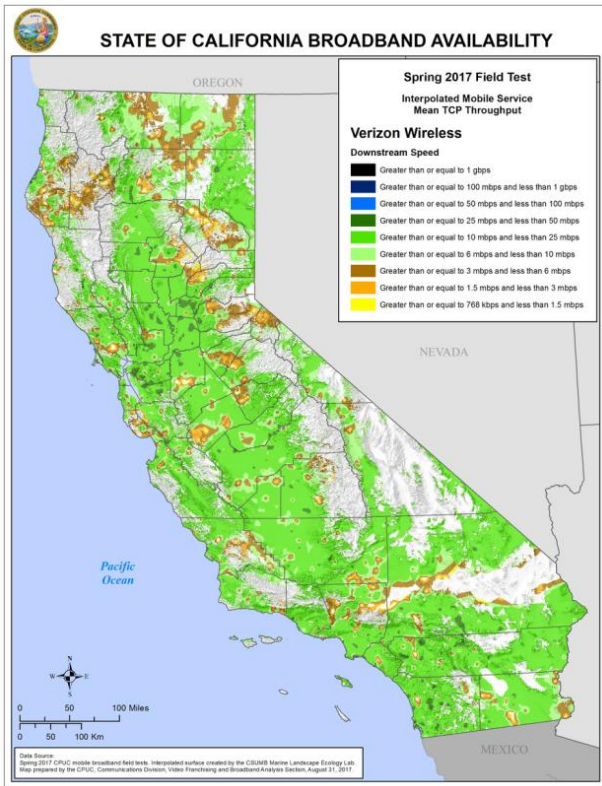
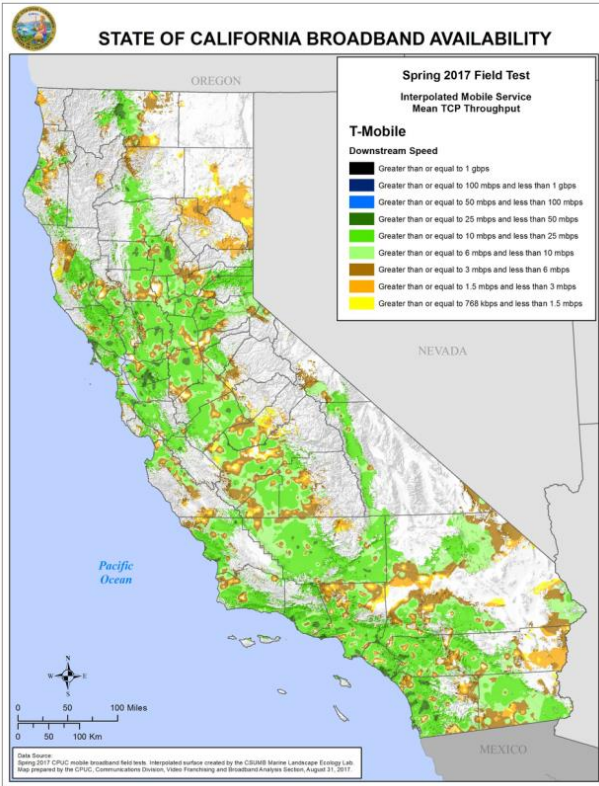
AT&T Downstream Coverage and Speed

Sprint Downstream Coverage and Speed



T-Mobile Downstream Coverage and Speed

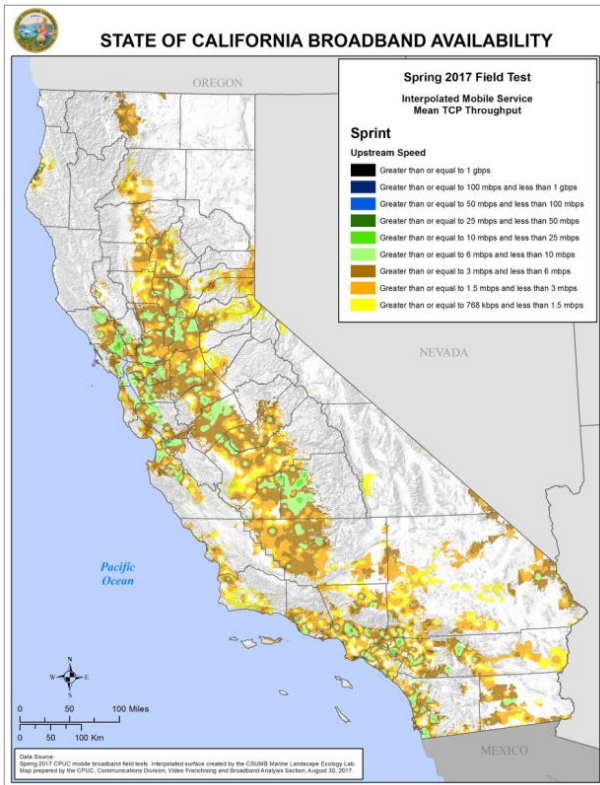
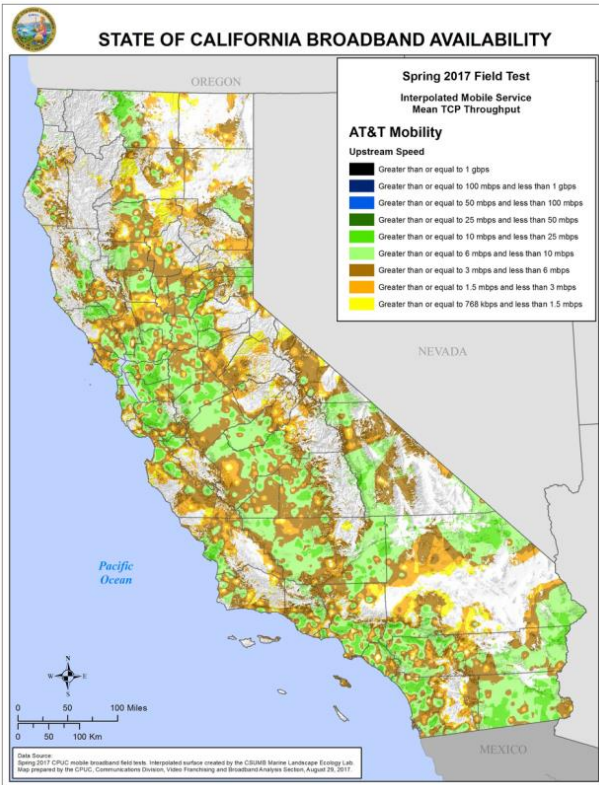
Verizon Downstream Coverage and Speed



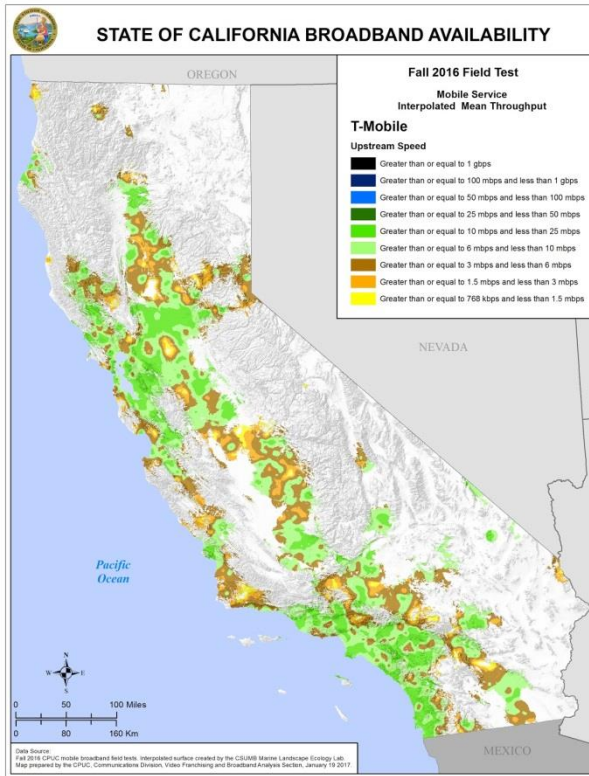
And then the estimated mean upstream throughput.

AT&T Upstream Coverage and Speed

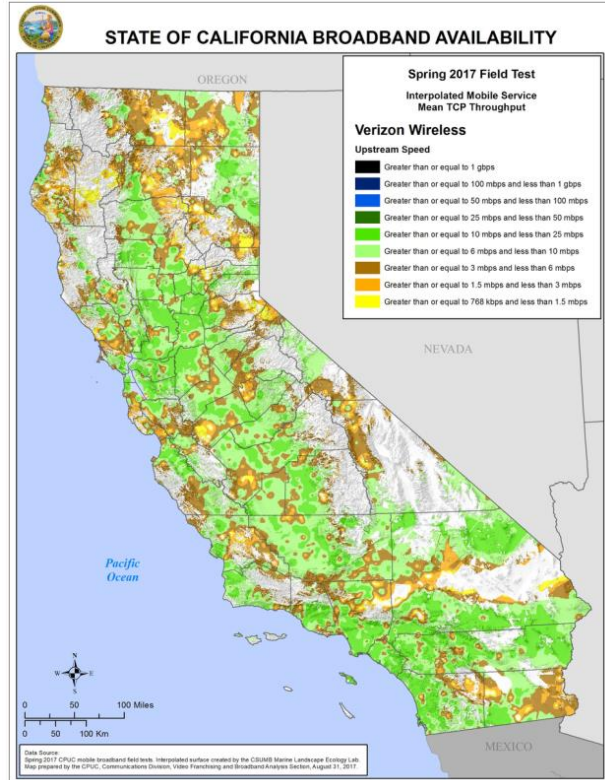
Sprint Upstream Coverage and Speed



T-Mobile Upstream Coverage and Speed



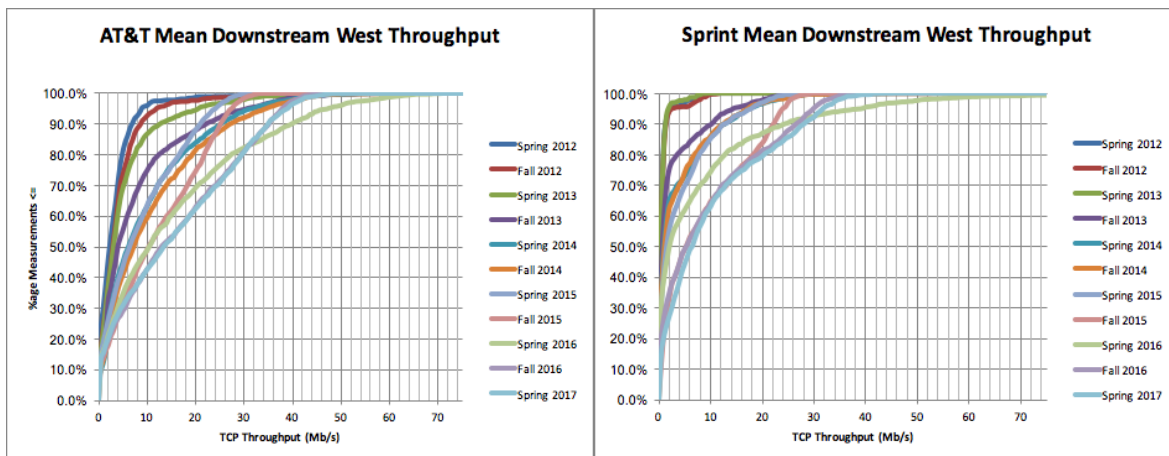
Verizon Upstream Coverage and Speed

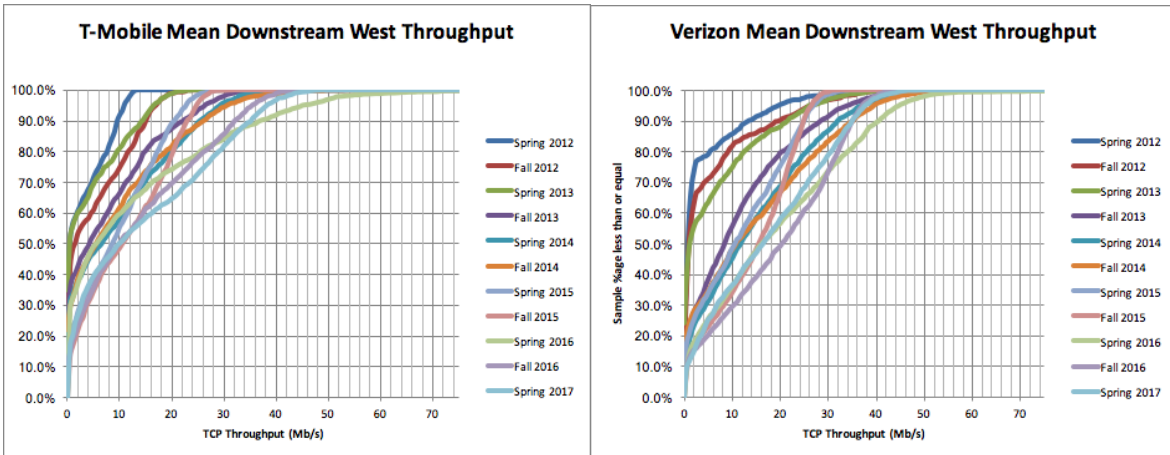


4. Performance Throttling Continues, Again

A pattern is appearing in possible high-end throughput throttling. The throttling observed in spring and fall 2015, which was not observed in spring 2016, now appears in fall 2016 and spring 2016.

This behavior can be seen in histogram plots of all the throughput measurements for each carrier.

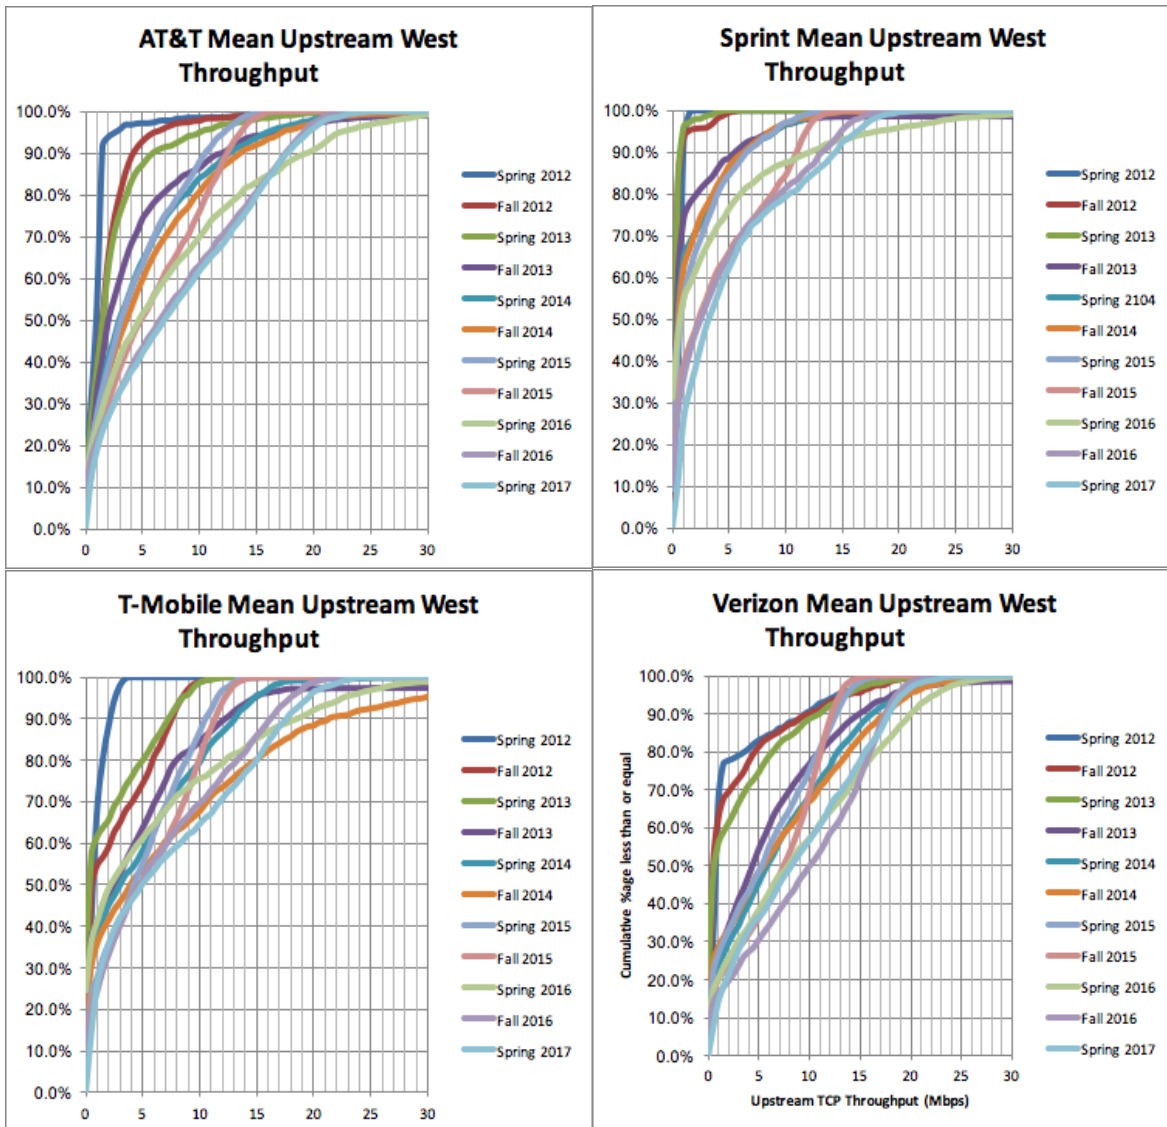




Spring and fall 2015 measurements showed clear downstream (and upstream) throughput throttling for all carriers, as demonstrated in the above cumulative histograms of downstream throughput for phones. We can see a clear change in the distribution for spring 2015. In prior measurement rounds we can see 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 with a clear common boundary of about 30 Mb/s. This strongly suggests intentional throttling of high performance throughput results.

This apparent throughput throttling no longer appeared in the spring 2016 measurements with strong tails to the high end of the throughput distribution with some very high performance results for all carriers. The throttling, for all carriers, had returned in the fall 2016 and continued in the spring of 2017. The downstream throttling threshold seems to have increased from the fall 2015 level of 30 Mb/s to 40 Mb/s in fall 2016.

The same pattern, at the same times, is also seen in the upstream throughput. The upstream throttling threshold seems have increased to 20 Mb/s (15 Mb/s for Sprint) in fall 2016/spring 2017 from 15 Mb/s (12.5 Mb/s for Sprint) in fall 2015.



5. Broadband Coverage Decreases

Let's first look at broadband coverage at the 25 Mb/s down/3 Mb/s up⁵ broadband benchmark, then at the 10/1 benchmark⁶.

In the spring of 2015, there was a dramatic (~50%) decrease for AT&T, T-Mobile and Verizon in the number of measurement locations meeting the 25/3 benchmark. For the spring 2016 measurement round a similar dramatic increase in the number of locations meeting that benchmark is seen. The fall 2016 measurement saw a material decrease in the number of locations meeting 25/3 throughput.

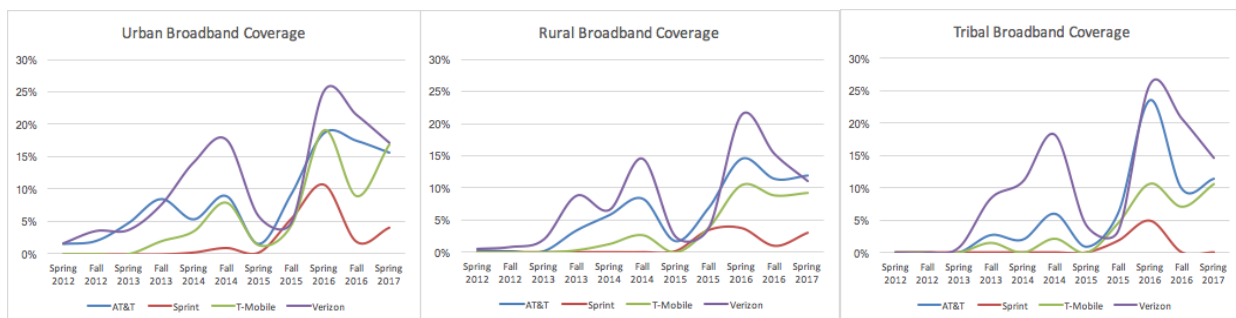
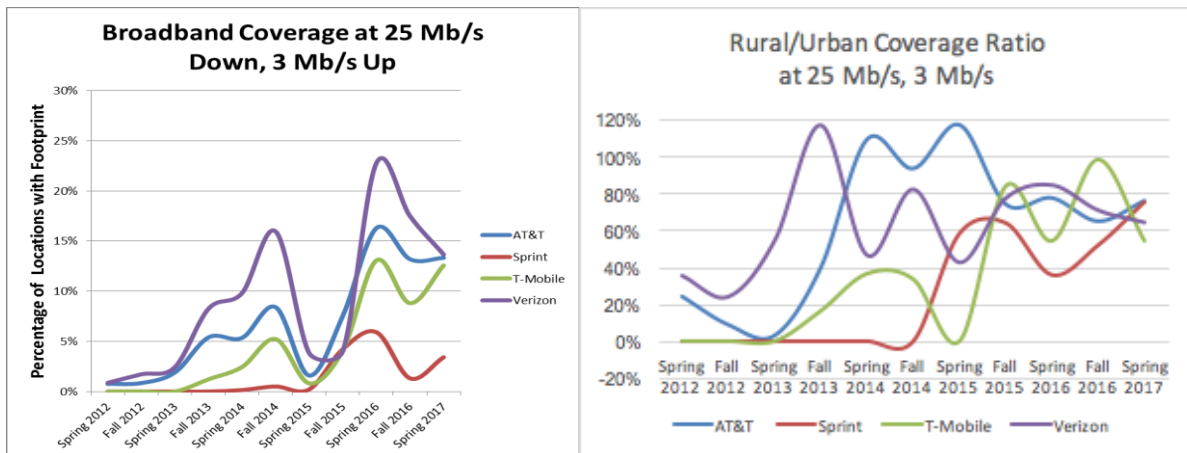
⁵ The FCC established the 25/3 speed benchmark in their 2015 Broadband Progress Report to measure progress in the deployment of fixed broadband. See FCC GN Docket No. 14-126.

⁶ Under the FCC's Connect America Fund Phase II rules, service providers must offer broadband at speeds of at least 10 megabits per second (Mbps) downstream and 1 Mbps upstream with latency at or below 100 milliseconds round trip. See <https://www.fcc.gov/consumers/guides/connect-america-fund-phase-ii-faqs>.

Spring 2016 measurements illustrate a convergence for AT&T, T-Mobile and Verizon at ~13% coverage.

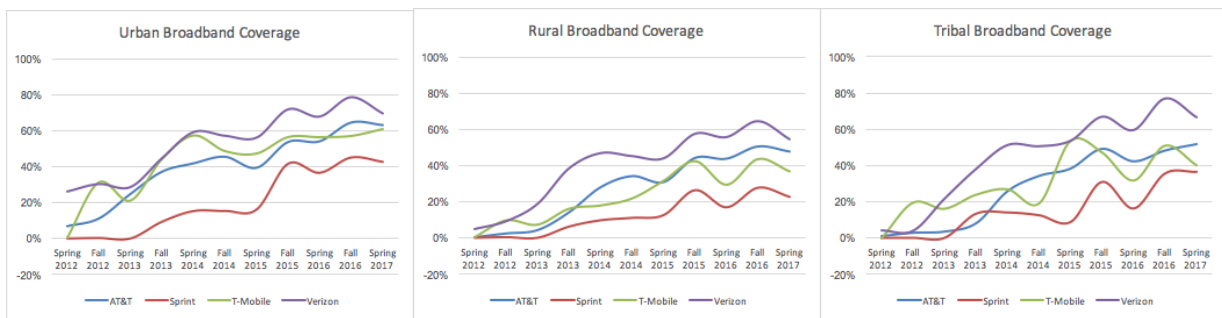
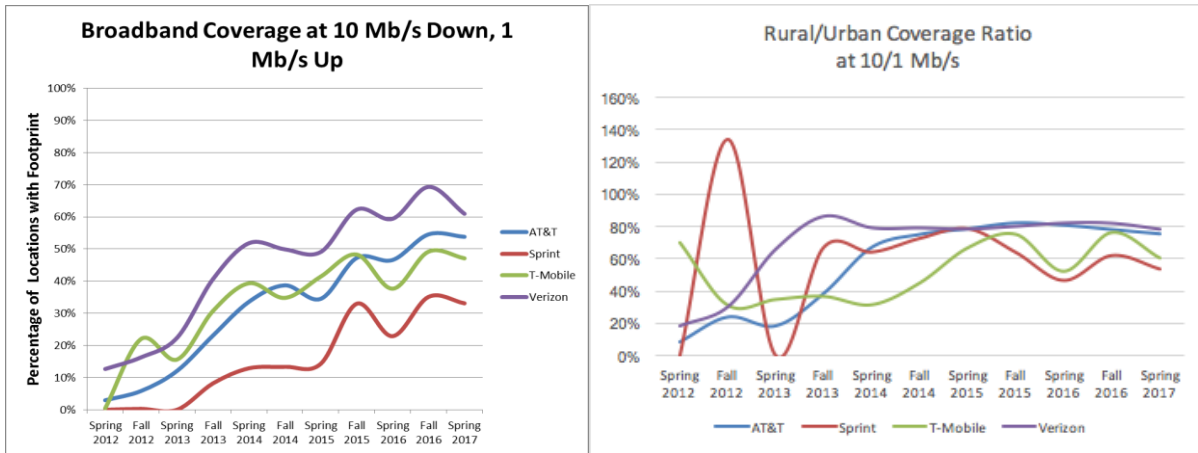
What is also interesting in the relative rural/urban coverage is the apparent stabilization at the ratio of ~70% for all carriers.

This same coverage pattern is observed for all carriers and for all demographics. Sprint uniquely has lower coverage in all demographics than the other three carriers.



When looking at the 10/1 benchmark, we see substantive differences from the 25/3 benchmark.

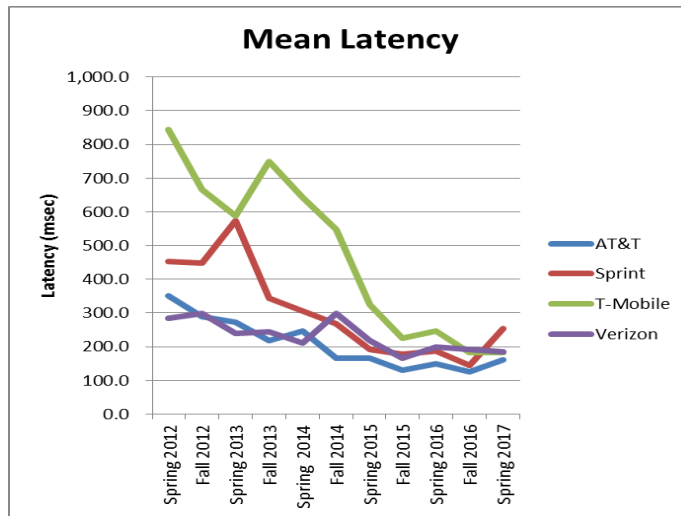
First, the average broadband coverage at 10/1 is substantially higher than that at 25/3, at ~50% rather than the ~13% for 25/3. Second, the coverage at 10/1 has grown over the last 5 years at a steady rate rather than the dramatic swings at 25/3. Third, the rural/urban ratio is surprisingly similar at 10/1 (~70%) to 25/3 (70%).

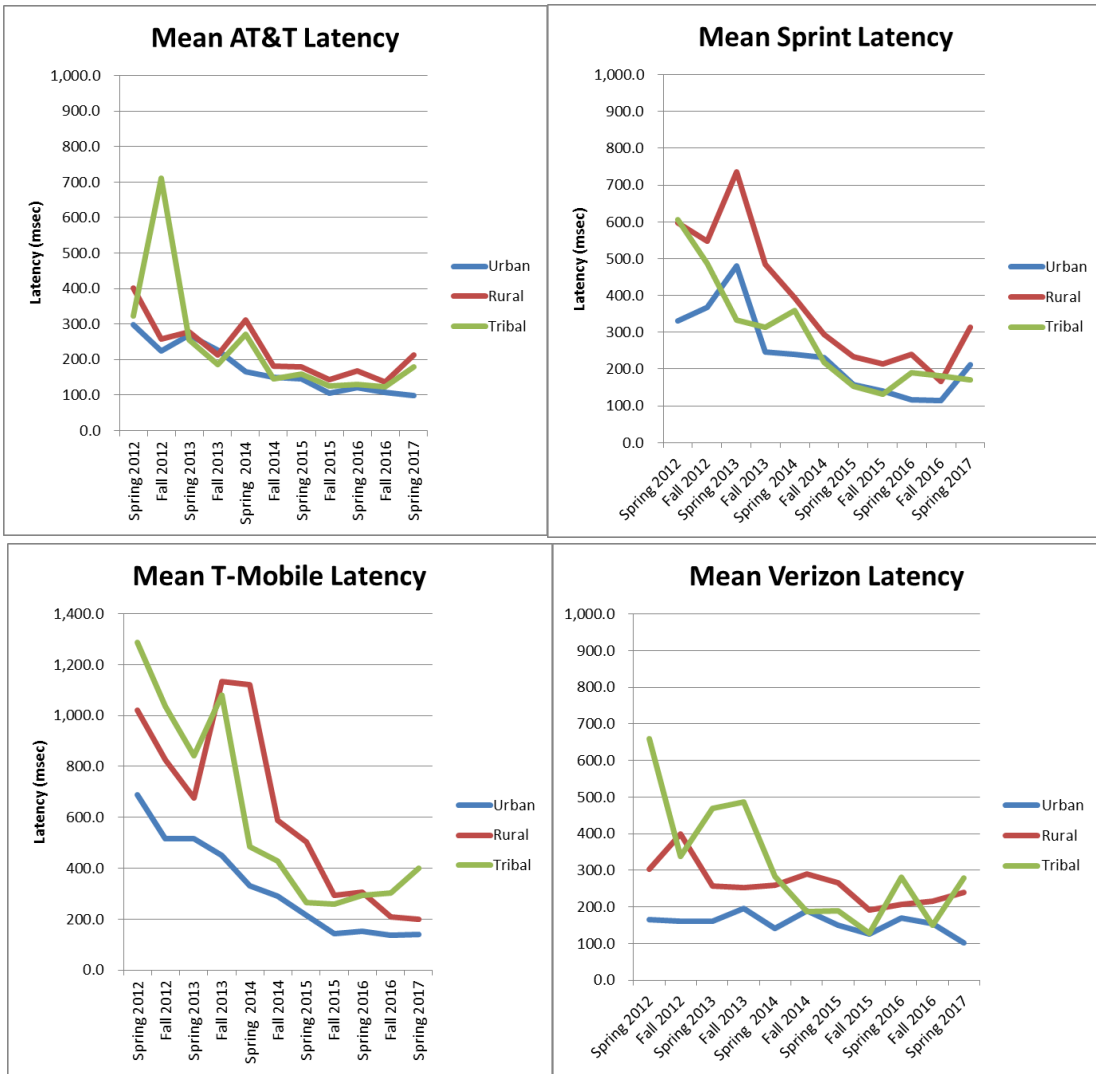


6. Measures of Service Quality Show Mixed Change

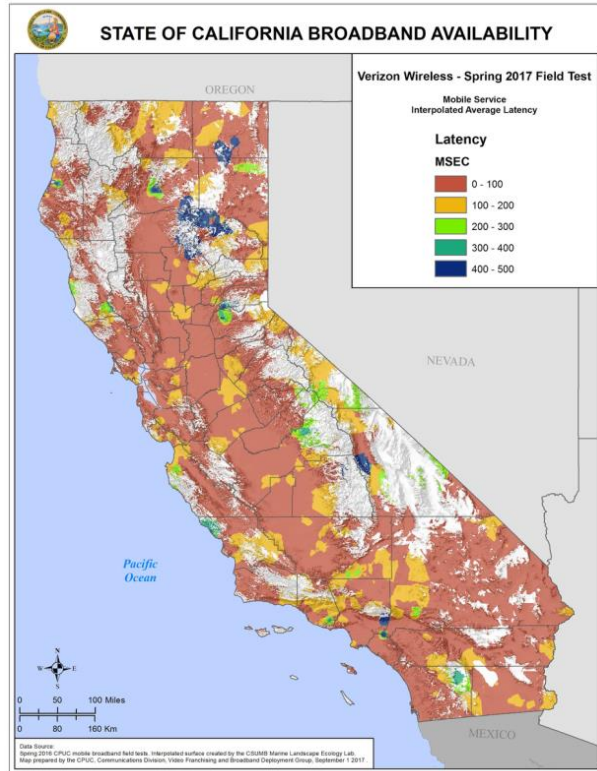
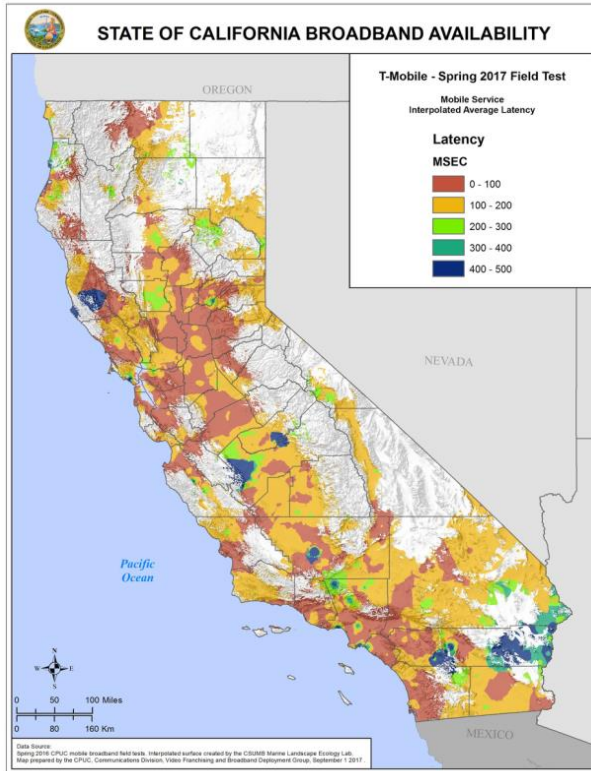
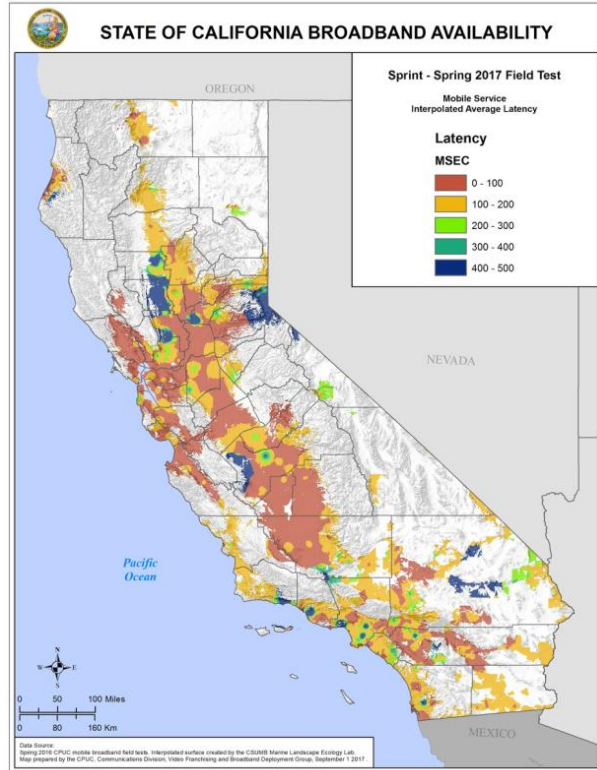
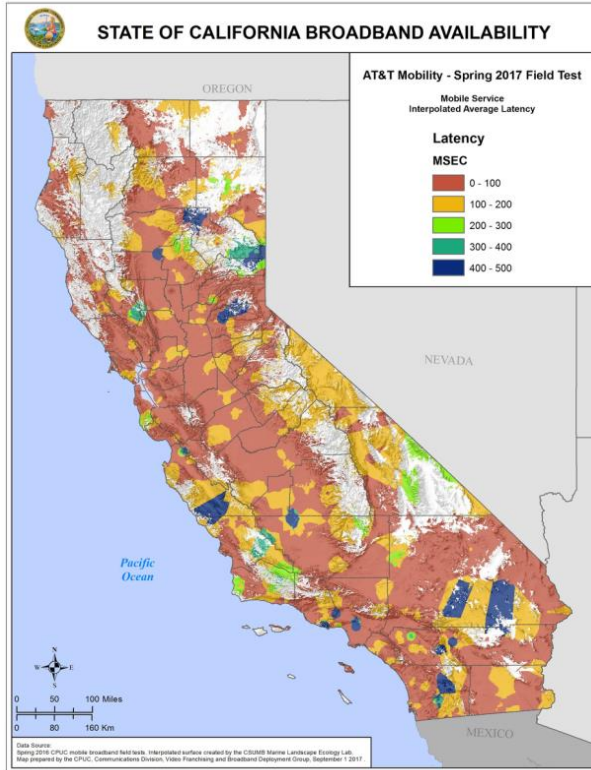
There is a mixed story in spring 2017 measurements of network quality.

Mean latency (measured in milliseconds) worsened for both AT&T and Sprint and remained about the same for T-Mobile and Verizon. Demographic variation was somewhat unique for each carrier. For AT&T and Verizon, urban latency improved with rural and tribal worsened. For T-Mobile, there was modest improvement in urban and rural latency while tribal latency materially worsened. Sprint exhibited dramatic degradation in rural and urban latency, while tribal remained essentially unchanged.

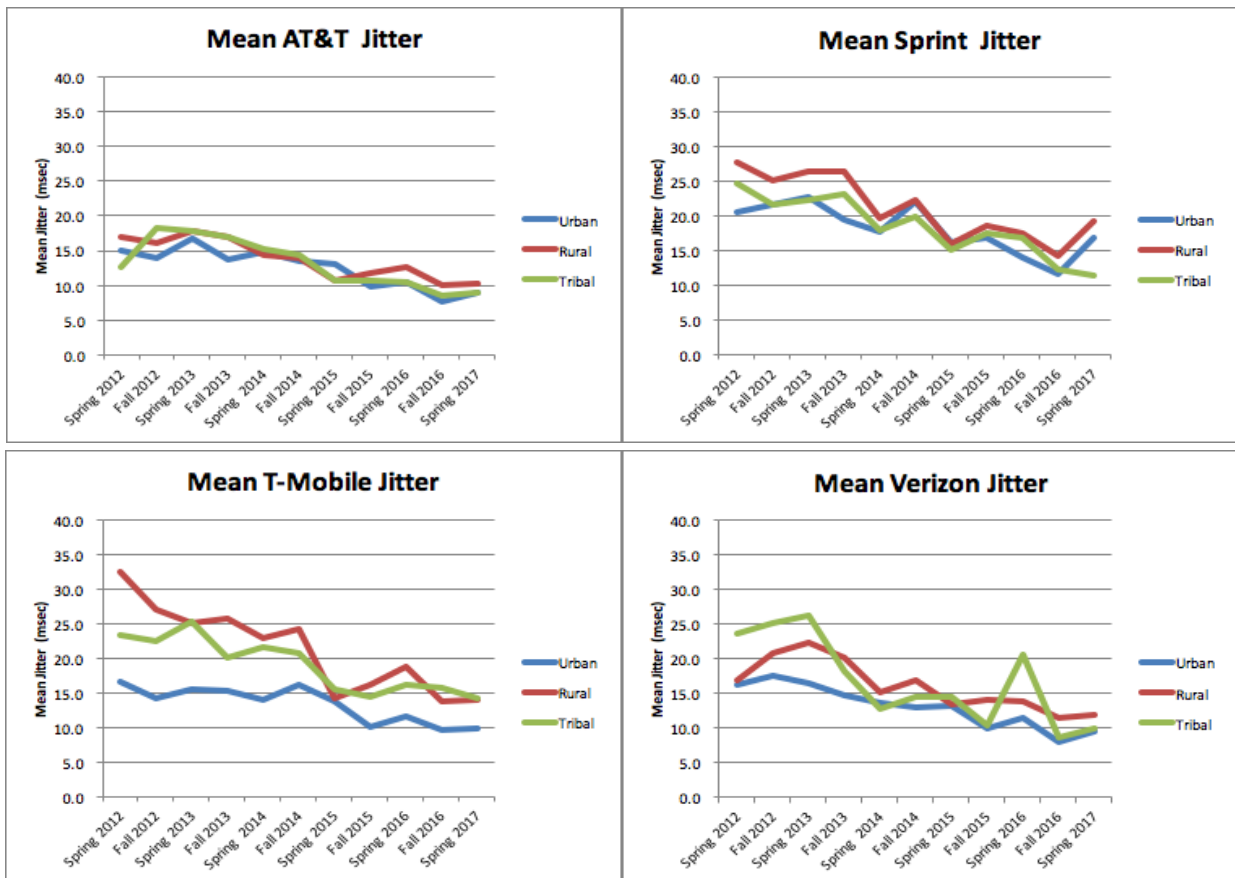
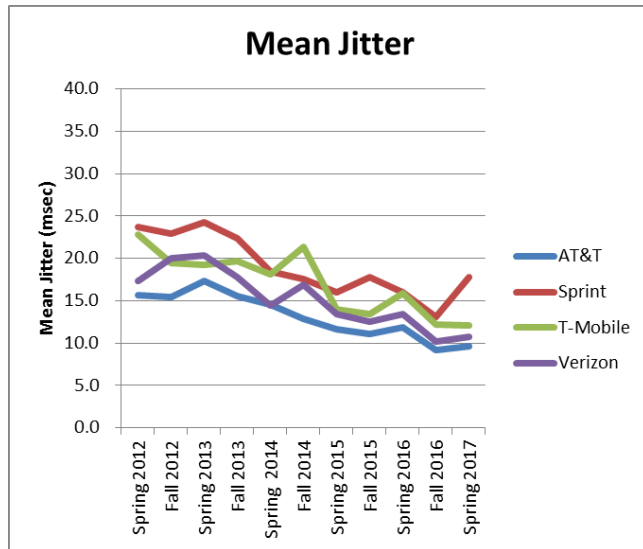




Latency varies across the state for all carriers. CalSPEED's latency maps are particularly useful for identifying specific areas of poor latency. In the following images, the warmer the color, the lower the latency




Mean jitter⁷ worsened for AT&T and Verizon for all demographics. Sprint demonstrated material degradation in urban and rural while showing improvement in tribal. T-Mobile showed degradation in tribal latency while staying constant in urban and rural areas.




⁷ Jitter is the measure of variation in latency. It is an important component of performance for real time streaming of audio and video.

Mean packet loss⁸ rates improved for AT&T, T-Mobile and Verizon and substantially worsened for Sprint. Urban packet loss degraded for AT&T and Sprint. Rural packet loss materially worsened for Sprint. Tribal packet loss improved for all carriers. On average rural has 50% more packet loss than urban for all carriers.

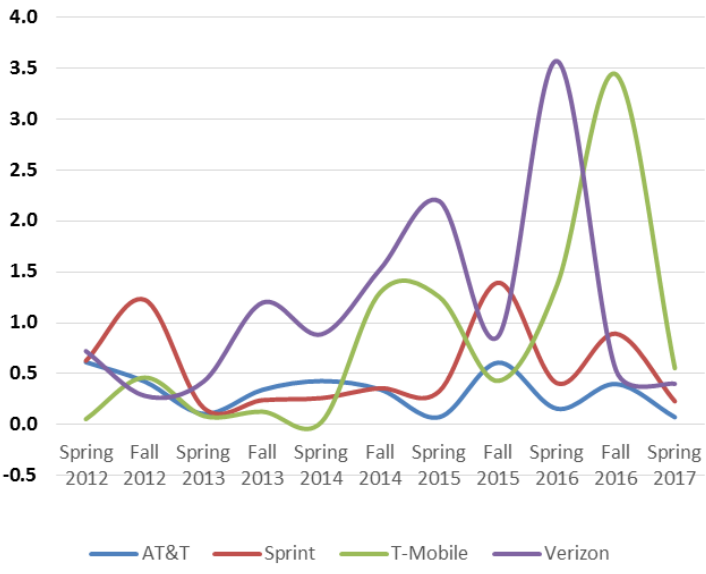



⁸ Mean packet loss is the average percentage of packets that are lost during transmission. Small increases in packet loss are particularly degrading for streaming media services.

Rural Packet Error Rate




Tribal Packet Error Rate








TCP connection quality⁹, as measured by failed TCP connection attempts¹⁰, worsened for all carriers. This degradation is almost entirely in rural and tribal TCP connection quality with modest improvements in urban quality. There remain profound differences between urban and rural areas with rural users having an almost 300-700% higher TCP connection failure rate over urban users even for AT&T and Verizon.




⁹ 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.

¹⁰ 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 that the call fails to connect to the destination.






7. Over-The-Top Voice MOS Service Quality Stable


The following graphs track the number of test locations where the imputed Mean Opinion Score (MOS), which is an estimate of voice call quality, equals or exceeds 4.0 for over-the-top (OTT) voice services like Skype. The mixed results of quality metrics -- particularly in latency, jitter and packet loss -- translated into a flattening of the curve between spring 2015 and spring 2016, we noticed a drop in MOS 4.0 locations, but this degradation reversed in fall 2016. The reversal reflects improvement in the underlying latency, jitter and packet loss metrics. However, it appears that OTT VoIP coverage has still not returned to the peak of spring 2015.



While urban MOS remained stable for all carriers, all carriers still show material degradation for rural and tribal users. In particular, AT&T and Sprint showed dramatic improvements in rural and tribal OTT MOS quality. Sprint, T-Mobile and Verizon rural and tribal users show some (partial) recovery but AT&T rural and tribal users show materially increased degradation.



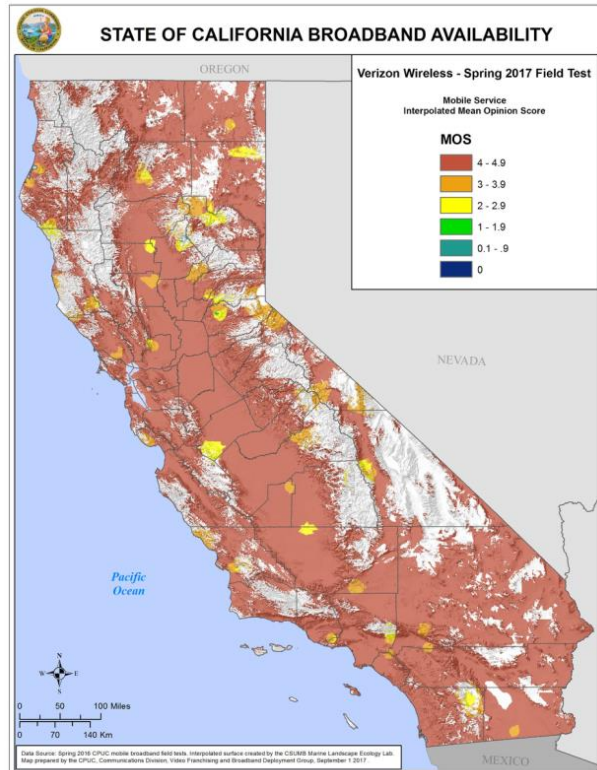
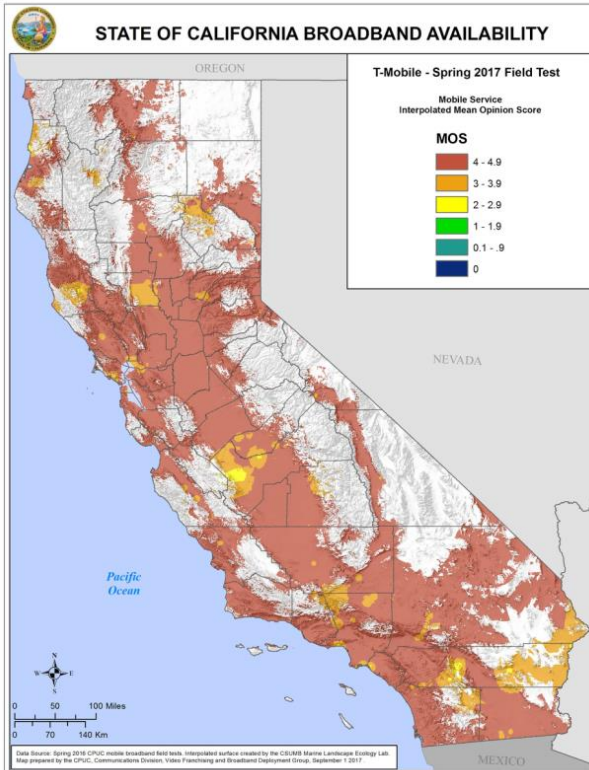
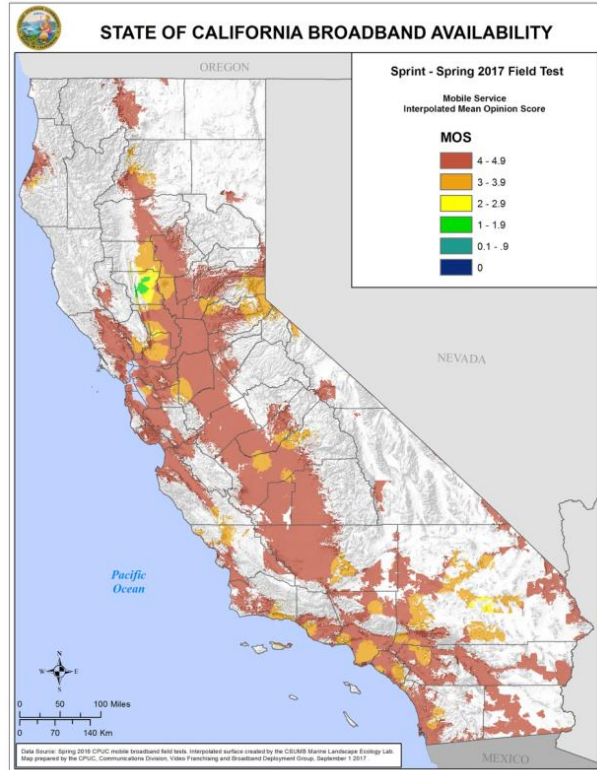
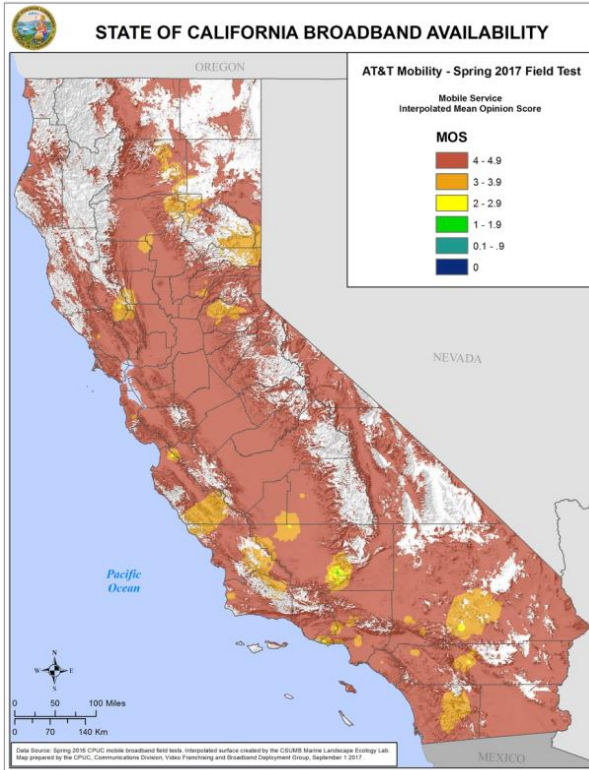




The ratio of average rural OTT MOS quality to urban has been stable during the entire time CalSPEED has been measuring OTT MOS - at ~85%.

CalSPEED maps OTT voice MOS across the state for each carrier. AT&T and Verizon both have extensive OTT voice service across the state, with Sprint and T-Mobile limited to dense urban areas and major transportation routes.

Note that every carrier has spots of poor service in which OTT voice (and likely video) will likely not be possible due to quality. OTT voice is likely not reliable in any location not colored red. As with the latency maps, the warmer the color, the better the performance, with red being voice-grade quality.



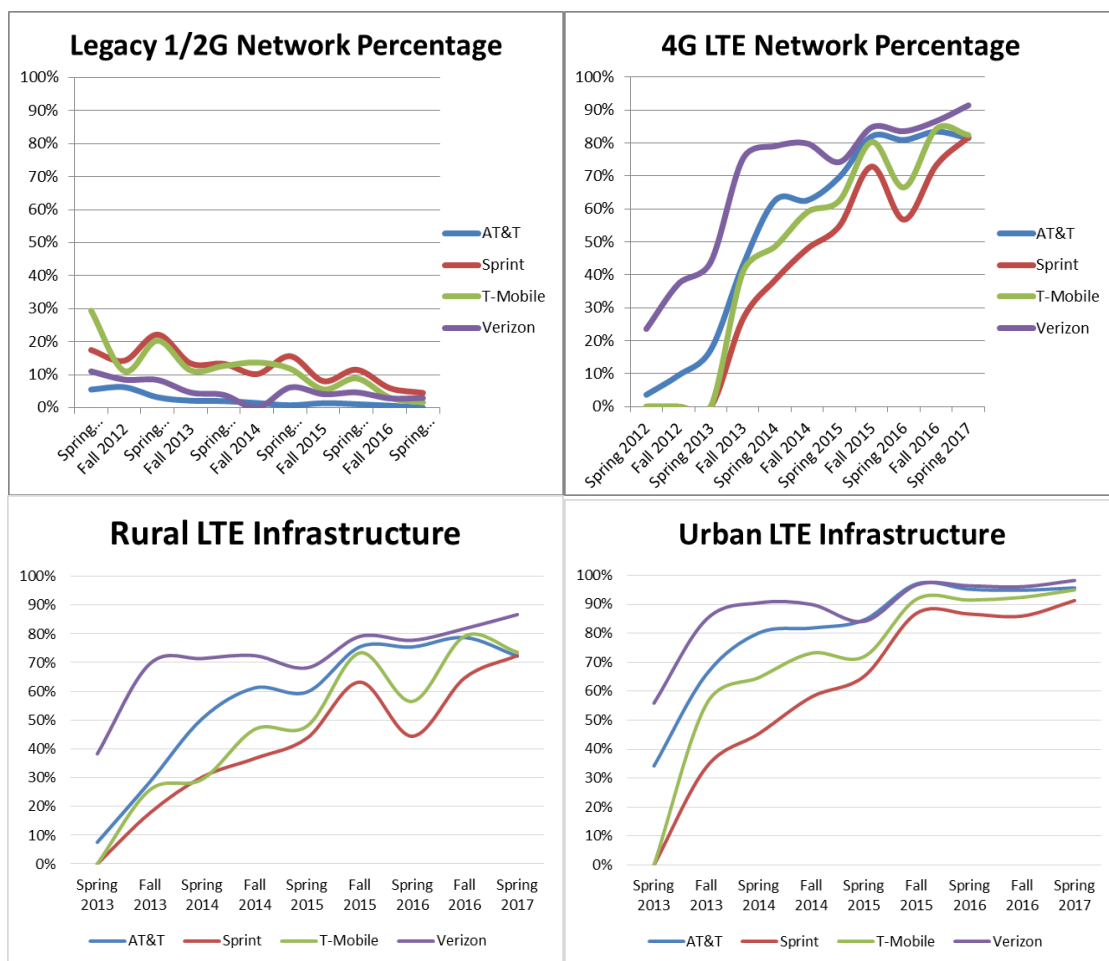
8. Technology Deployment Appears Stable

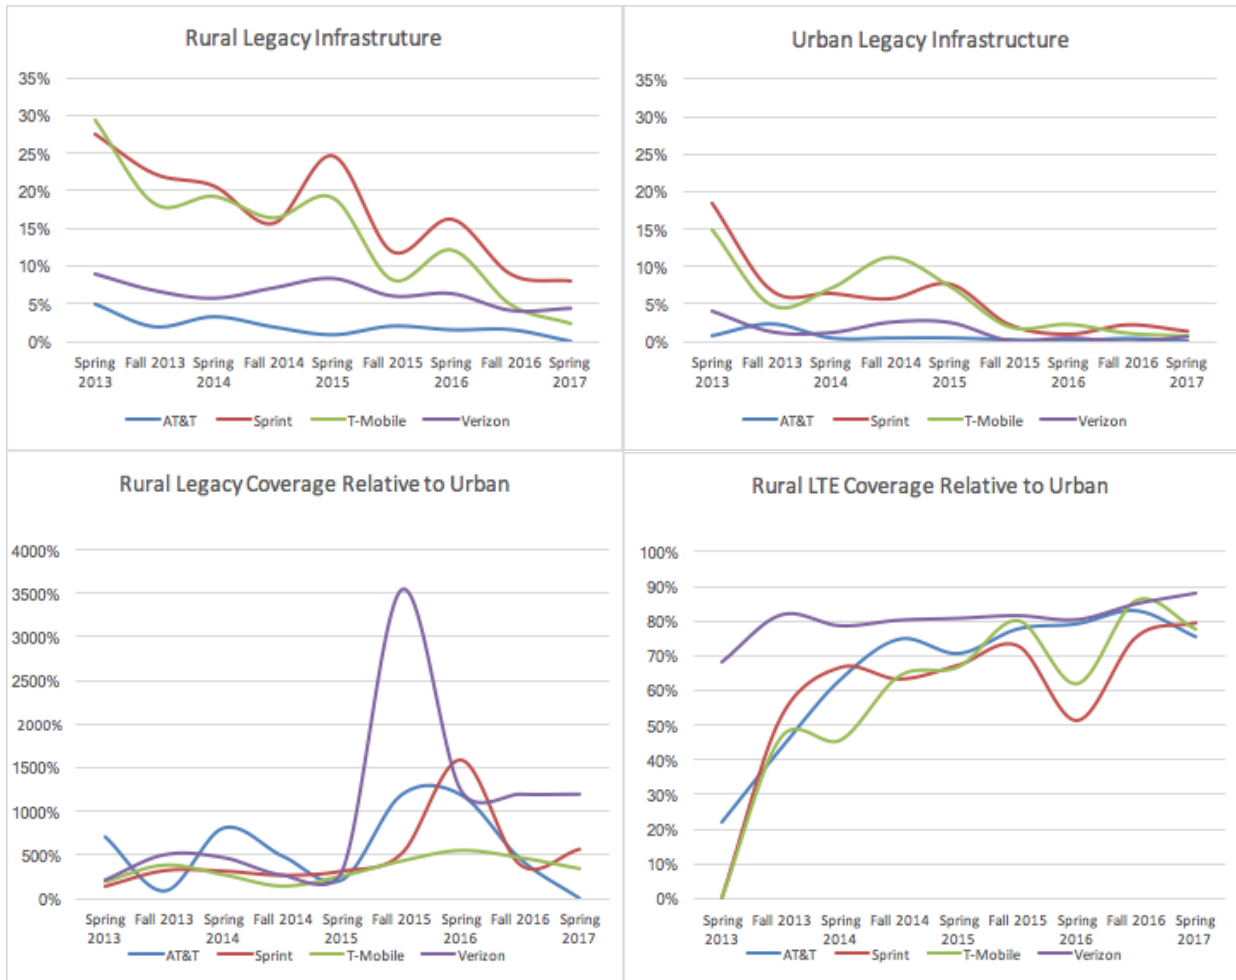
The deployment of LTE as well as the replacement of legacy 1G and 2G wireless access networks appears unchanged from the last four measurement rounds. The data suggests that as of the spring 2017 measurement, deployment of urban LTE is capped at between 90-95% of the measurement locations. Rural LTE deployment is at 70-85% of measured locations. Urban LTE deployment has remained stable since fall 2015.

Legacy 1 and 2G networks appear to have hit a floor of removal at between 1-3% of measurement locations while legacy rural deployments are still between 2 and 9% (Sprint being the hindmost). Notably, for spring 2017, AT&T has registered no legacy 1/2G networks in any demographic.

Rural users will encounter legacy service between 5 and 10 times more often than urban users for Sprint and T-Mobile, while for AT&T the likelihood is near zero. Rural users will experience LTE service only about 80% as often as urban users.

Advanced network services (such as VoIP and streaming video) are really only practical on LTE, and the extent to which LTE is not available, or legacy services are experienced instead, will dramatically degrade mobile broadband for rural and tribal users.






9. Video Quality Worsens

CalSPEED estimates over-the-top Internet video performance - both streaming (e.g. YouTube and Netflix) and interactive (e.g. Skype and FaceTime).

9.1 Video Streaming


Video streaming is the service of viewing a video, in real-time, stored on an Internet server. In order to reduce Internet traffic and improve video performance, most video content providers cache video content as close to the viewing user as possible to minimize latency and increase throughput. CalSPEED approximates this caching by estimating video streaming performance as downstream TCP throughput from the west server. The first analysis of CalSPEED video was for the spring 2015 measurement¹¹ and the analysis in this paper is the fifth update measurement from those conclusions.

¹¹ Biba, Ken, "Wireless Video in California - Spring 2015", Novarum, Inc. January 2016




Spring 2017 streaming video quality has overall worsened for all carriers, with fewer measurement locations reporting high definition (HD) quality and more reporting no service (NS), standard definition (SD), or low definition (LD) quality. HD streaming video coverage for tribal users has decreased for all carriers, but most dramatically for Sprint and T-Mobile.

Rural users appear most impacted with 5-10% increases in “No Service” locations for all carriers.

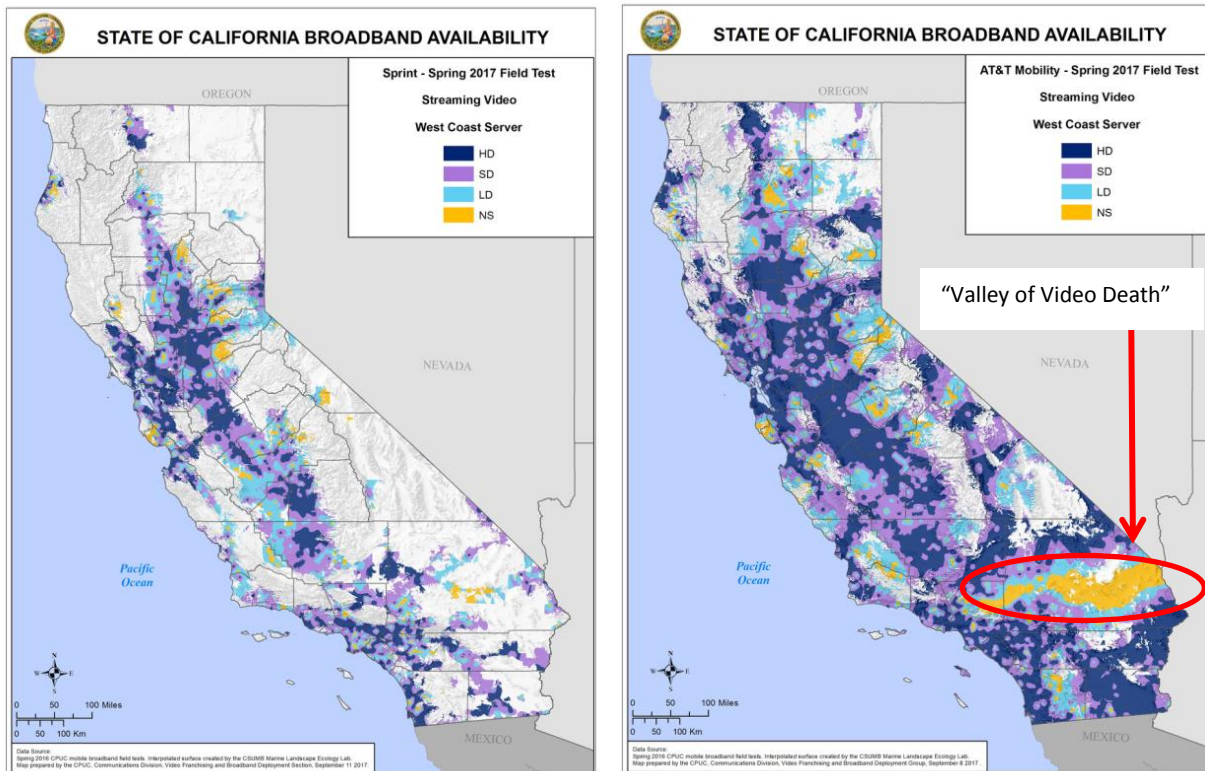


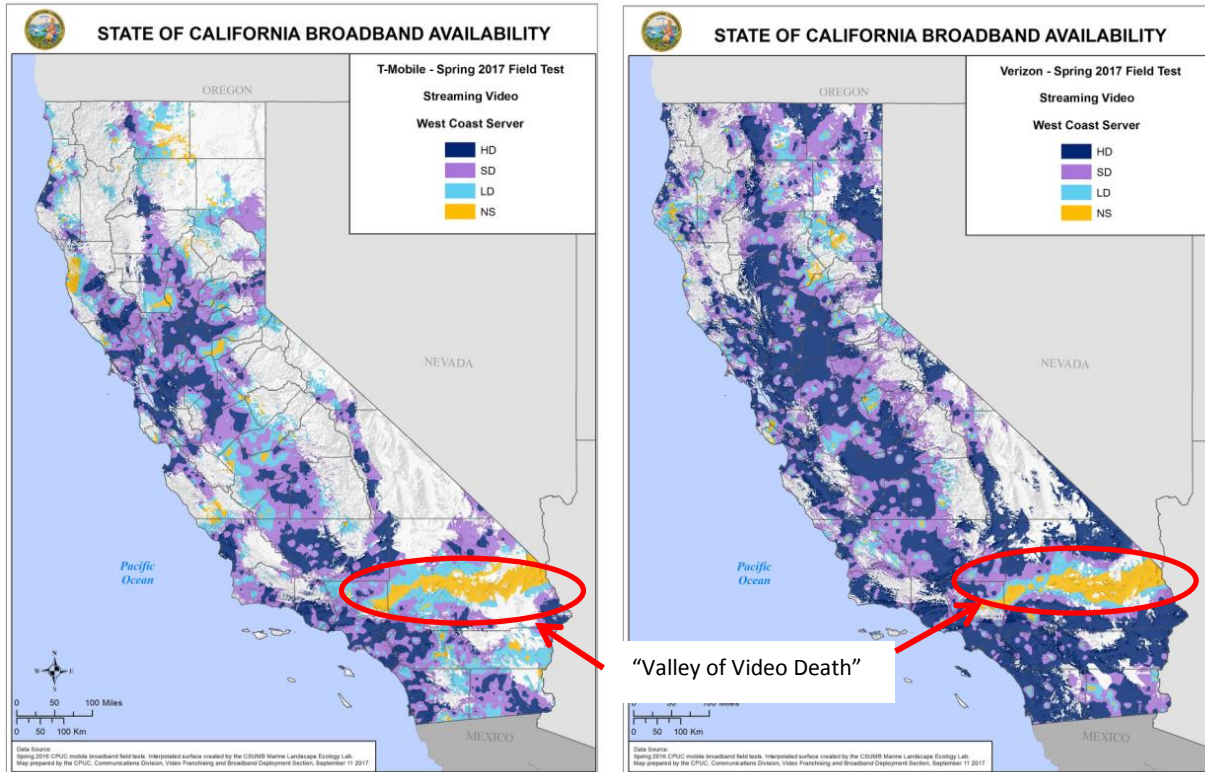
All four carriers continue to offer HD quality video streaming in over 50% of the urban measurement locations but all carriers show material drops in HD video streaming availability for rural and tribal users. Over the last five measurement rounds rural users, for all carriers, have about 70-80% of the HD streaming availability of urban users.



Rural users suffer from no HD streaming coverage ~2-3x more than urban users. Both HD availability and lack of coverage appear to be persistent over all six measurement rounds.

Mobile video streaming coverage is dramatically illustrated in the following coverage maps illustrating estimated mobile video service across the state for the four carriers. HD streaming coverage for all carriers is localized with Verizon giving the best overall state-wide coverage. A major difference is a dramatic band of poor (including No Service) across Southern California - the "Valley of Video Death" (in yellow) - for all four carriers.







9.2 Video Conferencing


Conference video is a two way audio/video stream between two users. CalSPEED uses two-way MOS and streaming video streaming quality to construct a metric for interactive video. Two estimates are made to evaluate interactive conversations throughout the US: a "west" estimate using the west server to emulate one side of the interactive conversation and an "east" estimate using the east server.

The video quality chart on the right illustrates that for all carriers and all geographies, conference video has modestly improved to the east server since spring 2016. Particular improvements have been for Sprint and T-Mobile for rural users.



The chart on the left above, illustrates that conference video degrades in availability when moving from urban to rural users and from west to east conversations (the performance of the Internet backbone degrades the interactive video). All carriers showed improvement in urban interactive video quality (particularly Verizon) at the same time as showing degradation in rural interactive video coverage.

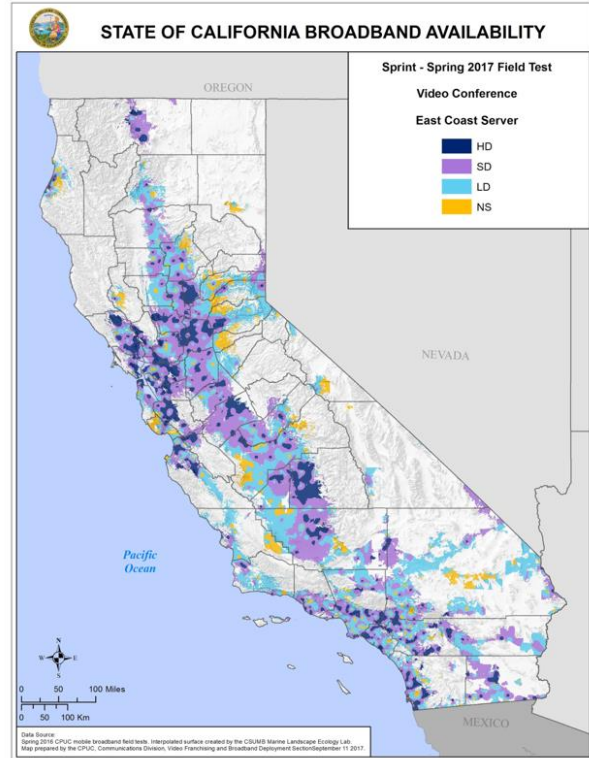
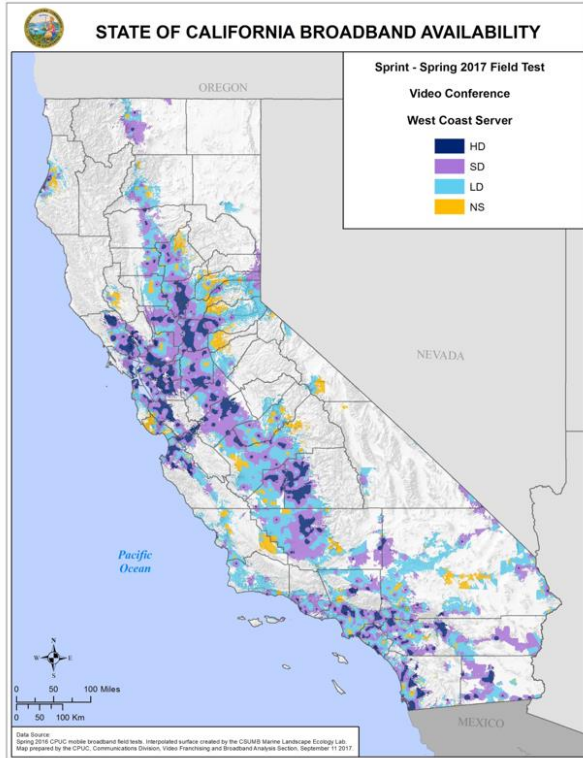
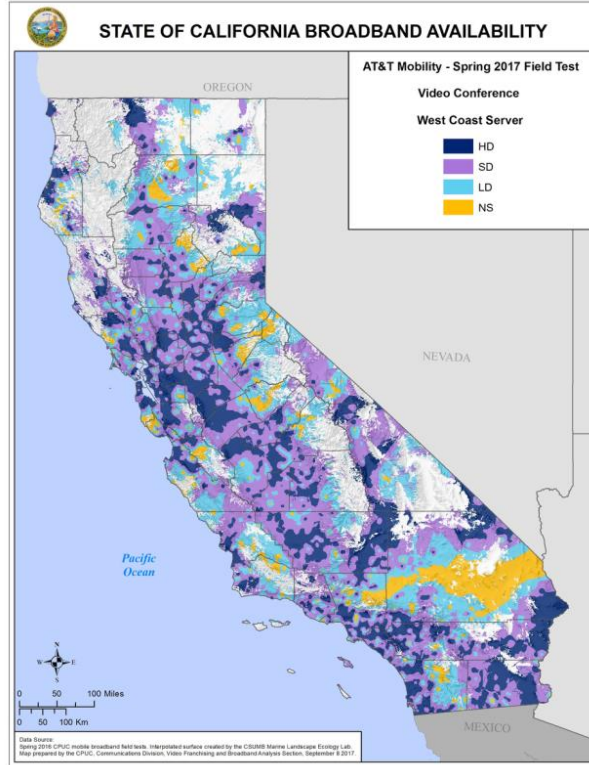
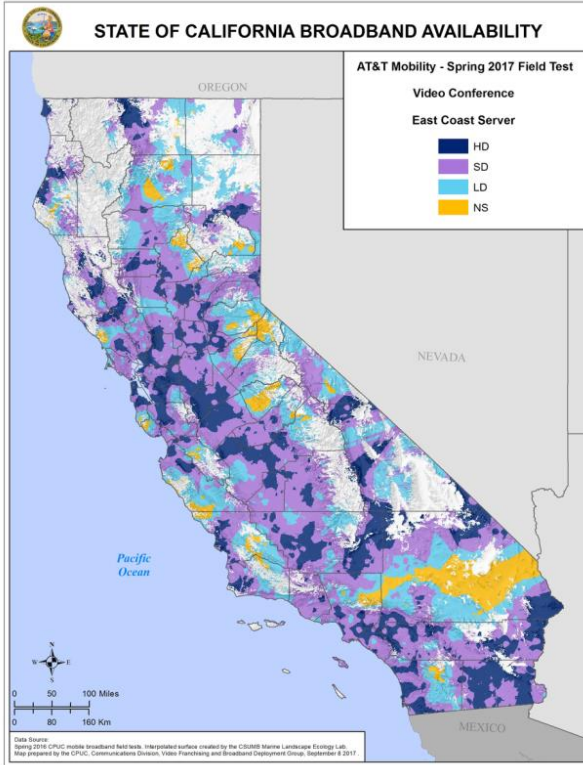
The Internet backbone, thru increased latency, can decrease the availability of HD quality conference video for all carriers. In measurements since fall 2015, HD conference availability is about equal between conference destinations - west or east.

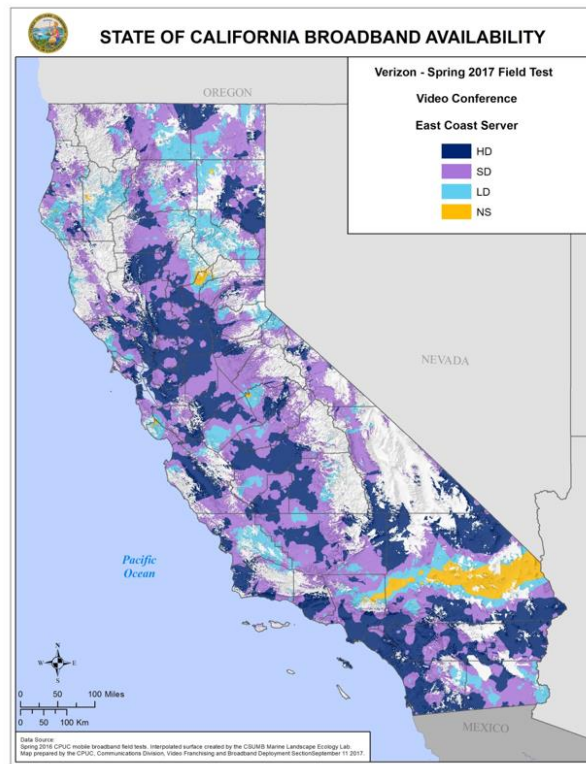
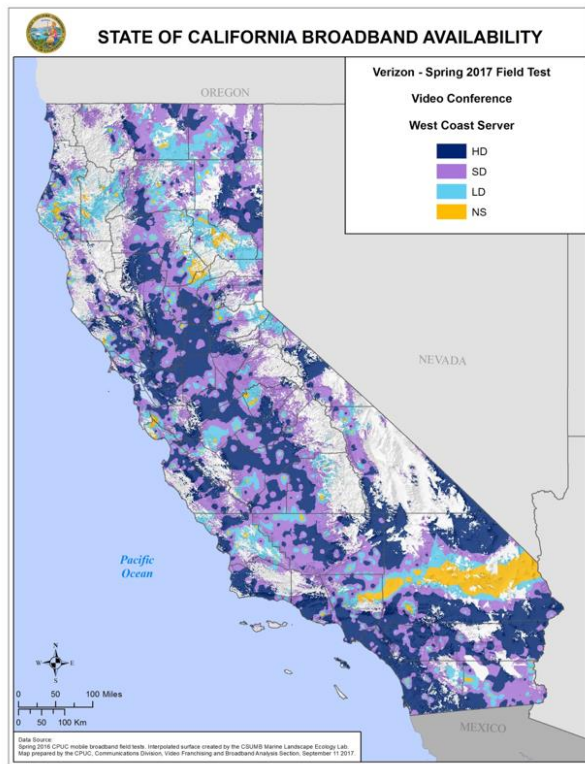
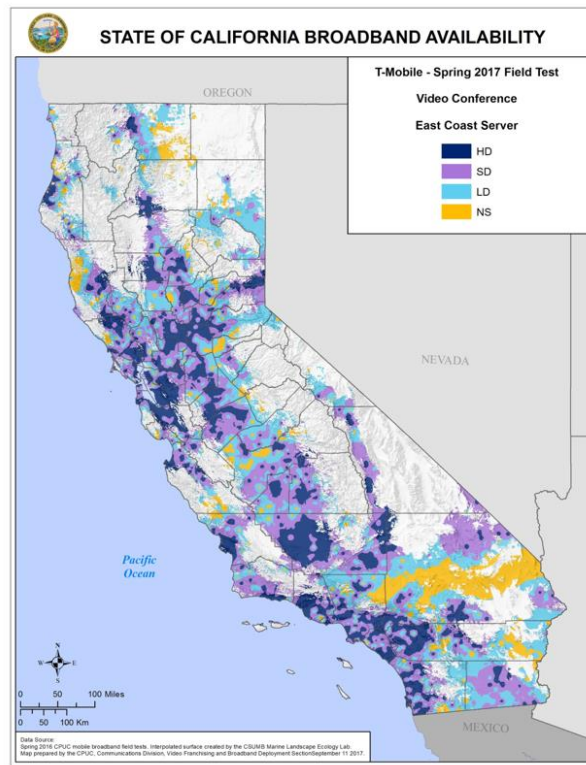
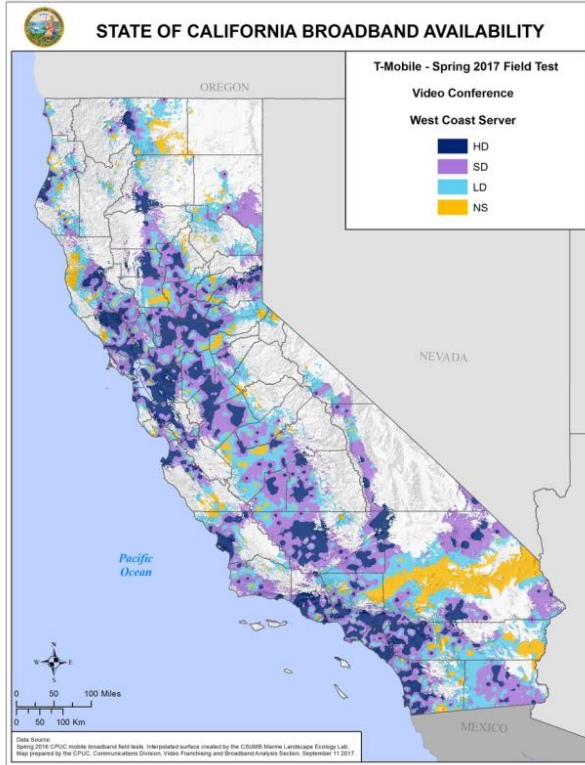


HD west and east conference availability are similar - ~50% of the measured locations (with Sprint around 30%). However (west illustrated above), rural users are materially less likely to have conference service availability. Only about 50-60% of rural users (compared to urban users) have access to HD conferencing and about 2x more rural users than urban users will have no conference ability at all (with the exception of Sprint, in which the lack of HD interactive video service is equally shared between urban and rural)

Interactive video service quality is mapped across the state for all four carriers in the following charts. With the increased east availability beginning in fall 2015, service quality east vs. west for all carriers is similar.


With spring 2017 we see a loss of interactive video service in Southern California - “Valley of Video Death” (the area marked in yellow).






10. Latency Difference Stabilizes, Throughput Difference Continues

CalSPEED measures performance to two geographically distinct servers to estimate the full range of Internet service - both to “local” servers and “distant” servers. Since users will be accessing Internet resources located not just geographically local, but distributed around the U.S. and the world, how each carrier chooses to integrate into 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 case, the physics of data transmission¹² adds about 80 milliseconds of latency to get from one coast to another - in addition to any local wireless access latency. Latency differences over 80 milliseconds suggest suboptimal 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 Kansas, where the Internet distance to either the east coast server or the west coast server is essentially the same. In the past 18 months, all carriers are converging on a close to optimum geographic mean latency penalty of ~80 msec.



TCP throughput is related to latency - the longer the latency, the lower the throughput¹³. 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 demonstrates this fact.

In recent years (fall 2014 thru spring 2015), we noticed a substantial decrease in the TCP throughput difference, and in the spring 2016 a dramatic increase in the TCP difference for all carriers. This dramatic decrease in TCP throughput to the east server has continued into spring 2017.

¹² The 80 millisecond calculation takes into account the speed of light

¹³ This is a consequence of TCP’s data reliability and congestion control mechanisms.

11. Relationships between Signal Strength, Signal/Noise and Throughput


There is much confusion about the relationship between “bars” of signal on a phone or tablet, and the quality of the service. Even today, a phone can be set to show RSSI (Received Signal Strength Indicator) rather than bars but there remains confusion on whether more RSSI or more bars, yields improved service.



Beginning with spring 2017, CalSPEED was updated to capture both RSSI and the effective signal quality (SNR - signal to noise ratio) during each measurement. Different mobile technologies represent these values differently with, unfortunately, no common standard - but the preponderance of LTE service in the CalSPEED dataset lets us analyze RSSI and SNR for LTE networks.

The following scatter plots show the relationship between both signal strength and SNR for both urban and rural demographics for each carrier. There is a shared set of conclusions for all carriers.





- There is a similar pattern for all carriers.
- There is no substantial difference between rural and urban.
- There is a weak correlation between signal strength (and perhaps bars of service) and TCP throughput.
- There is a moderate positive correlation between SNR and TCP throughput. Higher SNR, on average, generally results in higher TCP throughput.
- However, there is poor predictive relationship between either RSSI or SNR and TCP throughput. That is, a given RSSI or SNR can correspond with a wide range of values for throughput (from zero to 10s of Mb/s). Pick an SNR value and note the large range of possible TCP throughputs that were observed in the spring 2017 data - for all carriers.

AT&T

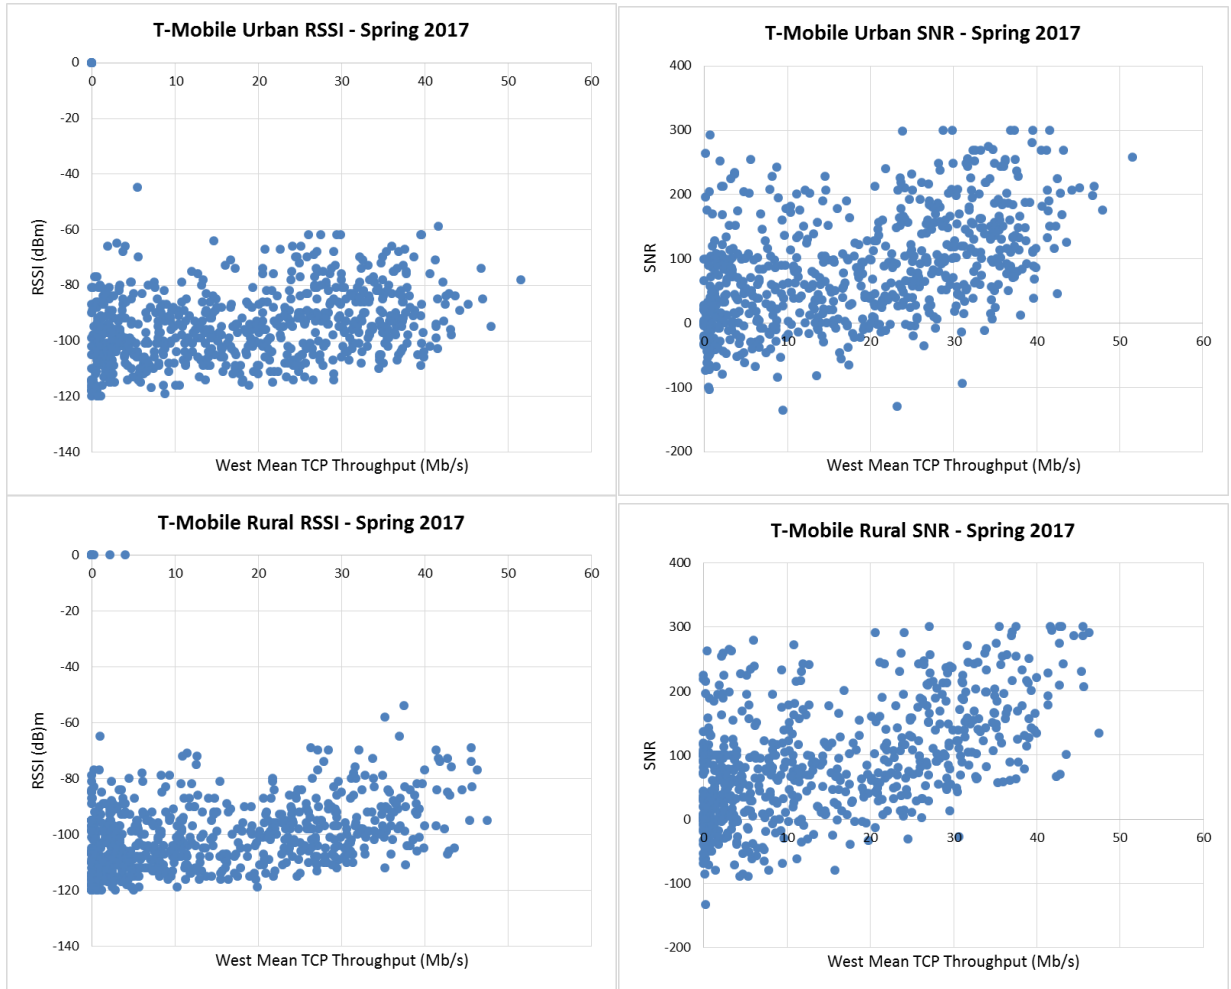





Sprint




T-Mobile




Verizon





Looking at the correlation coefficients for each carrier, urban and rural, for RSSI and SNR to west TCP throughput, a weak relationship of RSSI to throughput is apparent, and a much stronger relationship between SNR and TCP throughput appears.



While there is moderate correlation between SNR and throughput, the scatter plots demonstrate the weak predictive relationship between SNR and throughput. Better RSSI, and particularly lower SNR, is generally good ... but does not guarantee higher throughput.

12. Conclusions

This is an analysis of the measurement Round 11, spring 2017 dataset for CalSPEED. The key results for California mobile broadband include the following.

Persistent mobile digital divide	The rural/urban mobile digital divide appears to be persistent with rural users having material (about 1/3) poorer mobile broadband service over a period of years with little indication of improvement. This is true for throughput, latency, packet loss, MOS, jitter, broadband coverage and deployed technology. Spring 2017 shows additional incremental degradation.
Mean throughput worsens	Mean throughput worsens for all carriers (small increase for Sprint downstream and for Verizon upstream) and returns to levels last seen at least a year ago.
Carrier throughput throttling continues	Throughput throttling for all carriers continues from fall 2016 into spring 2017. The downstream throttling threshold continues from fall 2016 at ~42 Mb/s. The upstream throttling threshold continues from fall 2016 at ~20 Mb/s (15 Mb/s for Sprint).
Mobile broadband coverage decreases	Broadband coverage at 25 Mb/s down, 3 Mb/s continues below year ago levels at ~13%. Using the benchmark of 10 Mb/s down and 1 Mb/s up, improves “coverage” to ~50% for all carriers.
Underlying service quality mixed	Mean rural latency worsens, jitter worsens, packet error rate improves, and TCP connection reliability worsens for all carriers.
Rural and Tribal TCP connection failure rates 3-7x urban	TCP connection attempts fail much more often for rural users than for urban users for all carriers. An urban AT&T or Verizon user can expect 2% of TCP connection attempts to fail (often invisibly retried by the using applications) while ~20% or rural users can expect TCP connection failures.
Over the Top VoIP quality modestly improved, continued relative impairment for rural and tribal	VoIP quality continued improved after substantial degradation in spring 2016. Rural and tribal VoIP quality remains degraded.
LTE deployment coverage has peaked	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. Notably, there was no legacy 1/2G service detected for AT&T in California.

Internet latency stabilizes	The latency difference between east and west servers has stabilized at ~80 msec.
Internet video and conferencing results mixed	Internet OTT streaming video and conferencing availability modestly improved from fall 2016. A Valley of Video Death appears as a consistent area of No Service for OTT video in rural Southern California for all carriers.
East/West throughput differences continue at all time high	Ratio of east/west throughput continues pattern of last year with an all-time high of 60% for all carriers but Sprint at 40%.
Little difference between devices	There continues to be no material throughput or latency difference between current smartphones and tablets.
Weak relationship between signal strength and throughput	There is a modest correlation between Signal to Noise Ratio (SNR) and west throughput, but not for Received Signal Strength Indicator (RSSI). Moreover, the wide range in throughput values for any particular SNR or RSSI means both are poor predictors of throughput.

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

How CalSPEED Measures Broadband

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 smartphone and current tablet. Both are upgraded for each measurement round to match the latest wireless technology deployed by each carrier.

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 - iPerf¹⁴. 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 complete network path, from the client device, through the local access network, through the Internet backbone, to several ultimate server destinations.

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

While it is impossible to measure all Internet servers, 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 up to 2:1 differences in end to end latency and jitter and material difference in upstream and downstream throughput between the two servers.

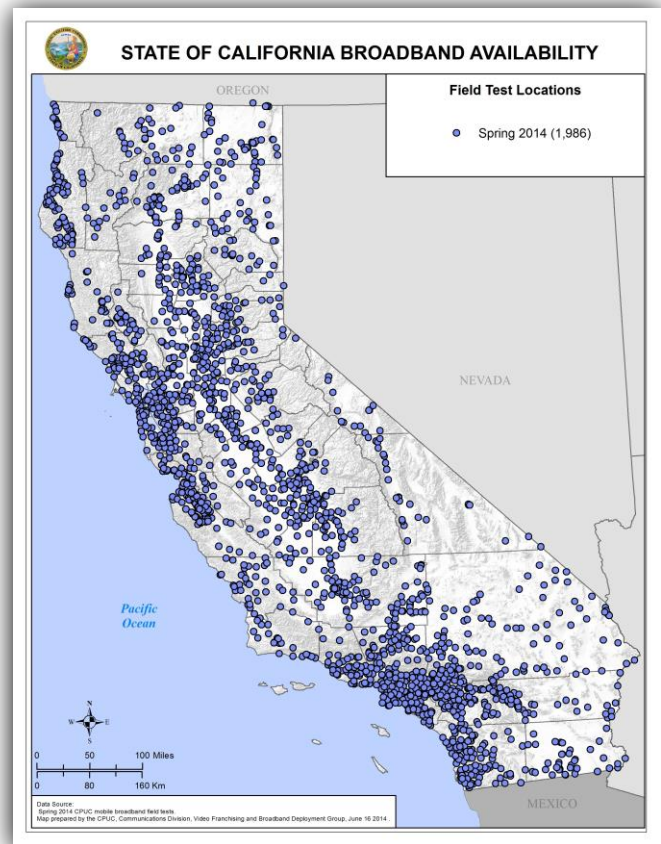
These differences in fundamental network performance illustrate that location matters - Internet performance delivered to the user - the Internet user experience - will vary based on where on the Internet the desired server is located. And desired servers are scattered across the Internet, not just close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

CalSPEED has two complementary methods of testing - the first is a structured sampling program of 1,986¹⁵ 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 tablet for each of the four major carriers. The use of multiple contemporary user devices gives a good snapshot of the best user experience.



¹⁵ Originally 1200, but later increased to improve predictive precision of the interpolation models.

The second method is the independent use of CalSPEED to provide crowd-sourced 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 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.

Not Just Data but Voice and Video. CalSPEED measures not only the underlying basic Internet data transmission of datagrams and TCP connections, but also interactive voice (the Internet's replacement for POTS), streaming video and interactive video (video conferencing).

CalSPEED constructs an over-the-top interactive voice model, using the LTE voice digitization method, that gives an estimate of the Mean Opinion Score (MOS) of the voice service.

CalSPEED uses a derivative of the Google's video quality metric¹⁶ to construct a metric of Internet video quality. CalSPEED measures both downstream streaming video (such as YouTube or Netflix) as well as interactive video (such as Skype or FaceTime). Streaming video is measured using downstream performance from CalSPEED's west server - assuming that most such video is cached closer to the user. Interactive video is measured both to the west and east servers (to assess the effect of the Internet backbone) and uses both upstream and downstream performance measures.

Maps for decision-makers not just for information. We then take the measurement data and create geospatial kriging¹⁷ 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 geo-statistical 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.

Massive Dataset. CalSPEED has now had nine rounds of sampling California (spring 2012, fall 2012, spring 2013, fall 2013, spring 2014, fall of 2014, spring 2015, fall 2015, and spring 2016) and is shortly to finish a tenth round (fall 2016). 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.

¹⁶ <https://www.google.com/get/videoqualityreport/#methodology>

¹⁷ <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 geo-statistical 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