CalSPEED: California Mobile Broadband - An Assessment - Fall 2016

Persistent Urban/Rural Digital Divide and Carrier Performance Throttling Returns

Ken Biba
Managing Director and CTO
Novarum, Inc.

This is an analysis of the Round 10, Fall 2016 dataset for CalSPEED. It shows substantial changes happening in the wireless broadband networks in California for all carriers.

Persistent mobile digital divide	The rural/urban mobile digital divide appears to be persistent with rural users having material (about 1/3), often dramatic, poorer mobile broadband service over a period of years with little indication of improvement.
Mean throughput recovers	Mean throughput recovers for all carriers from decreases noted in Spring 2016 and trends begin to demonstrate differences between carriers in matching network capacity to user demand.
Carrier throughput throttling returns	Throughput throttling for all carriers returns after temporarily disappearing in the Spring 2016 measurement. 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 upstream throttling threshold appears to have increased to 20 Mb/s (15 Mb/s for Sprint) from 15 Mb/s in Fall 2015.
Mobile broadband coverage decreases	Material decrease in broadband coverage at the 25 Mb/s down, 3 Mb/s up for all carriers and all demographics.
Underlying service quality improves	Mean latency, jitter and TCP reliability modestly improved for all carriers.

Over the Top VoIP quality remains degraded, particularly for rural users	VoIP quality improved after substantial degradation in Spring 2016. Rural 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.
Internet latency stabilizes, throughput penalty continues substantial	The latency difference between East and West servers is has stabilized at about 80 msec, while throughput differences have dramatically grown to 60% for all carriers.
Internet video and conferencing recovers	After decreases in both streaming and conferencing coverage in Spring 2016, coverage increased for both services in Fall 2016.

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 ten rounds of measurement over the entire state collecting over 20,000,000 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 ten rounds of measurement. Previous papers have analyzed the prior rounds of measurement[1][2][3][4][5]. The methodology has been rigorously analyzed with respect to other available mobile measurement tools[6]. This paper examines the incremental changes from the previous report extending thru the Fall of 2016.

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

[1] Ken Biba, "Assessment of California Mobile Broadband Spring 2014", Novarum, September 2014.

[2] Ken Biba, "Assessment of California Mobile Broadband Fall 2014", Novarum, June 2015.

[3] Ken Biba, "Assessment of California Mobile Broadband Spring 2015", Novarum

[4] Ken Biba, "Assessment of California Mobile Broadband Fall 2015", Novarum

[5] Ken Biba, "Assessment of California Mobile Broadband Spring 2016", Novarum

[6] Ken Biba, "Comparison of CalSPEED, FCC and Ookla", Novarum, Inc., September 2014.

2. Persistent Mobile Digital Divide

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

CalSPEED metrics document this persistent digital divide.

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

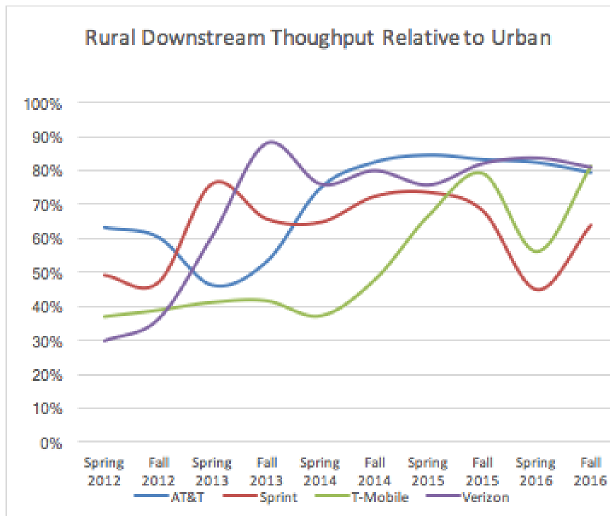

Rural mean upstream TCP throughput is consistently only 60-70% of urban.

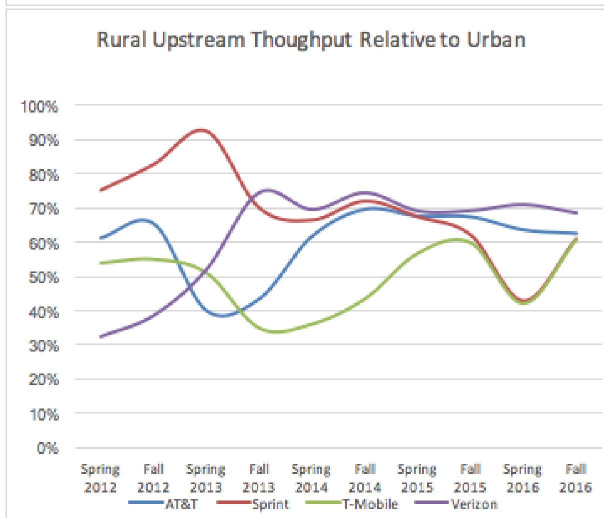

Rural mean latency is consistently at least 30% worse than urban.

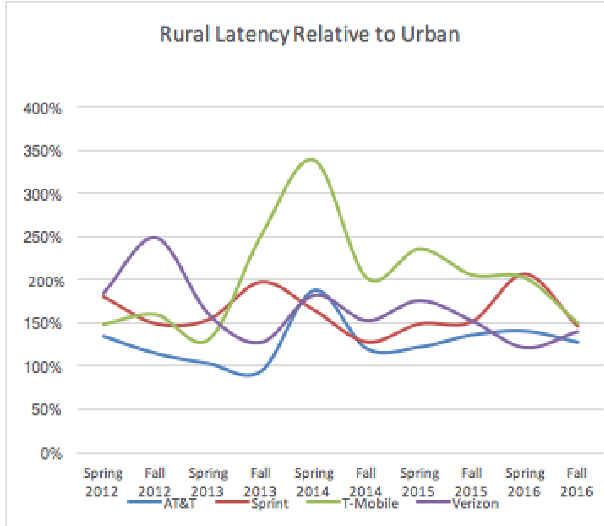

Rural mean jitter is consistently at least 20% worse than urban.

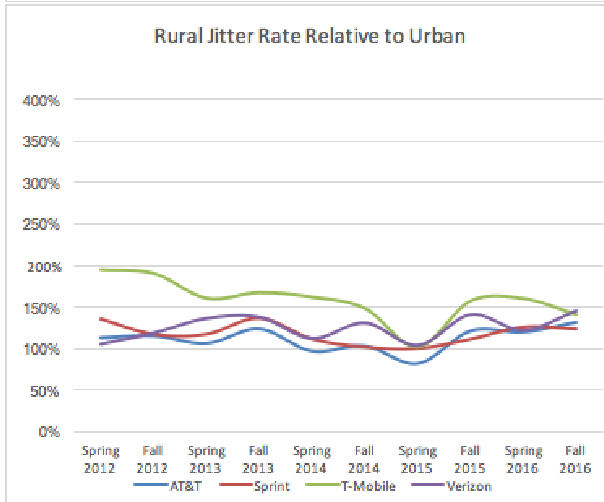

Rural packet loss is consistently 20-200% worse than urban. The poor performance of rural latency, jitter and packet loss strongly implies that real-time streaming voice and video services will perform materially worse in rural than urban areas.

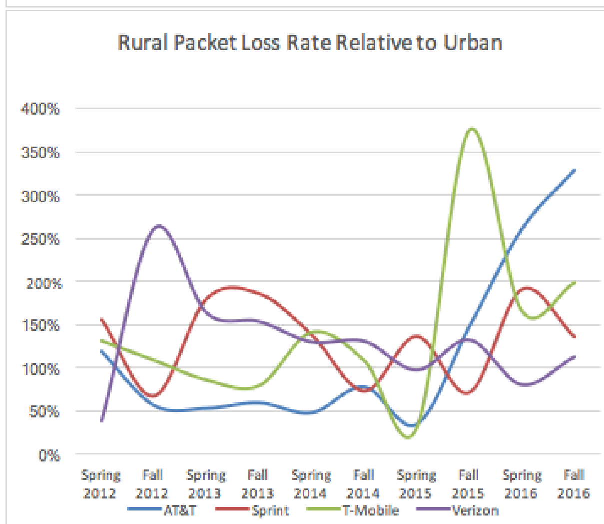

Rural TCP connection failure rate is consistently at least 5x worse than urban with indications that this disparity is getting worse. For rural and tribal users of all carriers, about 1 in 5 TCP connection attempts fail. Particularly disturbing is a trend for both AT&T and Verizon to do worse over time.

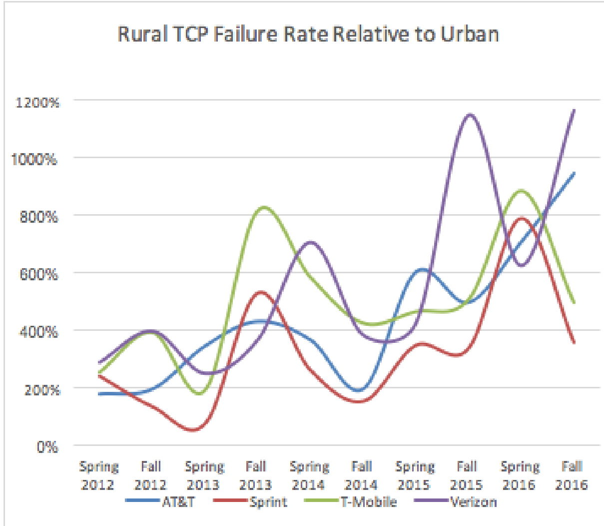

Rural OTT digital voice coverage is consistently only 85% of urban.

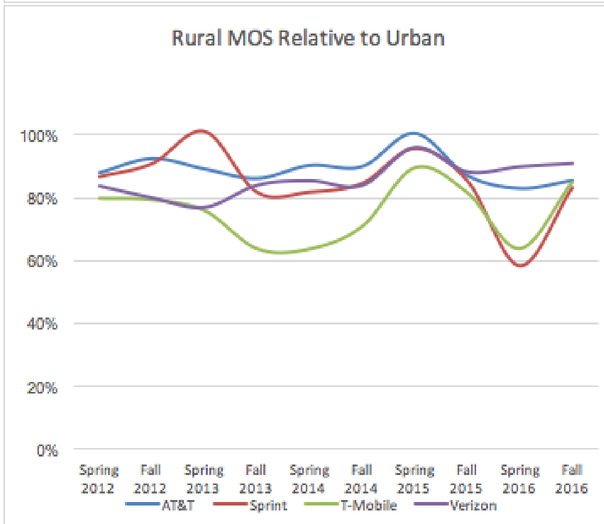

Rural streaming video coverage is consistently 60-80% of urban. Rural users will encounter no streaming availability 3x more frequently than urban users.

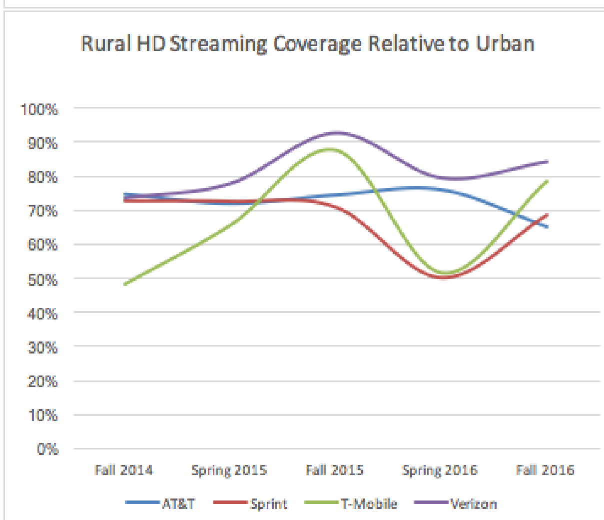

Rural conference video coverage is consistently 60-80% of urban. Rural users will encounter no conference availability 2x more frequently than urban users.

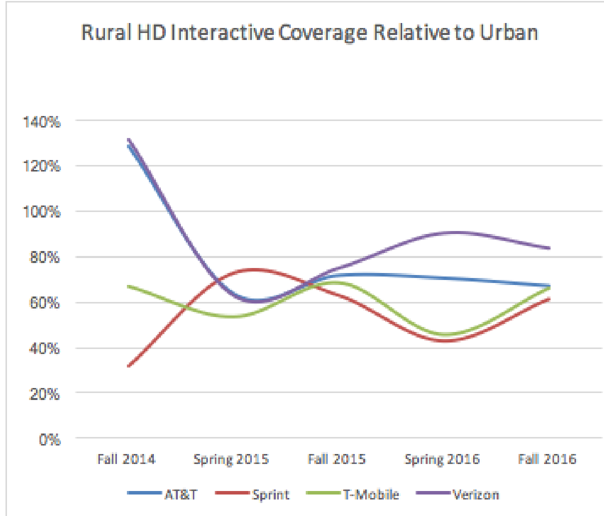

Rural mobile broadband coverage at the standard of 25 Mb/s down and 3 Mb/s up is consistently 60-80% of urban.

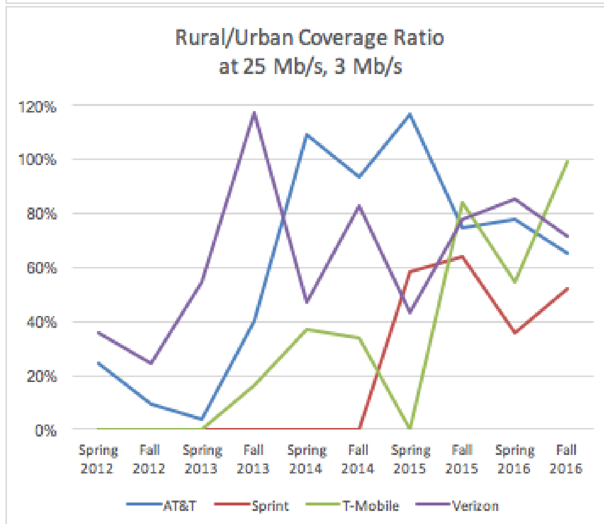

Rural LTE coverage is consistently 80% of urban. The trend suggests rural LTE availability will cap out below 85% of urban. This has substantial implications for the use of mobile broadband as a replacement for wired broadband in rural areas and for the capability of public services.

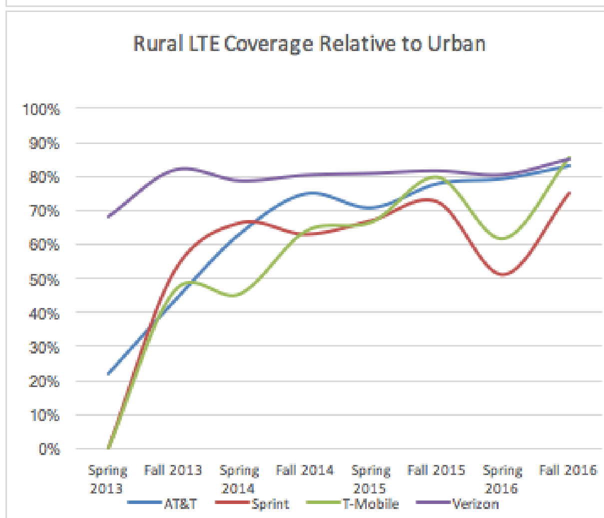

Rural mobile broadband consistently has more than 500% use of legacy 1G and 2G mobile broadband technology than urban. It is much more likely for a rural user to encounter very poor legacy mobile broadband service than an urban user.

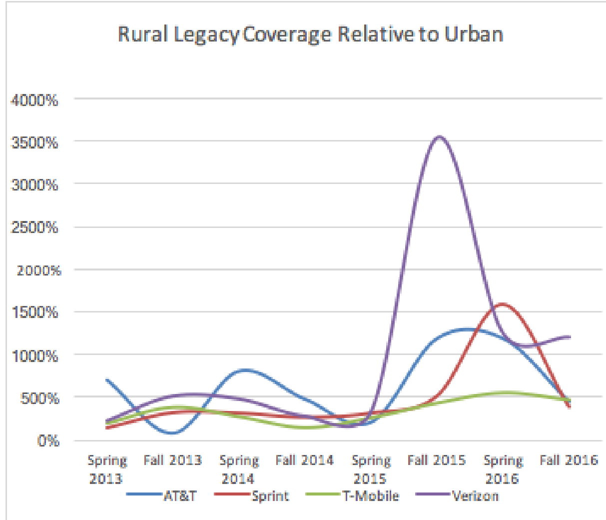

3. Mean Throughput Recovers

All mobile carriers show variation over time in 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. However, some variations in performance can be interpreted as correlated to changes (and mismatches) in network capacity and user demand.

Mean downstream throughput increased in Fall 2016 for all carriers. For Sprint and T-Mobile, it was a recovery from a marked decrease in the Spring 2016 measurement. Verizon and AT&T had a substantial downstream throughput decrease in Spring 2015 that they recovered from by Fall 2015. One speculation as to the cause of these decreases, followed by increases, is a mismatch between customer demand and network capacity as new technology and cell sites are deployed in response to increasing customer use of capacity.

Mean upstream throughput increased for all carriers in Fall 2016 after a marked decrease in Spring 2016 for all carriers.

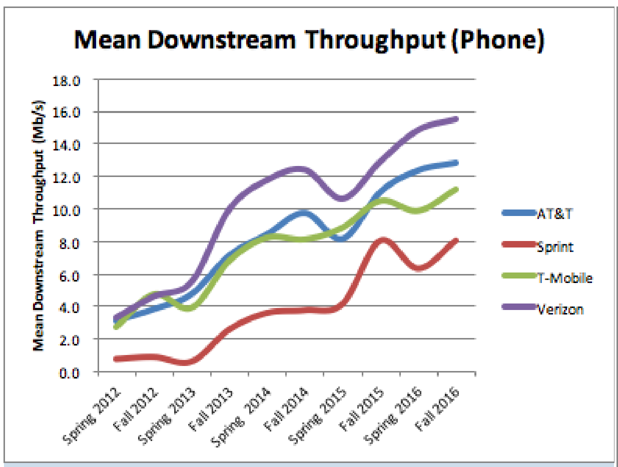

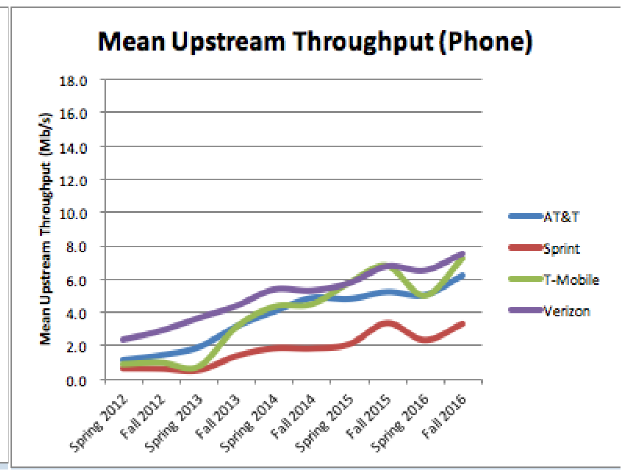

This trend of throughput variation varies between carriers and demographics. Both AT&T and Verizon had substantial (~20%) mean throughput decreases in the Fall of 2015 with recovery in the Spring of 2016 that each exhibited across urban, rural and tribal demographics. Sprint and T-Mobile exhibited substantial (~20-50%) decreases in mean downstream throughput in rural and tribal demographics in Spring 2016 that were recovered by Fall 2016.

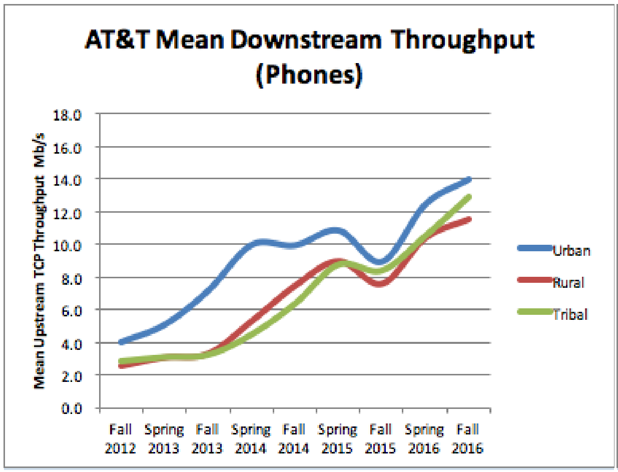

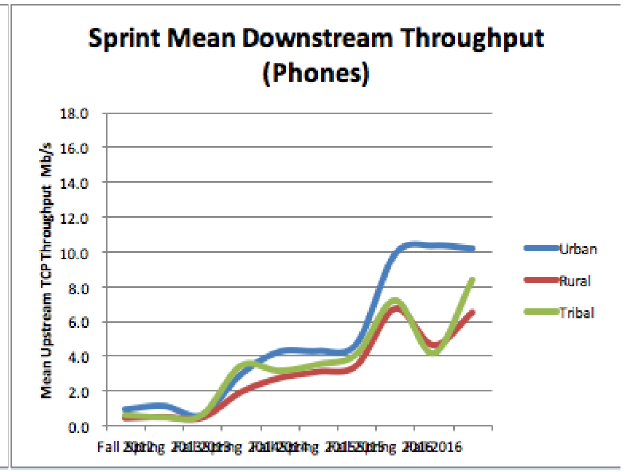

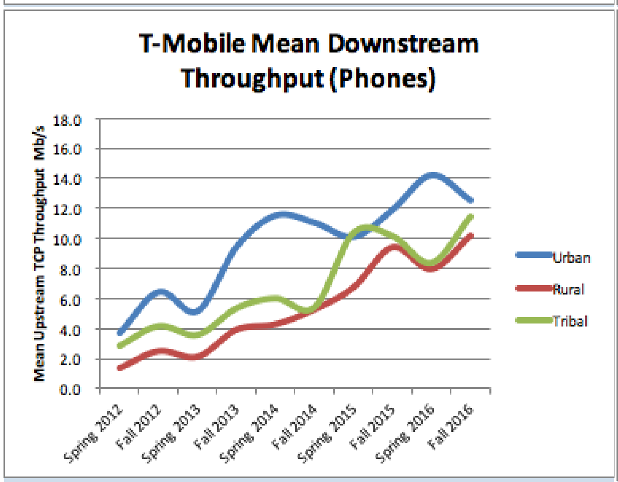

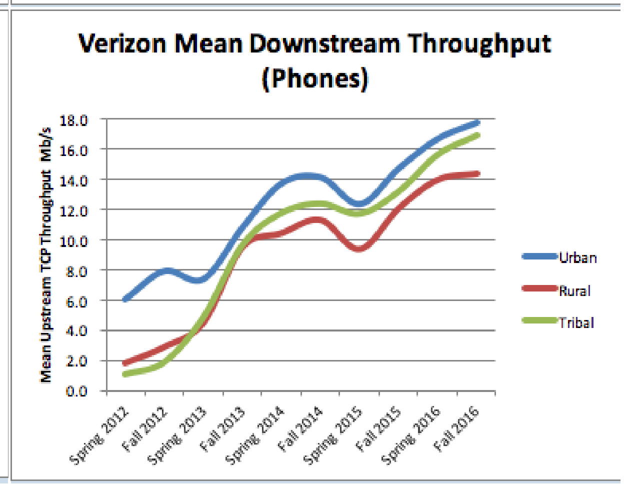

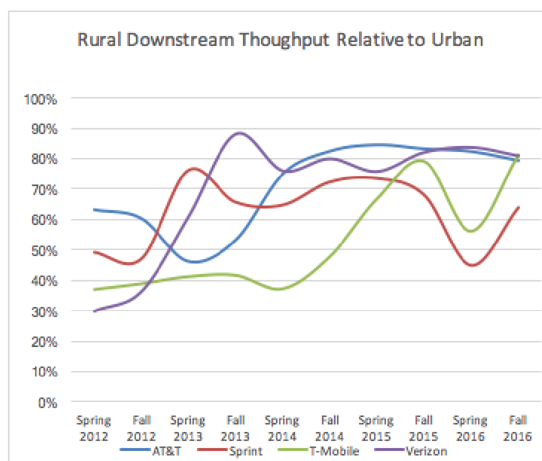

Comparing rural and tribal throughput to urban throughput demonstrates a persistent degradation compared to urban. The chart at the right demonstrates that mean rural downstream throughput, for all carriers and for all measurement rounds, is no more than 80% of mean urban throughput and often less. In the Spring of 2016, rural users of Sprint and T-Mobile had only about half the downstream throughput of urban users of these carriers.

All carriers showed a marked (~10-30%) decrease in mean upstream throughput in the Spring 2016 measurement. This decrease was consistent, for all carriers, across all three demographics (urban, rural, and tribal). This decrease was recovered in the Fall 2016 measurement for all carriers and for all demographics.

Using CalSPEED, we created maps showing the estimated mean downstream and upstream throughput across the state. The coverage maps are clipped to the carrier supplied service footprint.

First, estimated mean downstream throughput. AT&T and Verizon clearly have a greater coverage footprint in California compared to Sprint and T-Mobile. Their rural coverage is largely confined to major roads.

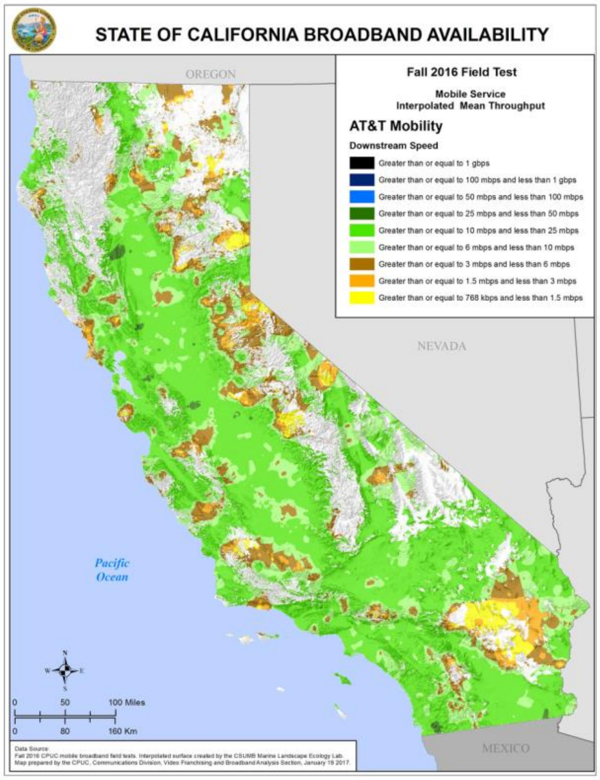

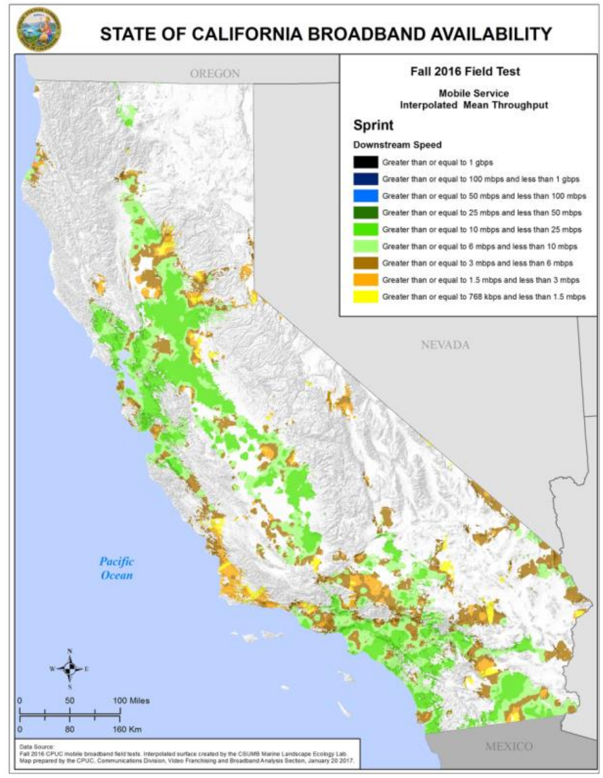

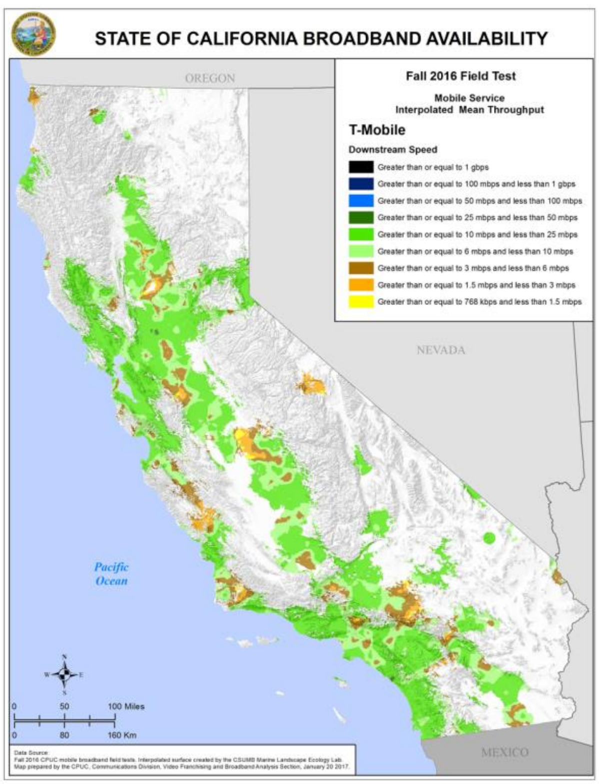

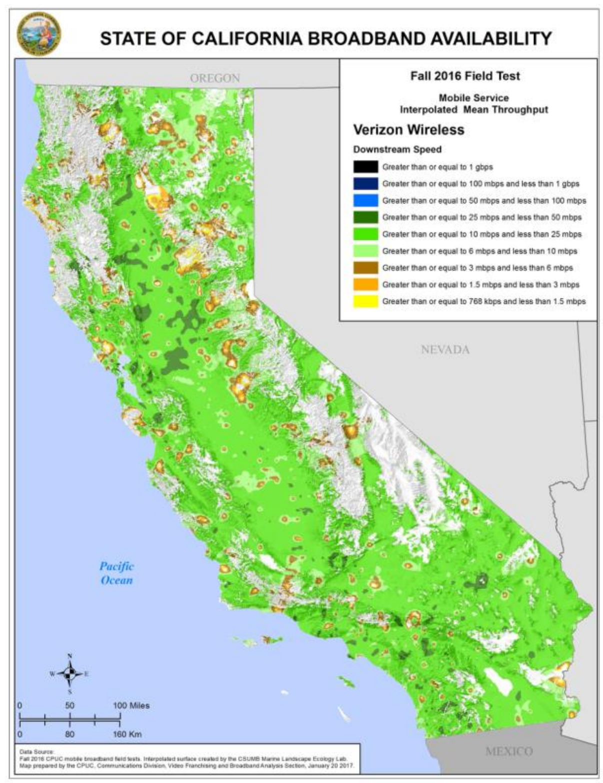

And then the estimated mean upstream throughput.

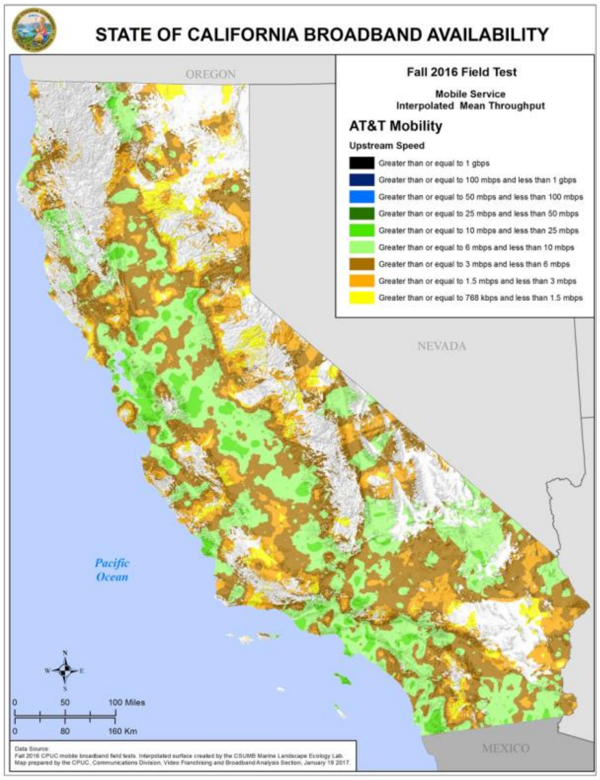

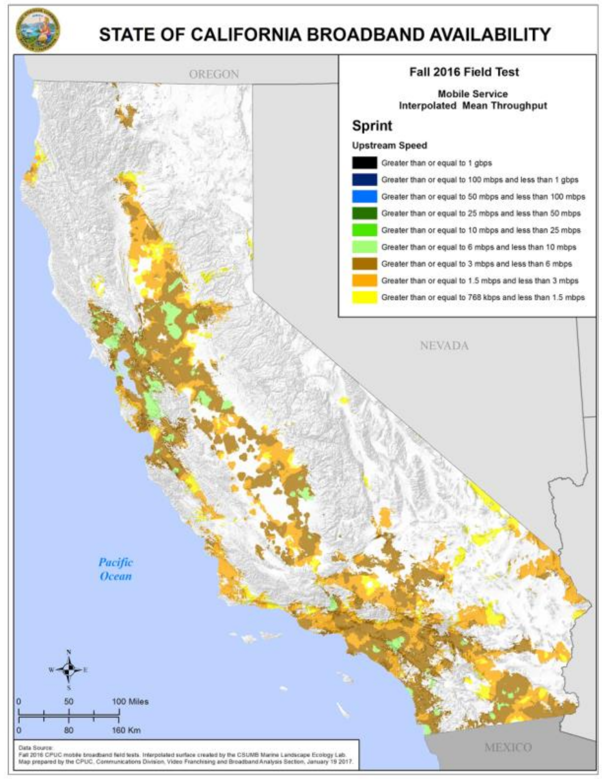

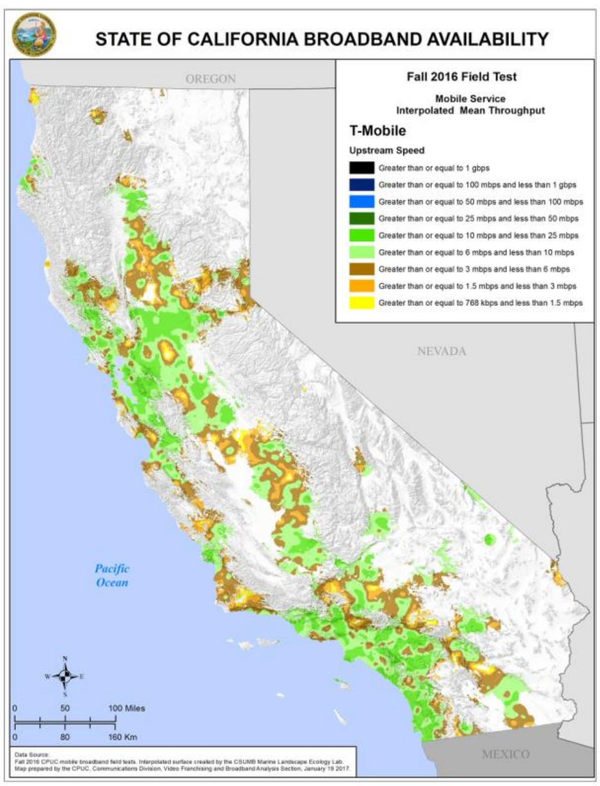

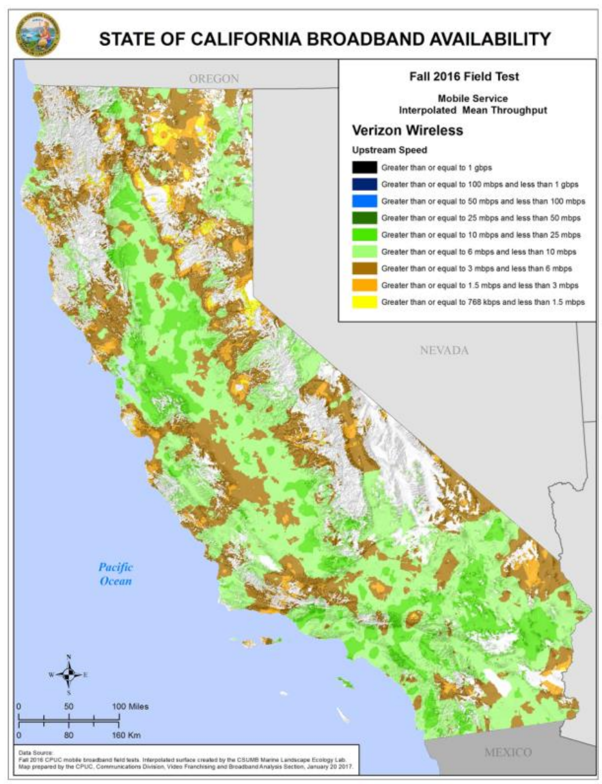

4. Performance Throttling Returns

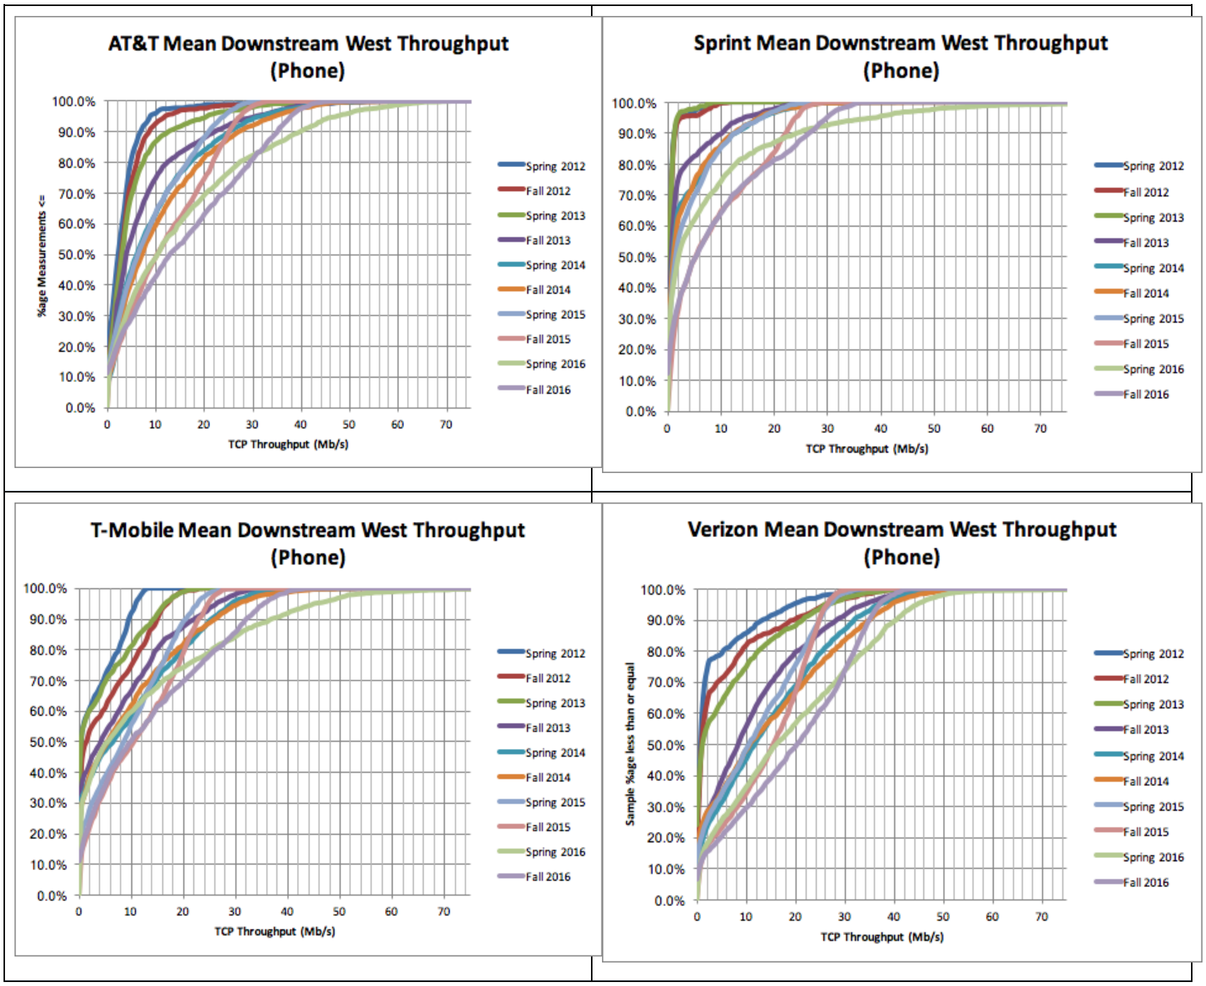

Spring and Fall 2015 measurements noted clear downstream (and upstream) throughout throttling for all carriers as demonstrated in the following cumulative histograms of downstream throughput for phones. We can see a clear change in the distribution for Spring 2015. In prior measurement rounds we can 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 – in 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 was 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, has returned in the Fall 2016. 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 upstream throttling threshold seems have increased to 20 Mb/s (15 Mb/s for Sprint) from 15 Mb/s (12.5 Mb/s for Sprint) in Fall 2015.

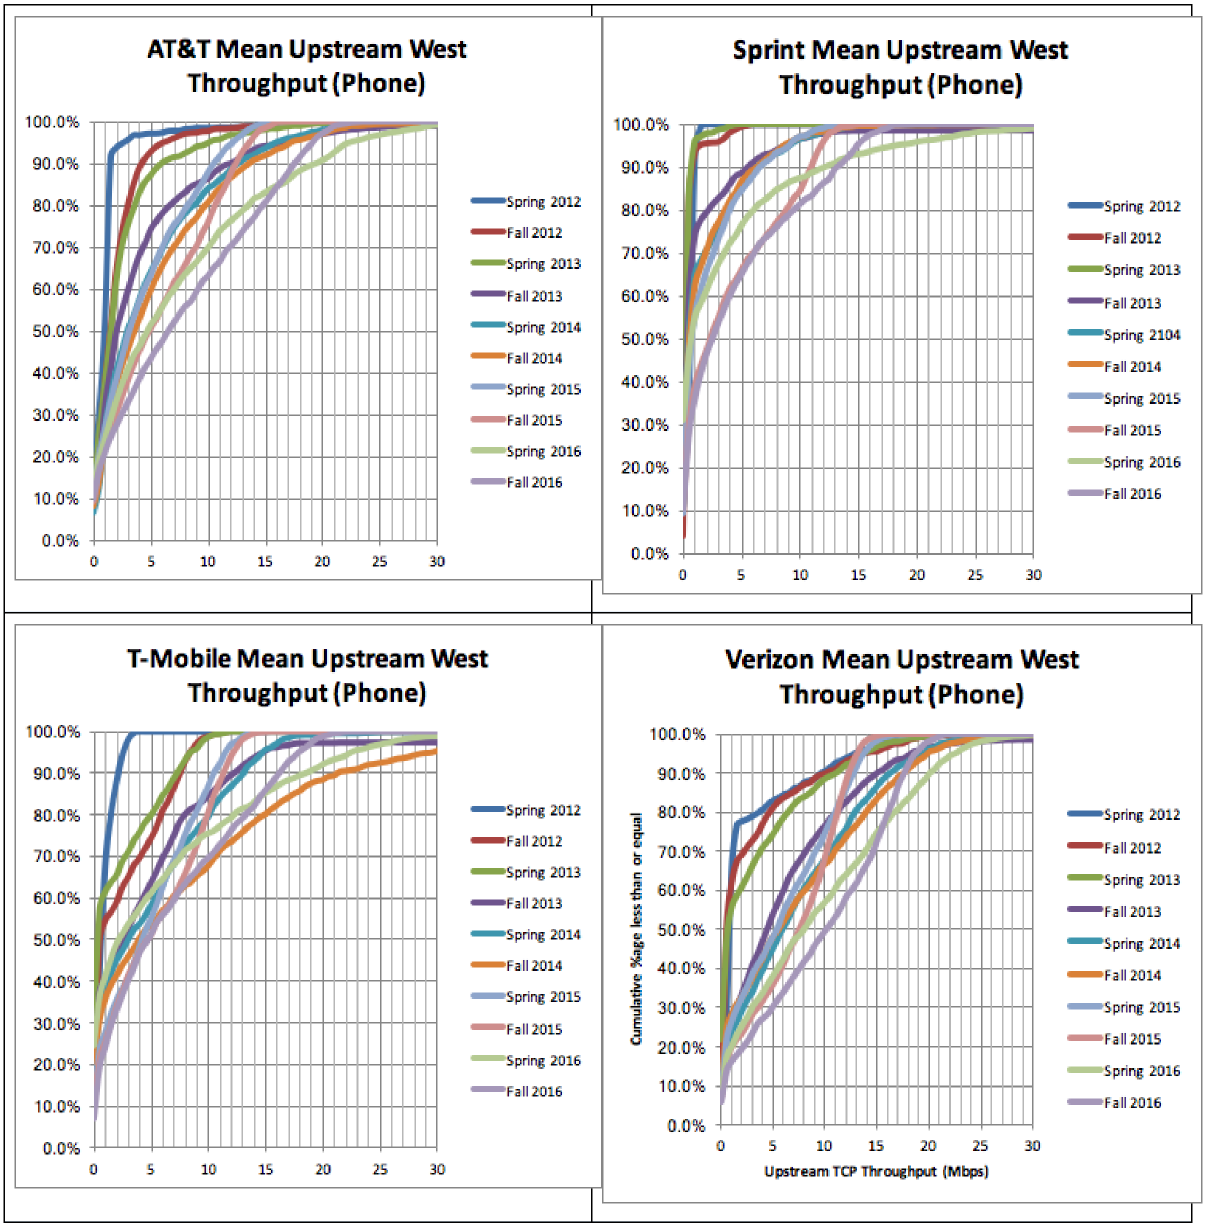
AT&T Mean Upstream West Throughput (Phone)
Sprint Mean Upstream West Throughput (Phone)
T-Mobile Mean Upstream West Throughput (Phone)
Verizon Mean Upstream West Throughput (Phone)
Cumulative %age less than or equal
Upstream TCP Throughput (Mbps)
Spring 2012
Fall 2012
Spring 2013
Fall 2013
Spring 2014
Spring 2104
Fall 2014
Spring 2015
Fall 2015
Spring 2016
Fall 2016

5. Broadband Coverage Decreases

The variations in throughput over time have consequences with the percentage of test locations meeting the 25 Mb/s down/3 Mb/s up broadband standard. 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 standard. For the Spring 2016 measurement round a similar dramatic <u>increase</u> in the number of locations meeting standard is seen. The Fall 2016 measurement sees a material <u>decrease</u> in the number of locations meeting the standard.

This decrease is observed for all carriers and for all demographics. Rural and tribal users saw a bigger decrease than urban users.

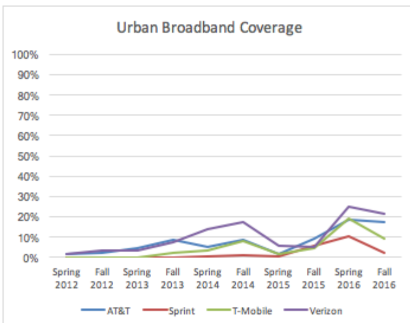

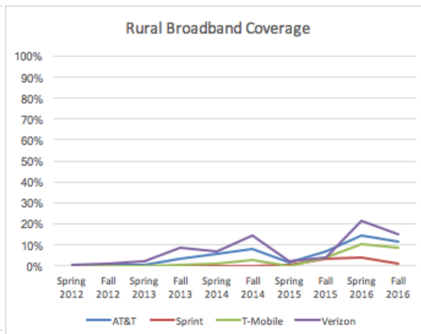

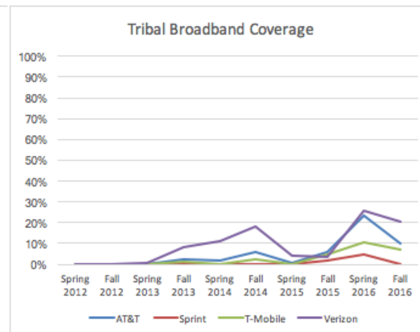

6. Many Measures of Service Quality Improve

Many, but not all, measures of service quality improved in the Fall 2016 measurements.

Mean latency materially improved for all carriers after degrading in the Spring 2016 measurement round. Latency modestly degraded for T-Mobile's tribal users and for Verizon's rural users.

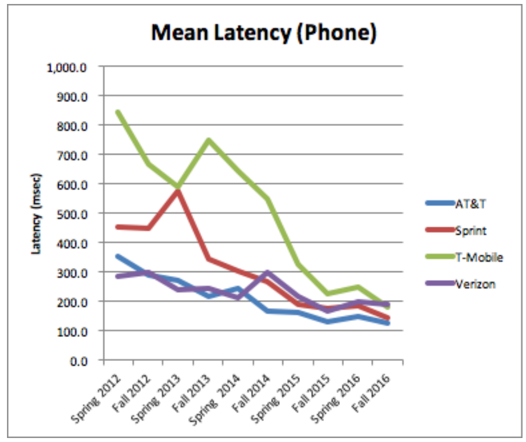

Latency varies across the state for all carriers.

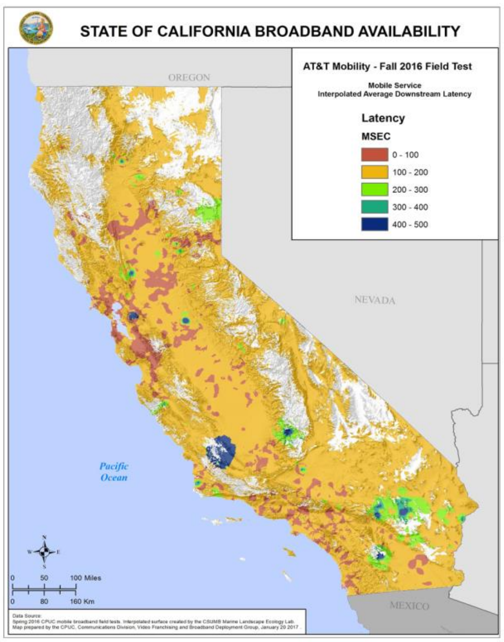

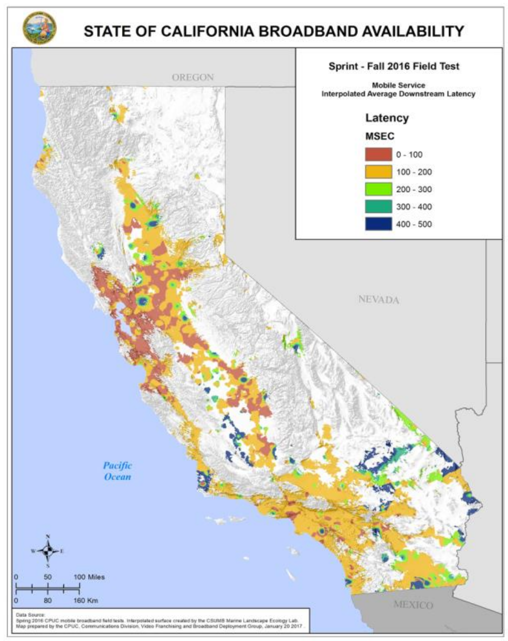

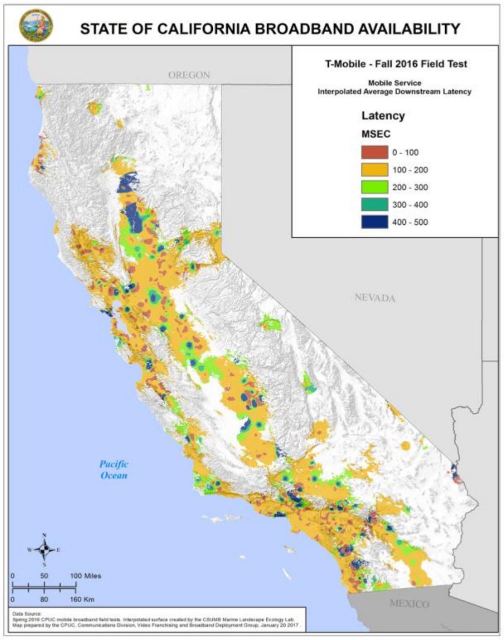

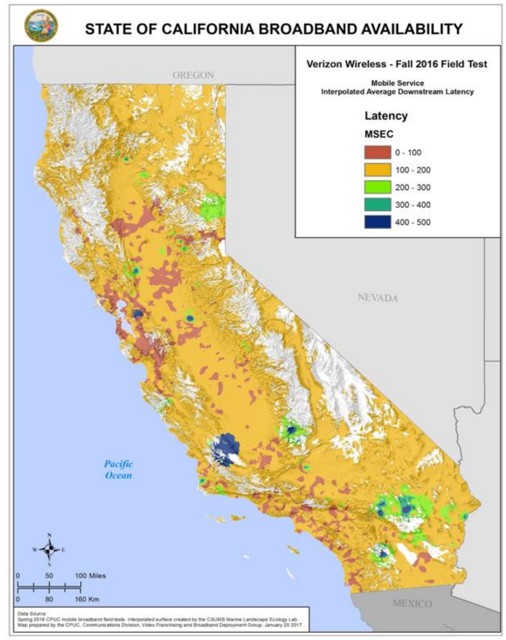

Mean jitter[7] improved for all carriers and for all demographics.

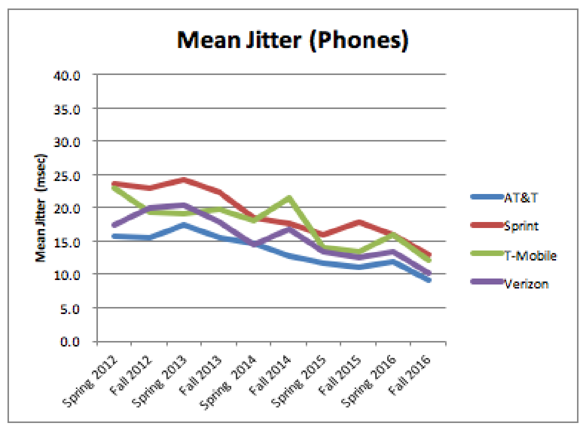

Mean packet loss[8] rates improved for AT&T and Verizon and degraded for Sprint and T-Mobile. Urban packet loss degraded Sprint and T-Mobile. Rural packet loss degraded for T-Mobile. Tribal packet loss degraded for AT&T, Sprint and T-Mobile.

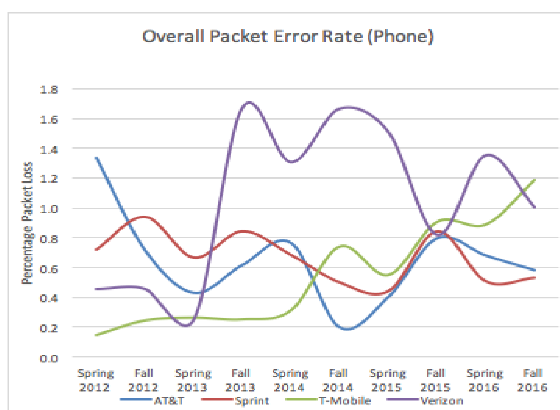

[7] Jitter is the measure of variation in latency. It is important component of performance for real time streaming of audio and video.
[8] Mean packet loss is the average percentage of packets that are lost during transmission. Small increases in packet loss are particularly degrading for streaming media services.

TCP connection quality[9], as measured by failed TCP connection attempts[10], improved for all carriers but dramatically for both T-Mobile and Sprint after dramatic degradation in the Spring 2015 measurement. This improvement is almost entirely in rural TCP connection quality. There remain profound differences between urban and rural areas with rural users having an almost 400% higher TCP connection failure rate over urban users even for AT&T and Verizon.

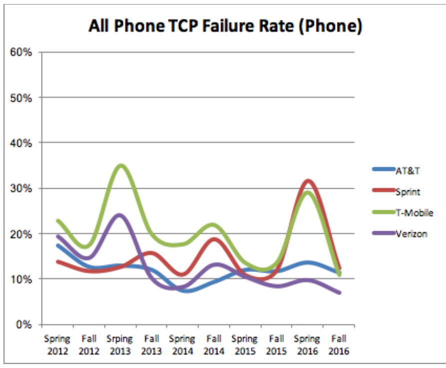

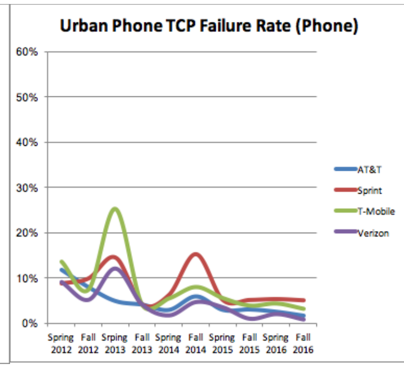

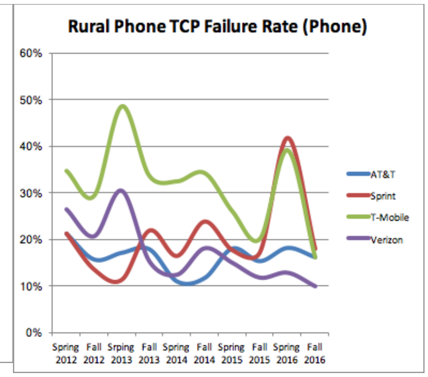

9 The fundamental reliable connection service for the Internet is TCP - Transmission Control Protocol. It provides reliable delivery of an ordered stream of bytes and is the foundation service for web browsing, most streaming media services, email, IM and most other user Internet services. CalSPEED measures 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, the call fails to connect to the destination.

7. Over-The-Top Voice MOS Service Quality Remains Degraded

The Spring 2016 degradation in underlying Internet quality metrics, particularly in latency, jitter and packet loss, means a substantial decrease in the OTT MOS quality for all four carriers. This degradation was largely recovered in the Fall 2016 measurement and reflects in the underlying improvement in quality of latency, jitter and packet loss. However, OTT VoIP coverage has still not returned to the peak of Spring 2015.

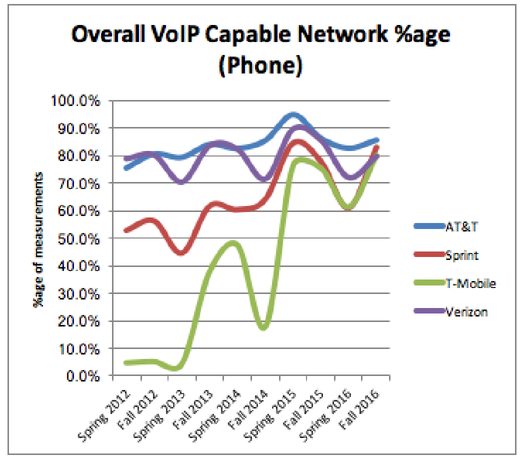

While urban MOS has largely recovered for all carriers, all carriers still show material degradation for rural and tribal users. Sprint, T-Mobile and Verizon rural and tribal users show some (partial) recovery but AT&T rural and tribal users show materially increased degradation.

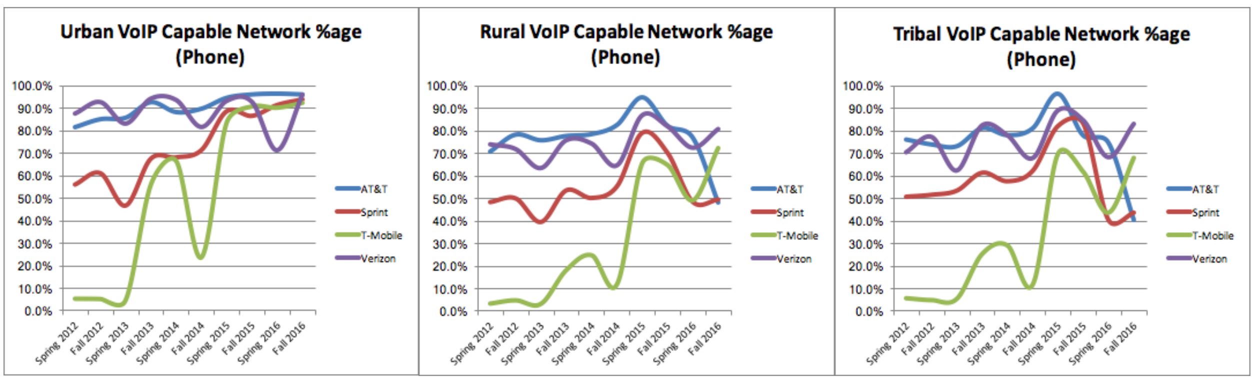

CalSPEED maps OTT voice MOS across the state for each carrier. AT&T and Verizon both have extensive OTT voice service across the state, with Sprint and T-Mobile limited to dense urban areas and major transportation routes.

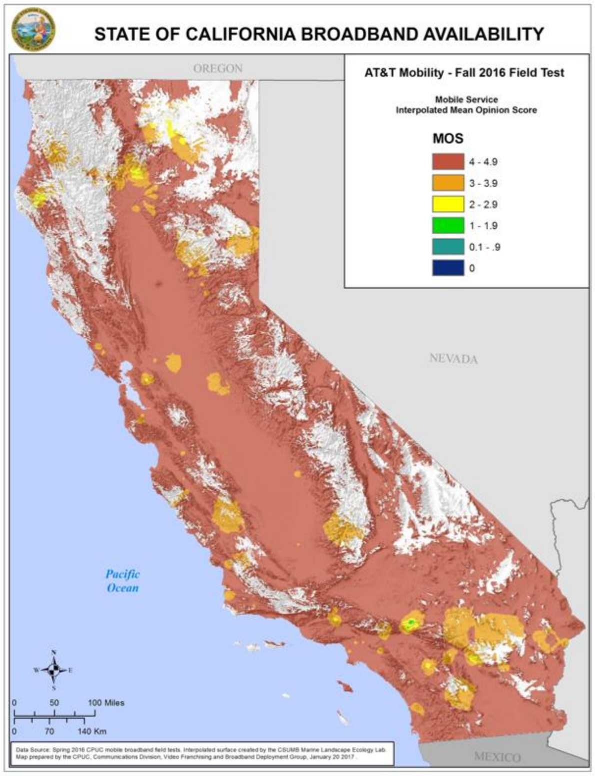
STATE OF CALIFORNIA BROADBAND AVAILABILITY
AT&T Mobility - Fall 2016 Field Test
Mobile Service
Interpolated Mean Opinion Score
MOS
4 - 4.9
3 - 3.9
2 - 2.9
1 - 1.9
0.1 - .9
0
OREGON
NEVADA
Pacific Ocean
MEXICO
0 50 100 Miles
0 70 140 Km
Data Source: Spring 2016 CPUC mobile broadband field tests. Interpolated surface created by the CSUMB Marine Landscape Ecology Lab.
Map prepared by the CPUC, Communications Division, Video Franchising and Broadband Deployment Group, January 20 2017.

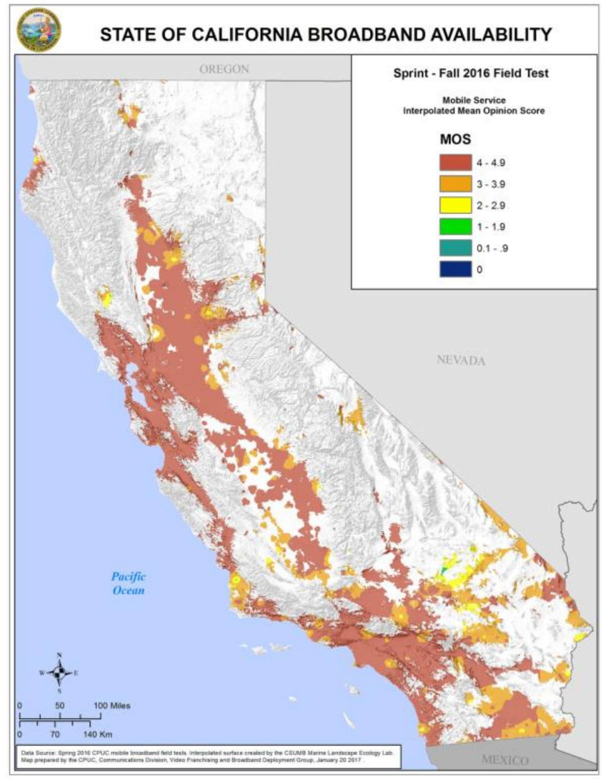
STATE OF CALIFORNIA BROADBAND AVAILABILITY
Sprint - Fall 2016 Field Test
Mobile Service
Interpolated Mean Opinion Score
MOS
4 - 4.9
3 - 3.9
2 - 2.9
1 - 1.9
0.1 - .9
0
OREGON
NEVADA
Pacific Ocean
MEXICO
0 50 100 Miles
0 70 140 Km
Data Source: Spring 2016 CPUC mobile broadband field tests. Interpolated surface created by the CSUMB Marine Landscape Ecology Lab.
Map prepared by the CPUC, Communications Division, Video Franchising and Broadband Deployment Group, January 20 2017.

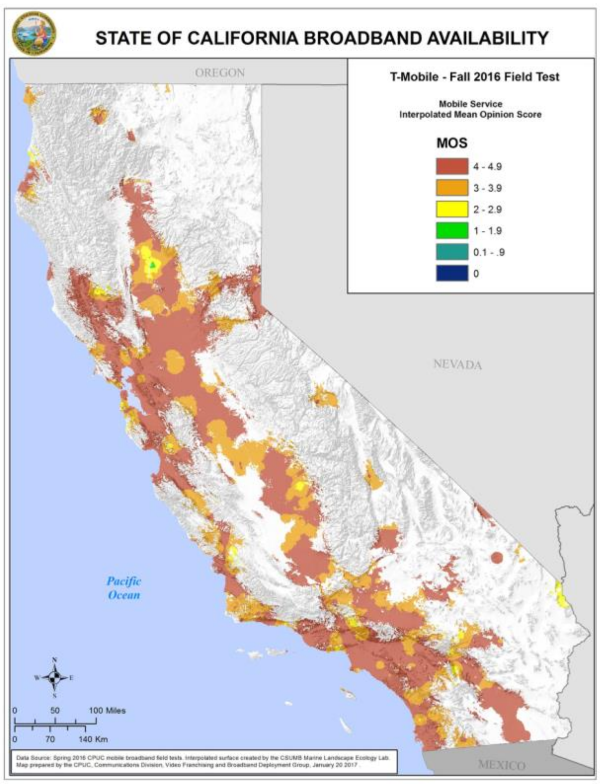
STATE OF CALIFORNIA BROADBAND AVAILABILITY
T-Mobile - Fall 2016 Field Test
Mobile Service
Interpolated Mean Opinion Score
MOS
4 - 4.9
3 - 3.9
2 - 2.9
1 - 1.9
0.1 - .9
0
OREGON
NEVADA
Pacific Ocean
MEXICO
0 50 100 Miles
0 70 140 Km
Data Source: Spring 2016 CPUC mobile broadband field tests. Interpolated surface created by the CSUMB Marine Landscape Ecology Lab.
Map prepared by the CPUC, Communications Division, Video Franchising and Broadband Deployment Group, January 20 2017.

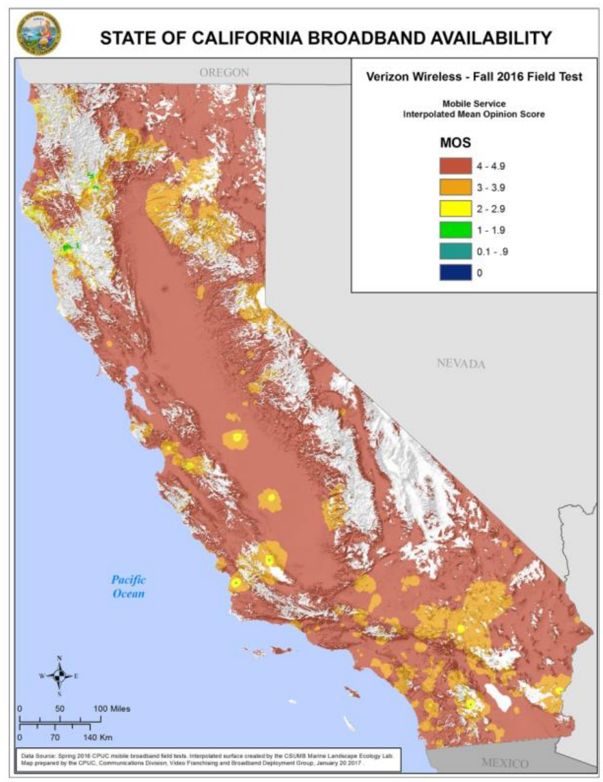
STATE OF CALIFORNIA BROADBAND AVAILABILITY
Verizon Wireless - Fall 2016 Field Test
Mobile Service
Interpolated Mean Opinion Score
MOS
4 - 4.9
3 - 3.9
2 - 2.9
1 - 1.9
0.1 - .9
0
OREGON
NEVADA
Pacific Ocean
MEXICO
0 50 100 Miles
0 70 140 Km
Data Source: Spring 2016 CPUC mobile broadband field tests. Interpolated surface created by the CSUMB Marine Landscape Ecology Lab.
Map prepared by the CPUC, Communications Division, Video Franchising and Broadband Deployment Group, January 20 2017.

8. LTE Deployment Coverage Appears Capped

The deployment of LTE as well as the replacement of legacy 1G and 2G wireless access networks appears capped as of Spring 2016 and that cap continues into Fall 2016. The data suggests that as of the Fall 2016 measurement deployment of urban LTE is capped at between 85-95% of the measurement locations. Rural LTE deployment is at 80% of measured locations (Sprint at 65%) with a suggestion of slight trend for modest improvement. Urban LTE deployment has remained constant 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).

Rural users will experience encountering legacy service between 5 (AT&T, Sprint and T-Mobile) and 10 times more often than urban users. 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.

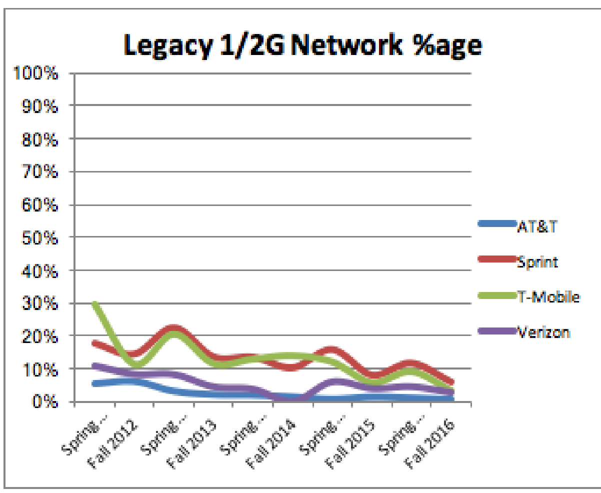

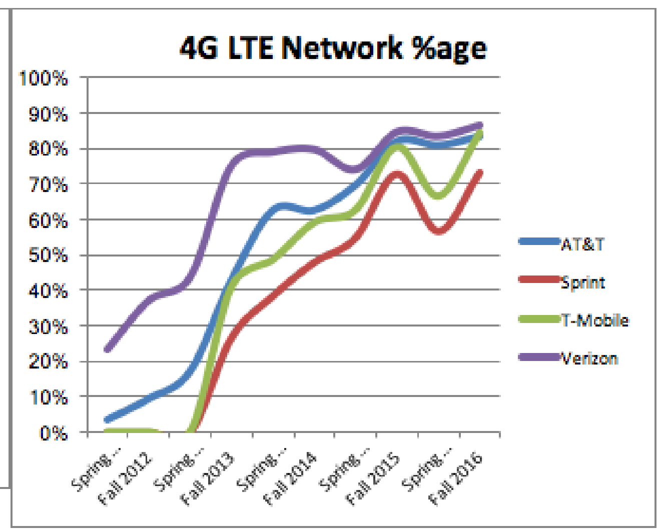

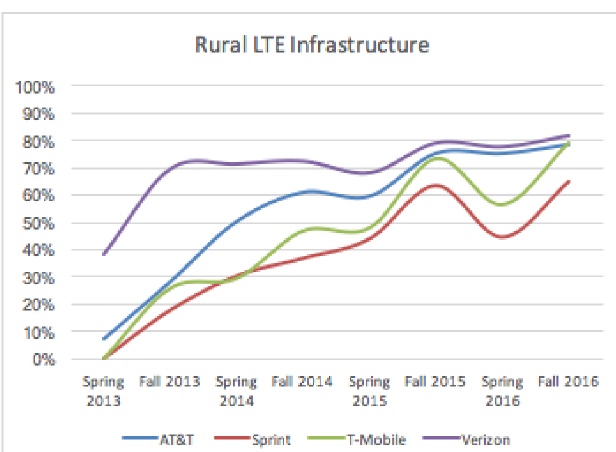

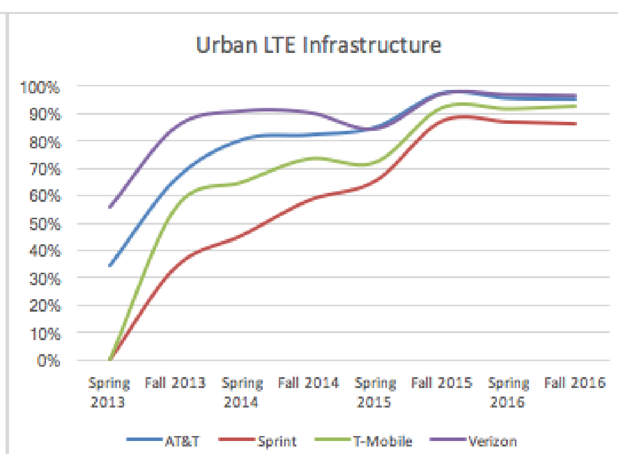

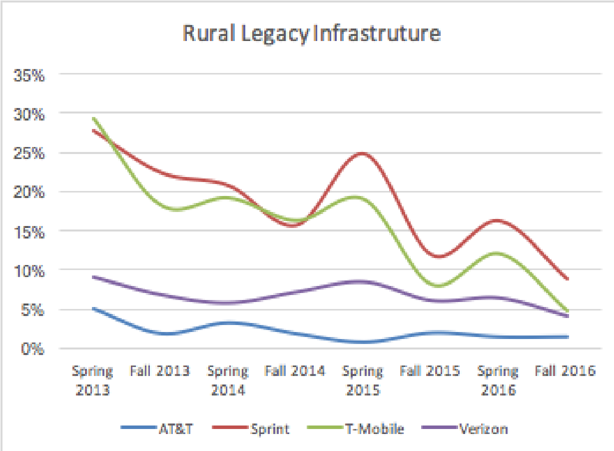
Rural Legacy Infrastruture
35%
30%
25%
20%
15%
10%
5%
0%
Spring 2013
Fall 2013
Spring 2014
Fall 2014
Spring 2015
Fall 2015
Spring 2016
Fall 2016
AT&T
Sprint
T-Mobile
Verizon

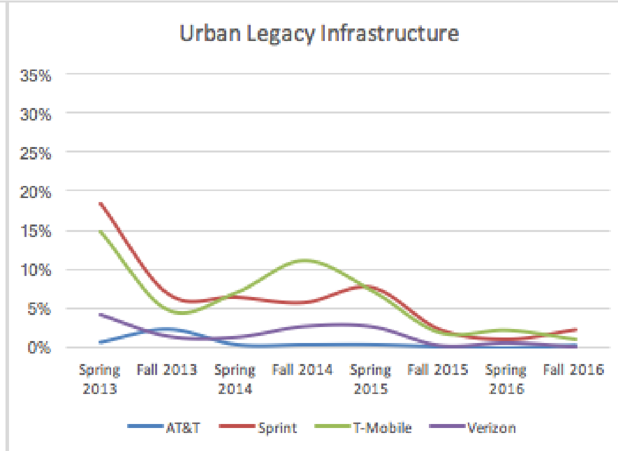
Urban Legacy Infrastructure
35%
30%
25%
20%
15%
10%
5%
0%
Spring 2013
Fall 2013
Spring 2014
Fall 2014
Spring 2015
Fall 2015
Spring 2016
Fall 2016
AT&T
Sprint
T-Mobile
Verizon

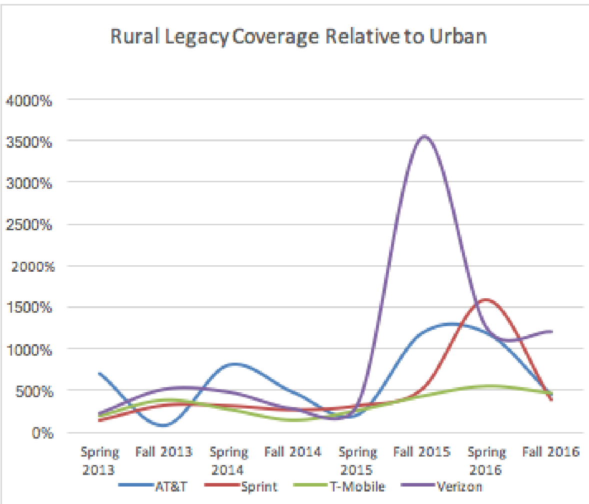
Rural Legacy Coverage Relative to Urban
4000%
3500%
3000%
2500%
2000%
1500%
1000%
500%
0%
Spring 2013
Fall 2013
Spring 2014
Fall 2014
Spring 2015
Fall 2015
Spring 2016
Fall 2016
AT&T
Sprint
T-Mobile
Verizon

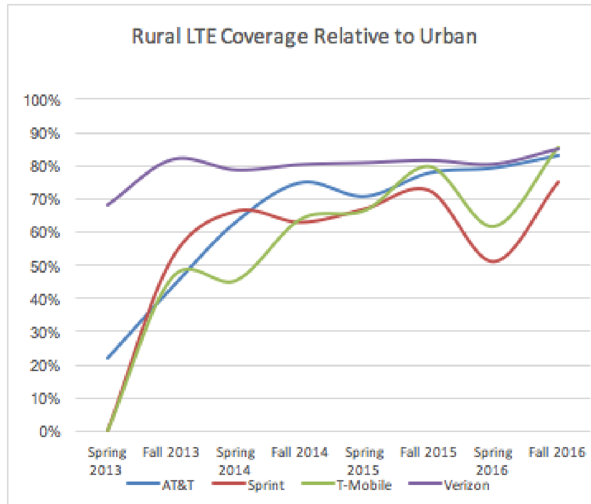
Rural LTE Coverage Relative to Urban
100%
90%
80%
70%
60%
50%
40%
30%
20%
10%
0%
Spring 2013
Fall 2013
Spring 2014
Fall 2014
Spring 2015
Fall 2015
Spring 2016
Fall 2016
AT&T
Sprint
T-Mobile
Verizon

9. Video Improves

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 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 measurement[11] and the analysis in this paper is the fourth update measurement from those conclusions.

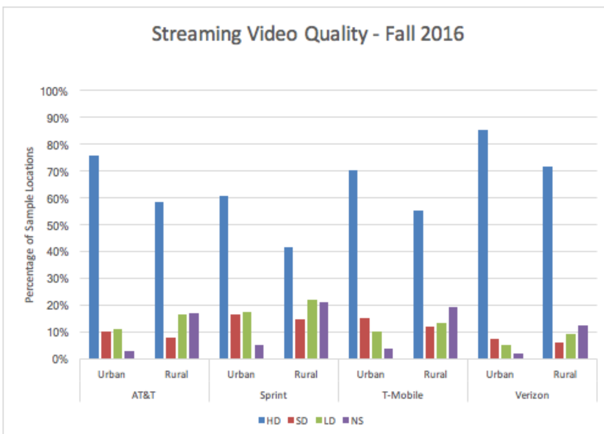

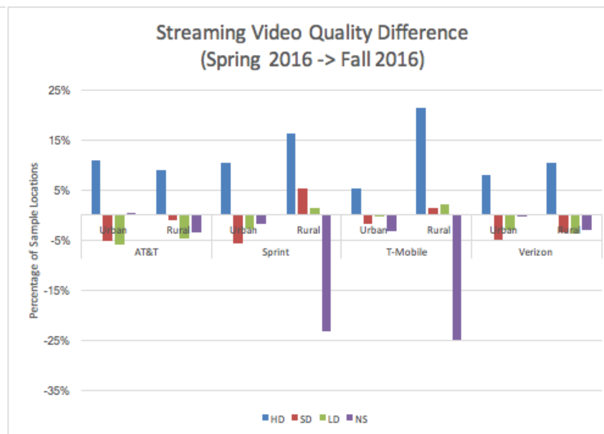

Fall 2016 streaming video quality has overall improved for all carriers, with more measurement locations reporting HD quality and fewer reporting no service or SD or LD quality. HD streaming video coverage for tribal users has decreased for all carriers, but most dramatically for Sprint and T-Mobile.

Sprint and T-Mobile show the greatest improvement, particularly in the ~25% decrease in the number of measurement locations reporting NoService.

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

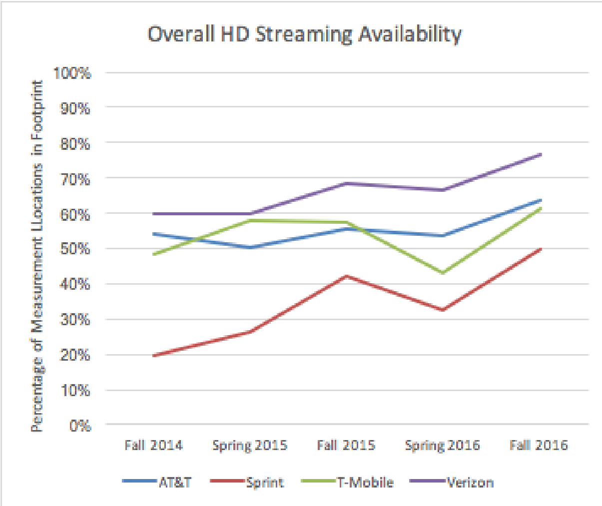

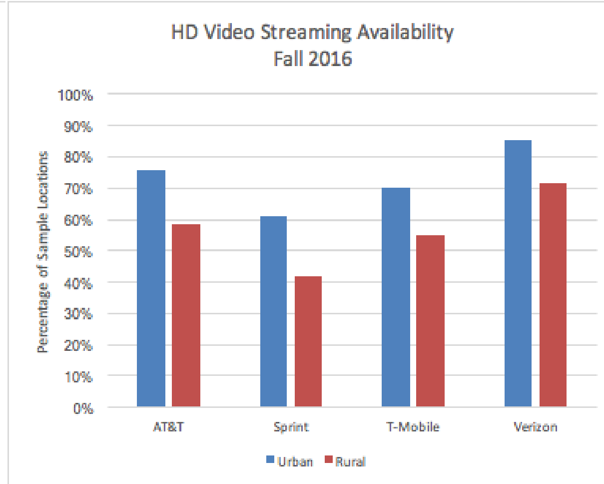

All four carriers offer HD quality video streaming in over 50% of the urban measurement locations but all carriers show material drops in HD video streaming availability for rural and tribal users. Over the last five measurement rounds rural users, for all carriers, have about 70-80% of the HD streaming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

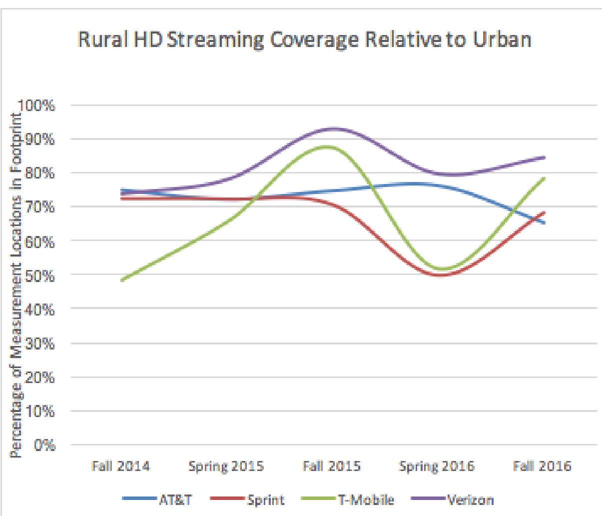

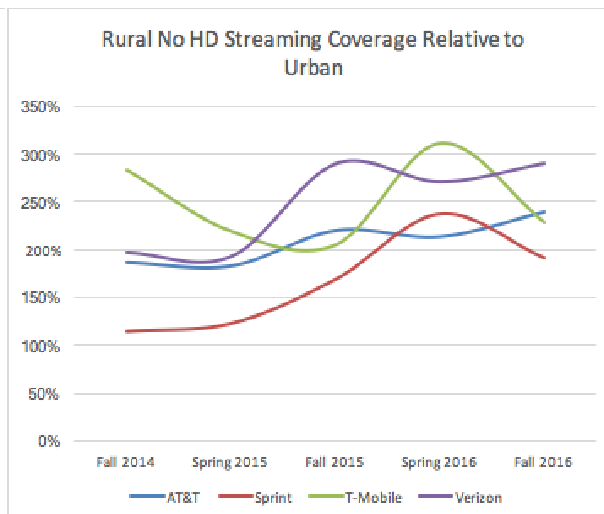

Rural users also suffer from no HD streaming coverage between 2-3x for urban users. Both HD availability and lack of coverage appear to be persistent.

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.

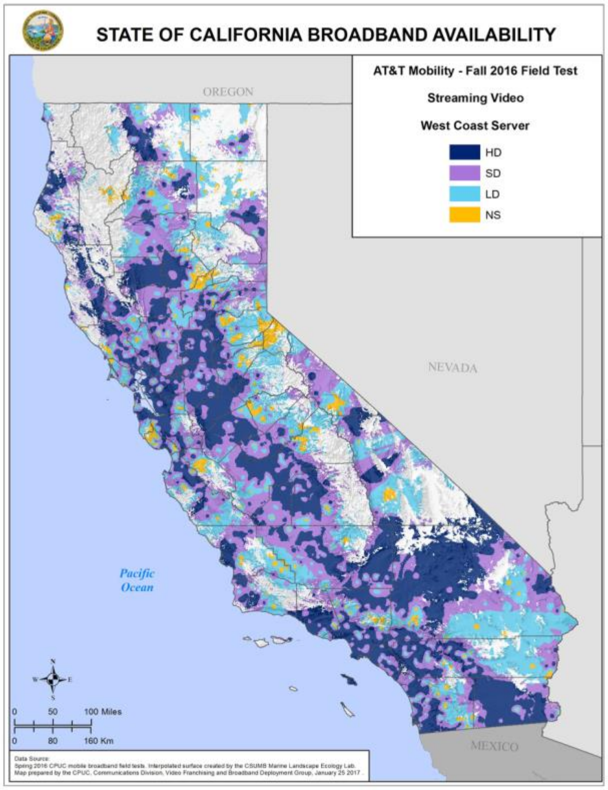

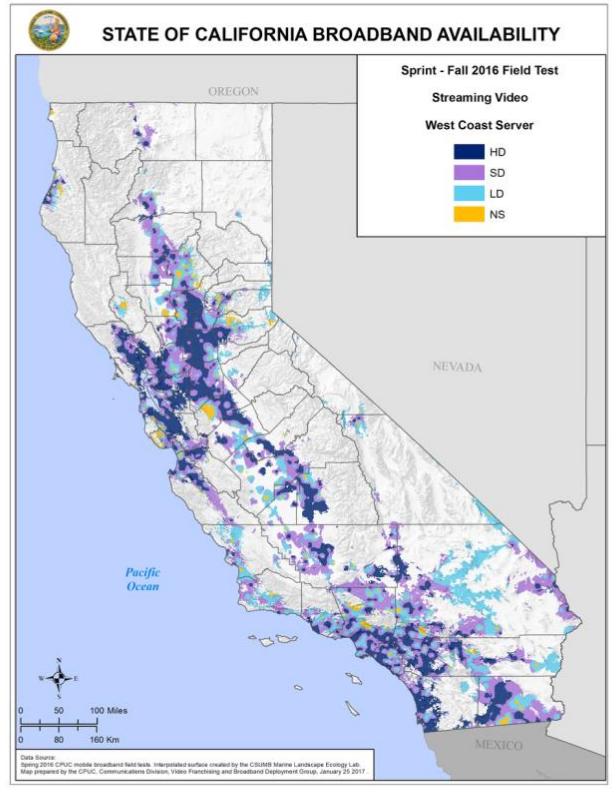

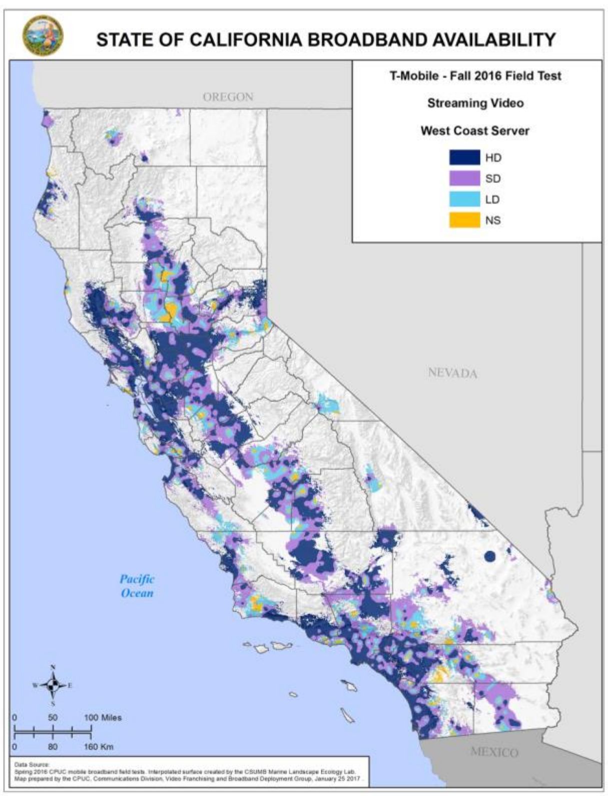

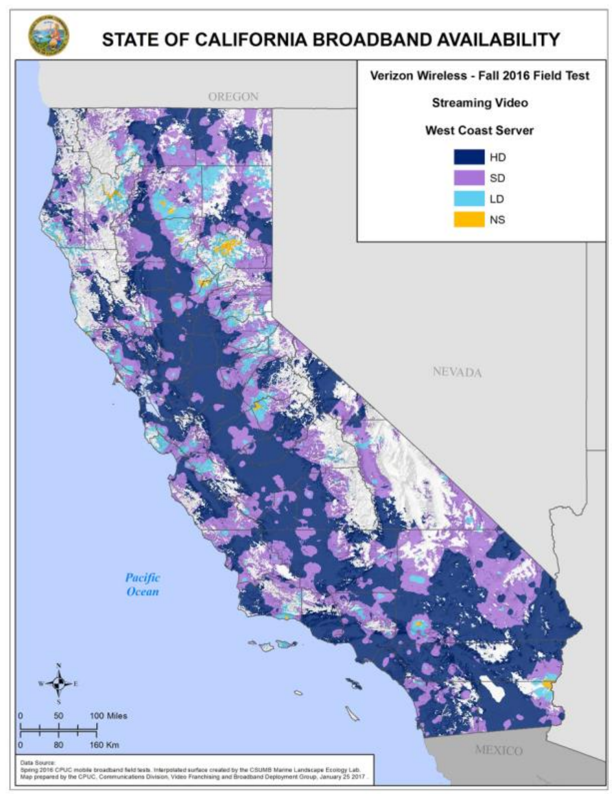

9.2 Conference Video

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.

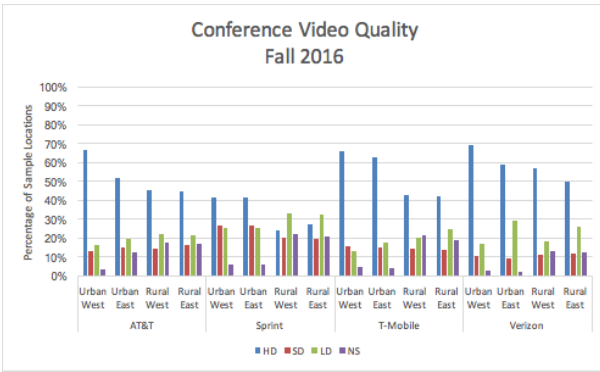

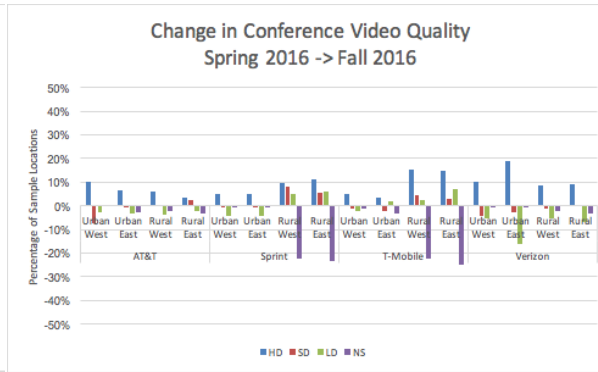

The chart on the left above, illustrates that conference video degrades in availability when moving from urban to rural users and from West to East conversations (the performance of the Internet backbone degrades the interactive video).

The Internet backbone, thru increased latency, can decrease the availability of HD quality conference video for all carriers. In recent measurements, HD conference availability is about equal between conference destinations - West or East.

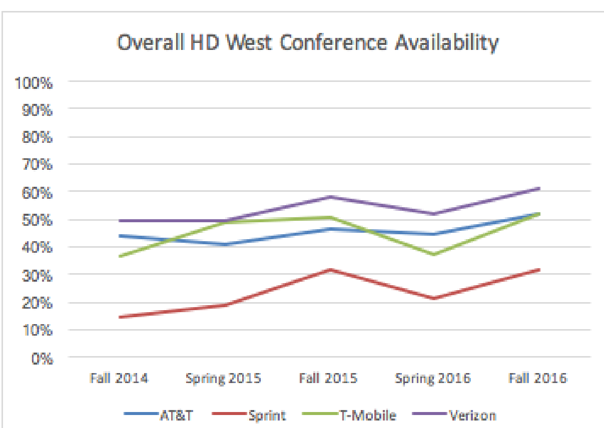

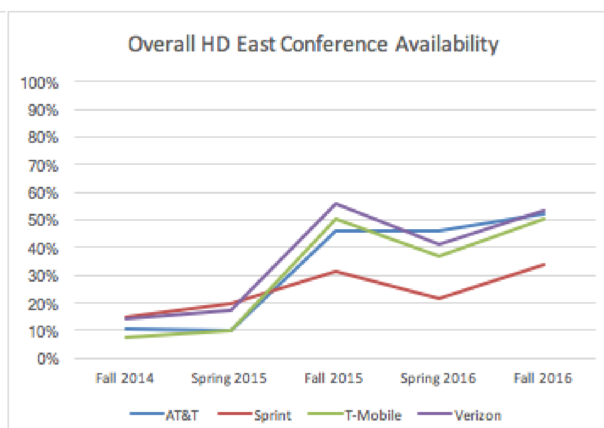

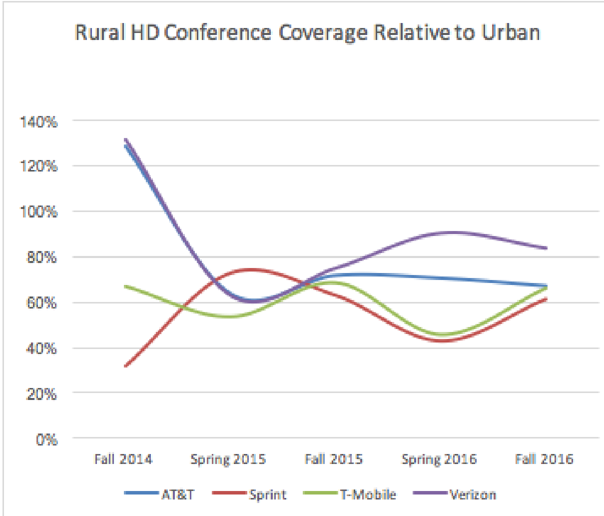

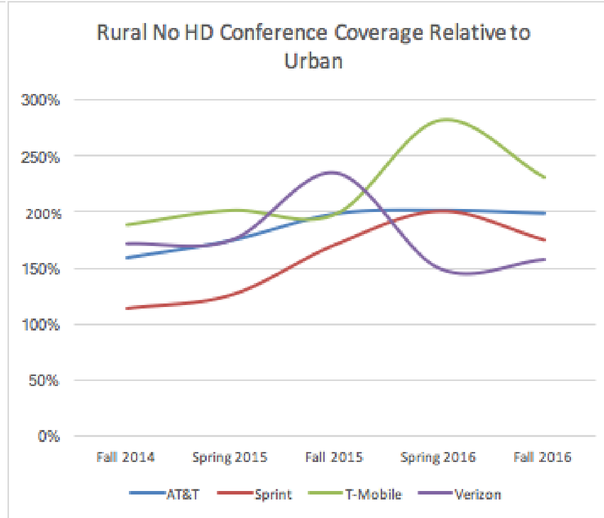

HD West and East conference availability are similar - about 50-50% of the measured locations (with Sprint around 30%). However (West illustrated above), rural users are materially less likely to have conference service availability. Only about 60-80% of rural users (compared to urban users) have access to HD conferencing and about 2x more rural users than urban users will have no conference ability at all.

Interactive video service quality is mapped across the state for all four carriers in the following charts. With the increased East availability beginning in Spring 2016, service quality East vs West for all carriers is similar.

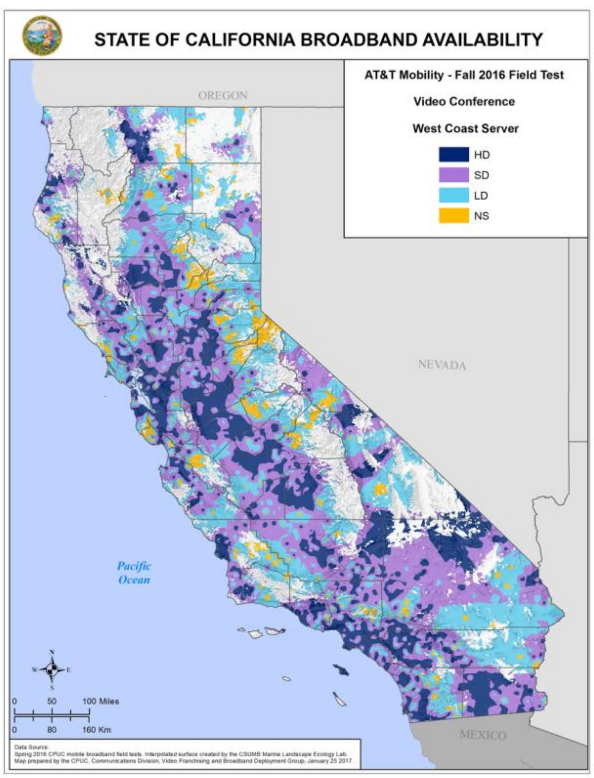

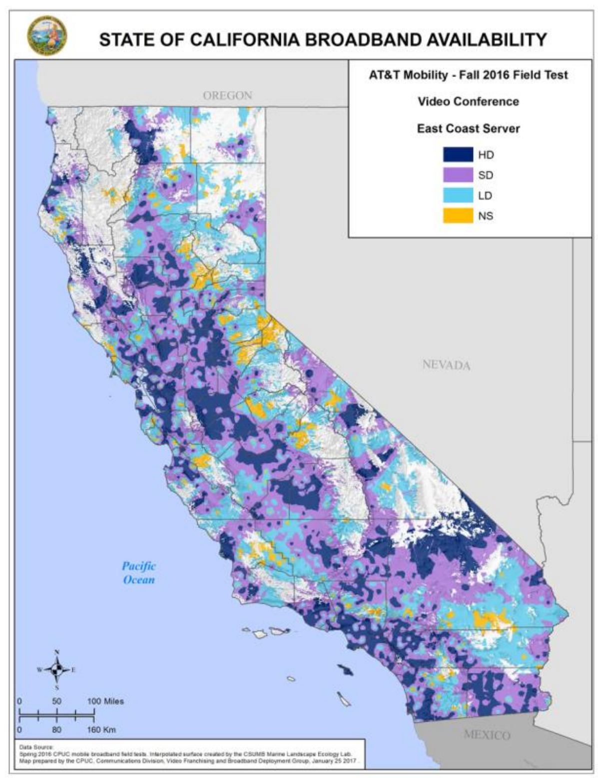

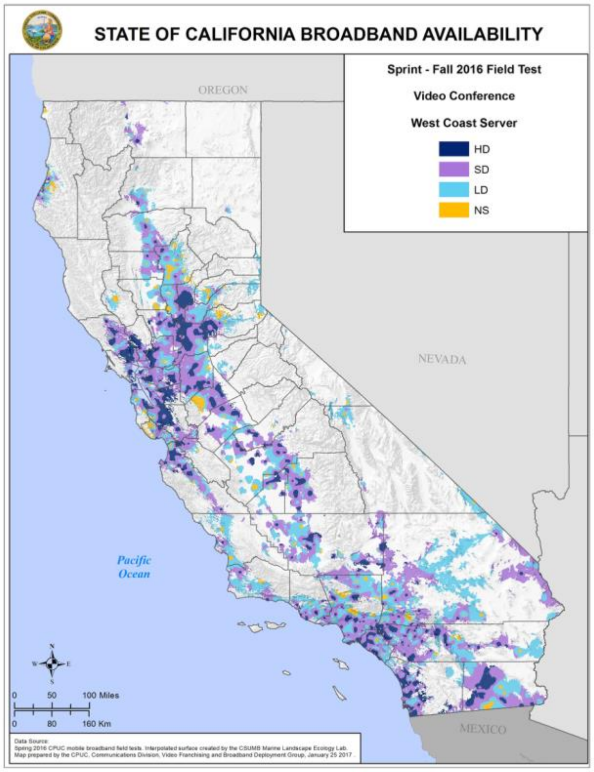
STATE OF CALIFORNIA BROADBAND AVAILABILITY
Sprint - Fall 2016 Field Test
Video Conference
West Coast Server
HD
SD
LD
NS
OREGON
NEVADA
Pacific Ocean
MEXICO
0 50 100 Miles
0 80 160 Km
Data Source:
Spring 2016 CPUC mobile broadband field tests. Interpolated surface created by the CSUMB Marine Landscape Ecology Lab.
Map prepared by the CPUC, Communications Division, Video Franchising and Broadband Deployment Group, January 25 2017

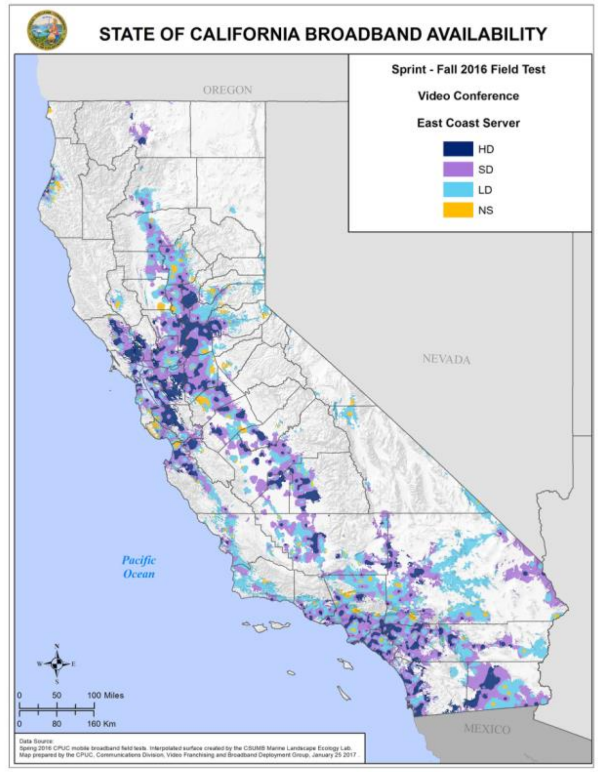
STATE OF CALIFORNIA BROADBAND AVAILABILITY
Sprint - Fall 2016 Field Test
Video Conference
East Coast Server
HD
SD
LD
NS
OREGON
NEVADA
Pacific Ocean
MEXICO
0 50 100 Miles
0 80 160 Km
Data Source:
Spring 2016 CPUC mobile broadband field tests. Interpolated surface created by the CSUMB Marine Landscape Ecology Lab.
Map prepared by the CPUC, Communications Division, Video Franchising and Broadband Deployment Group, January 25 2017

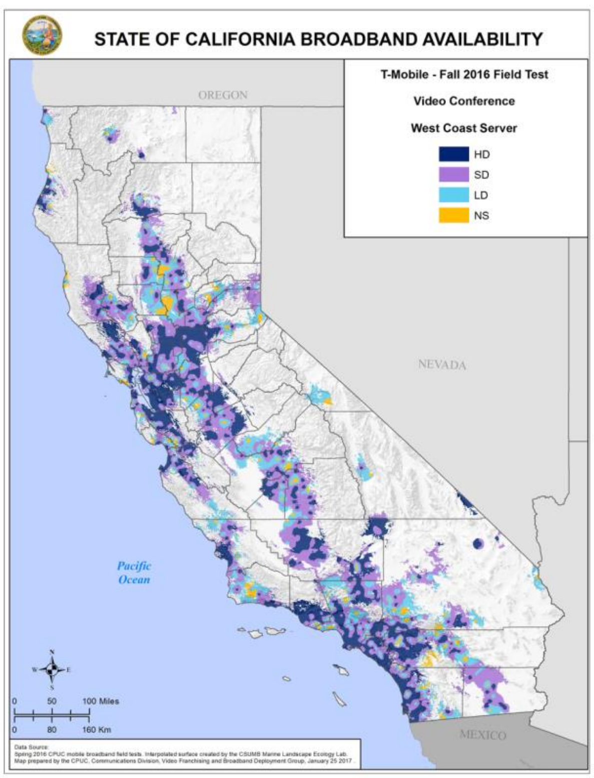
STATE OF CALIFORNIA BROADBAND AVAILABILITY
T-Mobile - Fall 2016 Field Test
Video Conference
West Coast Server
HD
SD
LD
NS
OREGON
NEVADA
Pacific Ocean
MEXICO
0 50 100 Miles
0 80 160 Km
Data Source:
Spring 2016 CPUC mobile broadband field tests. Interpolated surface created by the CSUMB Marine Landscape Ecology Lab.
Map prepared by the CPUC, Communications Division, Video Franchising and Broadband Deployment Group, January 25 2017

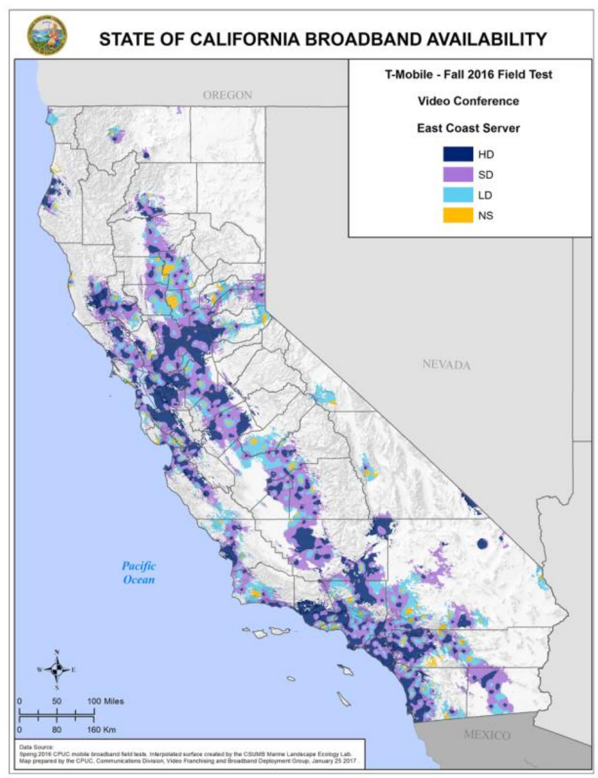
STATE OF CALIFORNIA BROADBAND AVAILABILITY
T-Mobile - Fall 2016 Field Test
Video Conference
East Coast Server
HD
SD
LD
NS
OREGON
NEVADA
Pacific Ocean
MEXICO
0 50 100 Miles
0 80 160 Km
Data Source:
Spring 2016 CPUC mobile broadband field tests. Interpolated surface created by the CSUMB Marine Landscape Ecology Lab.
Map prepared by the CPUC, Communications Division, Video Franchising and Broadband Deployment Group, January 25 2017

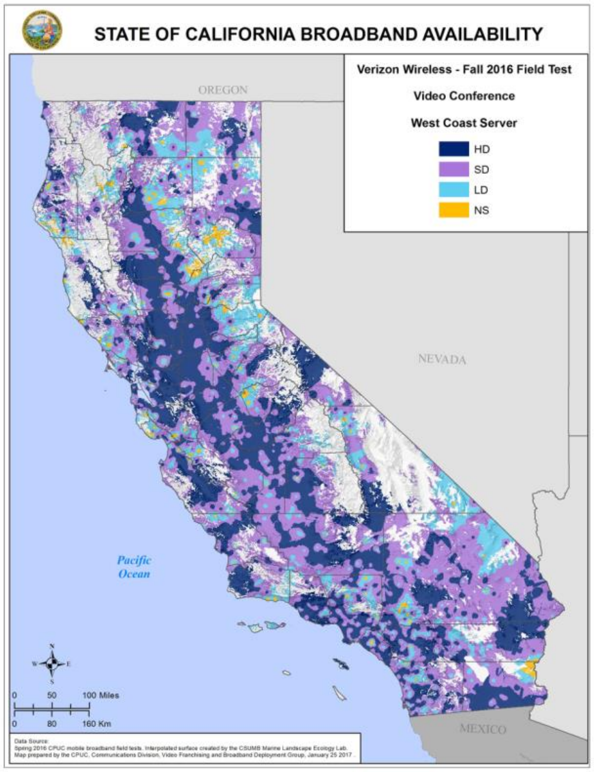
STATE OF CALIFORNIA BROADBAND AVAILABILITY
Verizon Wireless - Fall 2016 Field Test
Video Conference
West Coast Server
HD
SD
LD
NS
OREGON
NEVADA
Pacific Ocean
MEXICO
N
W
E
S
0 50 100 Miles
0 80 160 Km
Data Source:
Spring 2016 CPUC mobile broadband field tests. Interpolated surface created by the CSUMB Marine Landscape Ecology Lab.
Map prepared by the CPUC, Communications Division, Video Franchising and Broadband Deployment Group, January 25 2017.

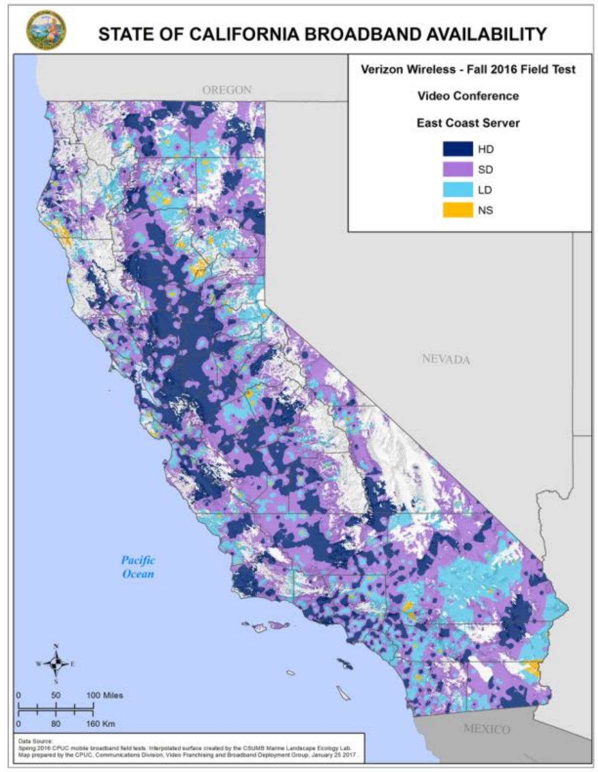
STATE OF CALIFORNIA BROADBAND AVAILABILITY
Verizon Wireless - Fall 2016 Field Test
Video Conference
East Coast Server
HD
SD
LD
NS
OREGON
NEVADA
Pacific Ocean
MEXICO
N
W
E
S
0 50 100 Miles
0 80 160 Km
Data Source:
Spring 2016 CPUC mobile broadband field tests. Interpolated surface created by the CSUMB Marine Landscape Ecology Lab.
Map prepared by the CPUC, Communications Division, Video Franchising and Broadband Deployment Group, January 25 2017.

10. Network Latency Difference Stabilizes, Network Throughput Difference Continues Poor

CalSPEED measures performance to two 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").

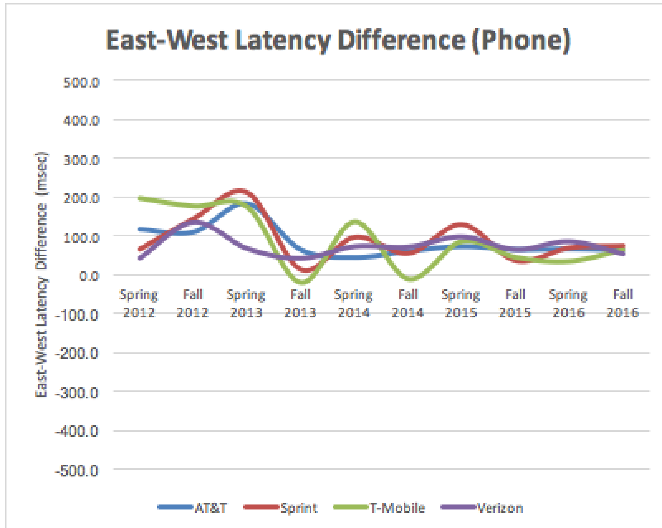

In the best case, the physics of data transmissions[12] adds about 80 msec of additional latency to get from one Coast to another - in addition to any local wireless access latency. Additional latency differences over 80 msec suggests suboptimal carrier Internet routing choices for traffic between East 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.

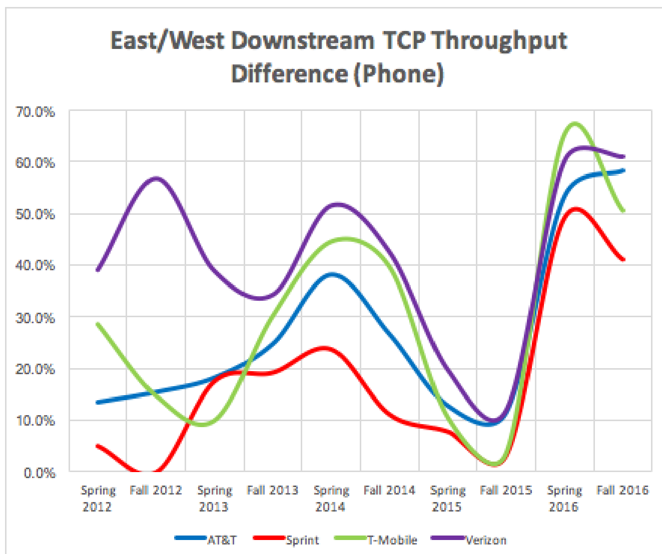

TCP throughput is technically related to latency … the longer the latency, the smaller the throughput[13]. Historically, we have seen that downstream throughput from the East server to California clients is 10-50% less than throughput from the West server. The chart to the right demonstrates this fact.

In recent years (Fall 2014 thru Spring 2015), note the substantial decrease in the TCP throughput difference and in the Spring 2016 the dramatic increase in the TCP difference for all carriers. This dramatic decrease in TCP throughput to the East server has continued into Fall 2016.

[12] Including the speed of light.
[13] A consequence of TCP's data reliability and congestion control mechanisms.

11. Conclusions

This is an analysis of the Round 10, Fall 2016 dataset for CalSPEED. It shows substantial changes happening in the wireless broadband networks in California for all carriers.

Persistent mobile digital divide	The rural/urban mobile digital divide appears to be persistent with rural users having material (about 1/3), often dramatic, poorer mobile broadband service over a period of years with little indication of improvement.
Mean throughput recovers	Mean throughput recovers for all carriers from decreases noted in Spring 2016 and trends begin to demonstrate differences between carriers in matching network capacity to user demand.
Carrier throughput throttling returns	Throughput throttling for all carriers returns after temporarily disappearing in the Spring 2016 measurement. 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 upstream throttling threshold appears to have increased to 20 Mb/s (15 Mb/s for Sprint) from 15 Mb/s in Fall 2015.
Mobile broadband coverage decreases	Broadband coverage (at the 25 Mb/s down, 3 Mb/s up threshold) is increasing for all carriers from a Spring 2015 low of under 5% to between 5% (Sprint) to 23% (Verizon) . Verizon increased from ~3% coverage in Spring of 2015 to ~23% coverage in the Spring of 2016.
Underlying service quality improves	Mean latency, jitter and TCP reliability modestly improved for all carriers.
Over the Top VoIP quality remains degraded, particularly for rural users	VoIP quality improved after substantial degradation in Spring 2016. Rural 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.
Internet latency stabilizes, throughput penalty continues substantial	The latency difference between East and West servers is has stabilized at about 80 msec, while throughput differences have dramatically grown to 60% for all carriers.
Internet video and conferencing recovers	After decreases in both streaming and conferencing coverage in Spring 2016, coverage increased for both services in Fall 2016.

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 smartphone and current USB datastick for laptops. 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 - iPerf[14]. 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

[14] http://en.wikipedia.org/wiki/Iperf

Internet backbone, to several ultimate server destinations.

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 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[15] 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.

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

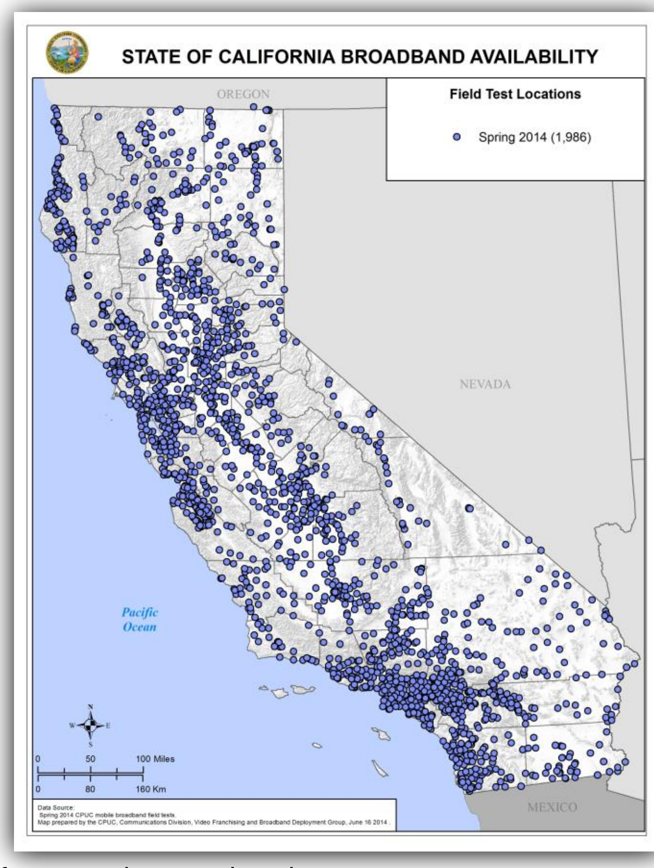

The second method is the independent use of CalSPEED to provide crowdsourced data. The structured sampling program avoids selection bias of when and where measurements are made, giving a full map that covers the entire state, including places not often visited by smartphone users but having mobile broadband service. The crowd sourced data adds additional detail to areas where 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.

<u>Not Just Data but Voice and Video.</u> 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 Googles' video quality metric[16] 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 affect of the Internet backbone) and uses both upstream and downstream performance measures.

<u>Maps for decision-makers not just for information.</u> We then take the measurement data and create geospatial kriging[17] maps interpolating CalSPEED measurements of (but not limited to)

[16] https://www.google.com/get/videoqualityreport/#methodology
[17] http://en.wikipedia.org/wiki/Kriging

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.

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.

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