Draft

CalSPEED Mobile: A Final Report on California Mobile Broadband - Six Years of Mobile Broadband Measurement

Improving Performance, Abundant Variability, Persistent Urban/Rural Digital Divide, and Carrier Traffic Shaping

Ken Biba
Managing Director and CTO
Novarum, Inc.

This is the Final Report and analysis of the six years of CalSPEED measurement of mobile broadband across California. CalSPEED Mobile consisted of twelve rounds of biannual state-wide measurements of mobile broadband quality for the four major mobile carriers in California - AT&T, Sprint, T-Mobile and Verizon. The CalSPEED Mobile project has documented the key transition from modest 3G quality service, to high performance 4G LTE service. There are key lessons to learn as the California mobile broadband network transitions to 5G services.

Mobile broadband quality improves overall

Mobile Internet quality has seen dramatic improvement from 2012 thru 2017:

- A 4x improvement in average throughput
- A 2x improvement in average latency
- A 1.6x degradation in packet loss rates
- A 1.7x improvement in packet jitter
- A 1.5x improvement in TCP connection reliability.

These changes have enabled the widespread deployment, across the state, of advanced communication services including voice over IP and streaming video.

Mobile broadband highly variable

Mobile broadband varies substantially by carrier (2x), infrastructure technology (10x), age of the user device (2x/3 years), location within state (10x) and destination information server (2x). But not much by time of day nor during a communication session (25%).

Rural/tribal mobile digital divide is substantial, persistent and likely to worsen

Rural and tribal users consistently experience about 3/5ths of the mobile broadband quality of their urban peers.

This appears for all broadband quality metrics of throughput, latency, packet loss, digital voice quality, digital video quality, jitter, connection reliability, broadband coverage and deployed modern digital communications technology.

Notably, 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 Internet application) while ~20% of rural users of those same carriers can expect TCP connection failures.

The deployment of 5G technology will likely improve urban users experience much more than rural users, increasing the broadband divide.

Mobile broadband is traffic shaped

Mobile broadband quality appears shaped, for all carriers, to have the following qualities:

- Capped maximum throughput upstream and downstream
- Improved median and average throughput
- Degraded performance outside of the regional Internet.

Mobile broadband coverage is modest

Broadband coverage at the 25 Mb/s down, 3 Mb/s is available at ~13% of all CalSPEED measurement locations at the end of 2017. At the degraded “broadband” standard of 10 Mb/s down and 1 Mb/s up, “coverage” improves to ~50% for all carriers.

But whichever way broadband is defined, rural users see about 50-75% of the service availability, at that standard as do urban users.

The End of 4G, the Hype of 5G, but 2G Lingers On

4G LTE has achieved high penetration for both urban and rural users. Some legacy 1G and 2G service areas remain particularly for rural users of Sprint, T-Mobile and Verizon. Notably there was no legacy 1/2G service detected for AT&T in California.

Urban users may see substantial increase in performance (5-10x?) and decrease in latency from the deployment of 5G that will make interactive and streaming services even more effective. These are unlikely to be available to rural and tribal users where population density and geography make 5G mmwave deployments uneconomical or physically impossible.

It is likely that the current mobile broadband digital divide between urban and rural/tribal users will not only widen, but widen dramatically.

Older Devices Mean Slower and Lower Quality Internet

CalSPEED measurements were largely made with the latest generation user devices in order to assess the deployed infrastructure. But CalSPEED has also measured with a variety of other devices in parallel.

What can be learned from this data?

- Different devices can have dramatically different performance.
- This difference could easily be ~25%/year which means a 3 year old user device (a phone with 3 year old technology) can easily be a factor of 2 lower in performance than a state of the art device.

Many users do not choose, or cannot afford, the latest technology. These users will not have the performance and quality of service documented by CalSPEED - but rather something substantially less. These user device choices amplify other differences of carrier and location.

The Signal is Not the Message

The signal bars on a smartphone are not a good predictor of mobile performance. At best, RSSI explains about 20% of the variation in throughput (Verizon urban), and at worst, explains 0% of the variation in throughput (AT&T urban). At best, SNR explains about 45% of throughput variation (Sprint rural) and at worst, about 6% of throughput variation (AT&T rural).

There is little ability of signal strength to predict Internet performance.

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 (or wrist!) wherever we go.

A Pew Research study[1] emphasizes that increasingly the Internet is mobile, rather than fixed.

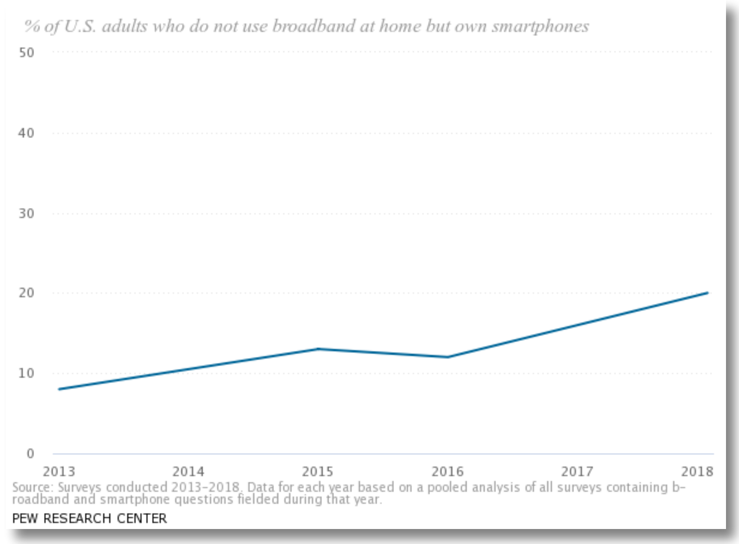

Knowing the quality of mobile broadband is a vital piece of our modern ecosystem much in the same way as we research the brand of car we drive or the type of house we own. With multiple mobile Internet providers, an independent third party assessment of this quality allows consumers and policy makers 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 and then 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. It calibrates a number of metrics of the user Internet experience and not only presents them as numbers - but maps them onto the state of California.

[1] "Mobile Fact Sheet", Pew Research Center, February 2018, http://www.pewinternet.org/fact-sheet/mobile/

CalSPEED has measured mobile broadband in California for six years with twelve rounds of mobile broadband measurement over the entire state. It has collected close to 40,000,000 measurements, at the same locations, across California of the four major mobile broadband carriers: AT&T Mobility, Sprint, T-Mobile and Verizon Wireless. CalSPEED will now turn its attention to measuring residential wired and WiFi[2] broadband.

This paper is the final report of the CalSPEED's mobile broadband measurement. Previous papers have analyzed the prior rounds of measurement[3][4][5][6][7][8]. The methodology has been rigorously analyzed with respect to other available mobile measurement tools[9]. The CalSPEED measurement methodology is discussed in greater detail in Appendix A.

[2] Ken Biba, "WiFi Internet in California - for Every Household", Novarum, January 2019.

[3] Ken Biba, "Assessment of California Mobile Broadband Spring 2014", Novarum, September 2014,https://www.dropbox.com/s/iwt8vhssv548vt5/Assessment%20of%20California%20Mobile%20Broadband%20Spring%202014.pdf?dl=0

[4] Ken Biba, "Assessment of California Mobile Broadband Fall 2014", Novarum, June 2015, https://www.dropbox.com/s/3kjkqqglpmeoa18/Assessment%20of%20California%20Mobile%20Broadband%20Fall%202014.pdf?dl=0

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

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

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

[8] Ken Biba, "Assessment of California Mobile Broadband Fall 2016", Novarum

[9] Ken Biba, "Comparison of CalSPEED, FCC and Ookla", Novarum, Inc., September 2014,https://www.dropbox.com/s/4awyr948lfz4zvv/CalSPEED%20Mobile%20Measurement%20Comparison.pdf?dl=0

2. Mean Mobile Broadband Quality Improves, but Intermittently

Mobile broadband service is delivered by mobile Internet carriers largely to hand held computing devices across the entire state. These Internet services include reliable data streams, web access, messaging, interactive voice, streaming video and conferencing. CalSPEED is unique is that its measurements are always made in the same locations.

Let's look at some of the key qualities of mobile broadband: throughput, service quality, over the top streaming voice and over the top video. Three consistent themes emerge -

- There is substantial variation in broadband quality over time, location and carrier.
- There is general trend of improvement in broadband quality with occasional regressions.
- Carriers differ dramatically in coverage across California.

Our measurements have generally used two different user devices per carrier - the first the "best" smartphone recommended by the carrier and the second an alternative device - either a laptop USB stick or a tablet. To report consistently, this analysis uses the the data collected on smartphones for each carrier.

CalSPEED measures all its metrics twice at each location - once to the "local" server (West - located in San Jose) and once to the "far" server (East - located in Northern Virginia). The measurements reported in our analyses, in general, average these measurements - attempting to mimic a user experience of accessing Internet resources across the Internet.

2.1 Throughput

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 giving higher capacity service.

Mean downstream throughput has increased roughly 3x from about 3 Mb/s in 2012 to over 10 Mb/s for all carriers by the end of 2017. Mean upstream throughput has increased from under 2 M/bs in 2012 to over 6 Mb/s. Verizon has consistently delivered the the highest mean throughput (downstream and upstream) across the state, while Sprint has consistently been the poorest throughput carrier throughout the state. AT&T and T-Mobile have dueled for middle position, each slightly poorer, on average, than leader Verizon. However, coverage maps will illustrate that AT&T covers substantially more of the state, with T-Mobile with a smaller coverage area concentrating on population centers.

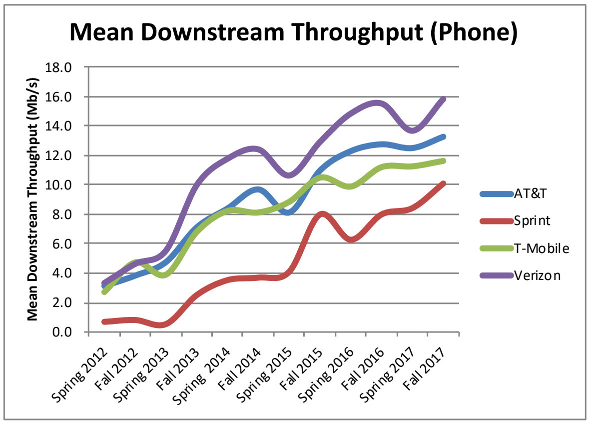

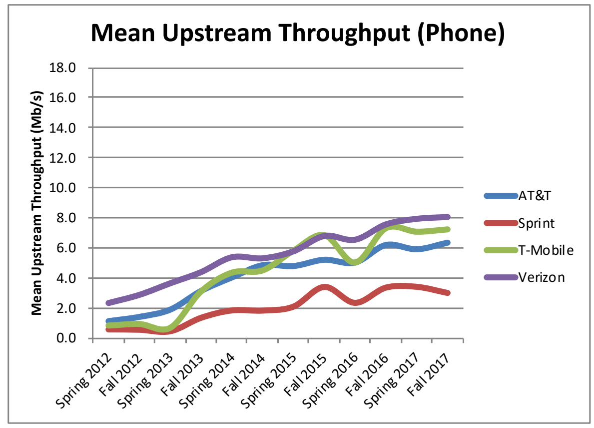

Each carrier has had substantial variation (up to 25%) in year-to-year aggregate mean throughput. We can speculate that these substantial variations in aggregate mean throughput are temporal mismatches of capacity and demand.

Looking at each carrier for each demographic[10], there is a consistent pattern of urban upstream and

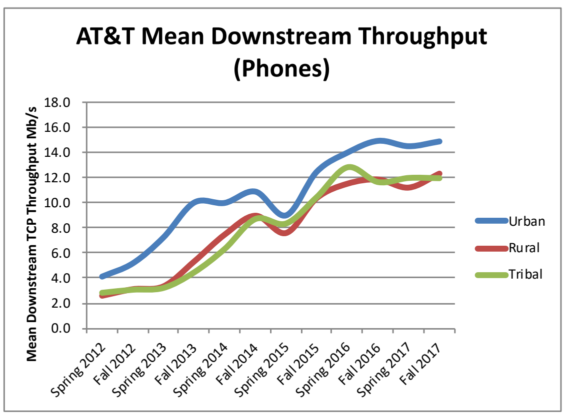

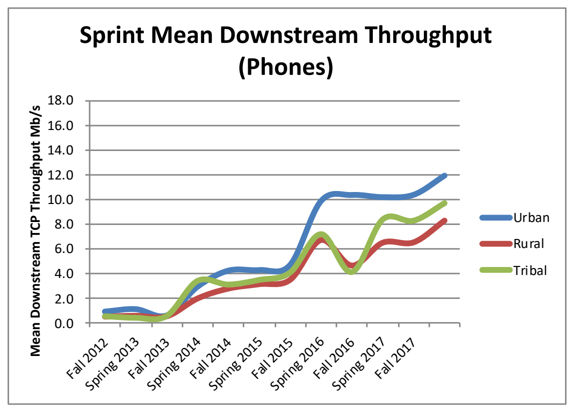

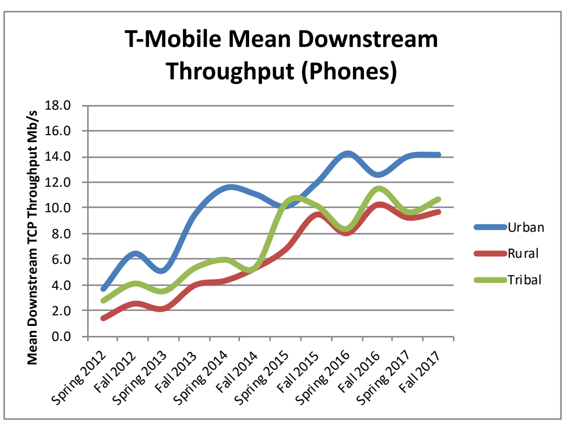

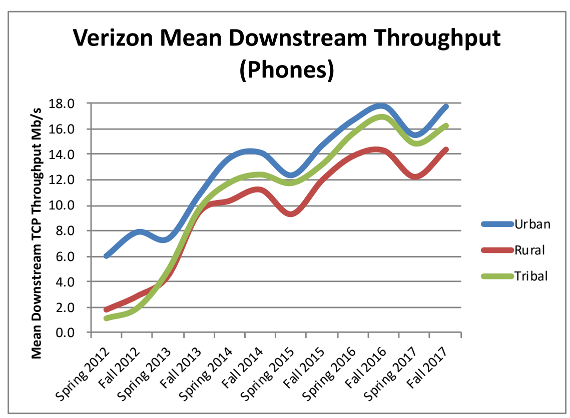

[10] CalSPEED records the GPS coordinates of each measurement and so our analysis partitions each measurement into one of three demographics: urban, rural and tribal. CalSPEED creates broadband quality maps using this same location information.

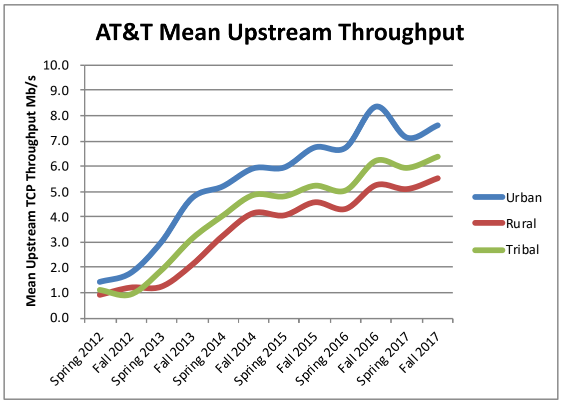

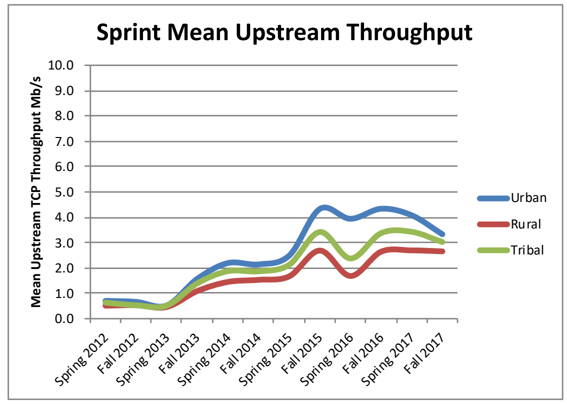

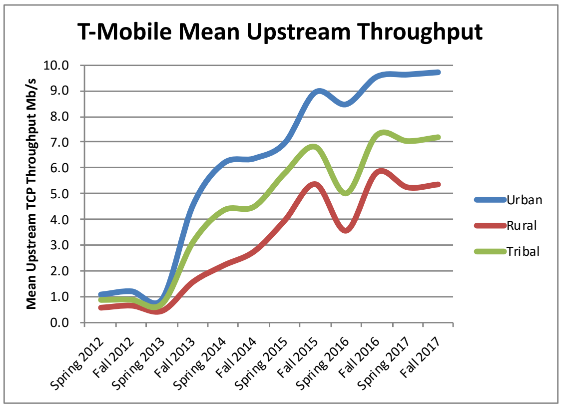

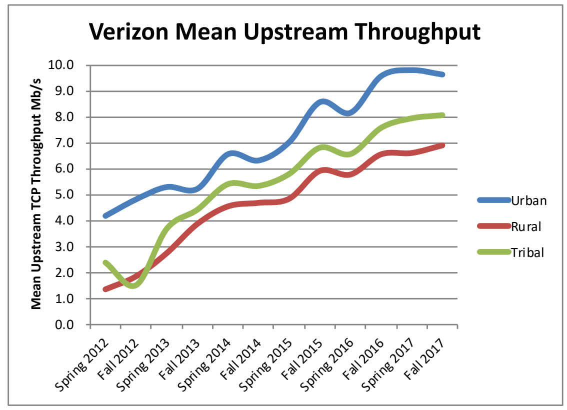

downstream mean throughput always exceeding rural and tribal performance. Rural and tribal have been quite similar with mean tribal performance occasionally exceeding rural. We will examine this difference between urban and rural mobile broadband quality in greater detail later in this paper.

2.2 Throughput Variation over Location

CalSPEED maps the estimated mean downstream and upstream throughput across the state using an interpolation that gives good fidelity to a few kilometers of resolution. The coverage maps are clipped to the carrier supplied service footprint.

Looking at the following maps for Fall of 2017 for both mean downstream and upstream throughput by carrier a few conclusions leap out.

- AT&T and Verizon have substantially larger (and similar) coverage areas, essentially most of the state, while T-Mobile and Sprint have much smaller coverage areas focussed on transportation corridors and major urban areas with minimal rural or tribal coverage.
- AT&T and Verizon - in their respective coverage areas - offer a consistent service across the state with small pockets of very good service and very poor service.
- T-Mobile's downstream service, concentrated in transportation corridors and urban areas, has extremes of service - either very high performance (generally major population centers) or very poor performance.
- Sprint's upstream service is particularly poor - offering not only the smallest coverage, but the lowest performance across the state.
- Neither T-Mobile or Sprint offer good coverage or service in rural or tribal areas.

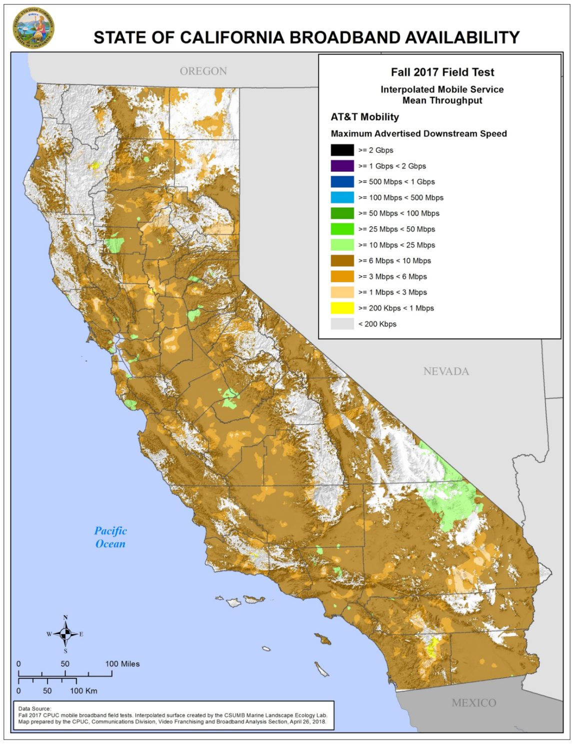
STATE OF CALIFORNIA BROADBAND AVAILABILITY
Fall 2017 Field Test
Interpolated Mobile Service Mean Throughput
AT&T Mobility
Maximum Advertised Downstream Speed
>= 2 Gbps
>= 1 Gbps < 2 Gbps
>= 500 Mbps < 1 Gbps
>= 100 Mbps < 500 Mbps
>= 50 Mbps < 100 Mbps
>= 25 Mbps < 50 Mbps
>= 10 Mbps < 25 Mbps
>= 6 Mbps < 10 Mbps
>= 3 Mbps < 6 Mbps
>= 1 Mbps < 3 Mbps
>= 200 Kbps < 1 Mbps
< 200 Kbps
OREGON
NEVADA
MEXICO
Pacific Ocean
N
W
E
S
0 50 100 Miles
0 50 100 Km
Data Source:
Fall 2017 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 Analysis Section, April 26, 2018.

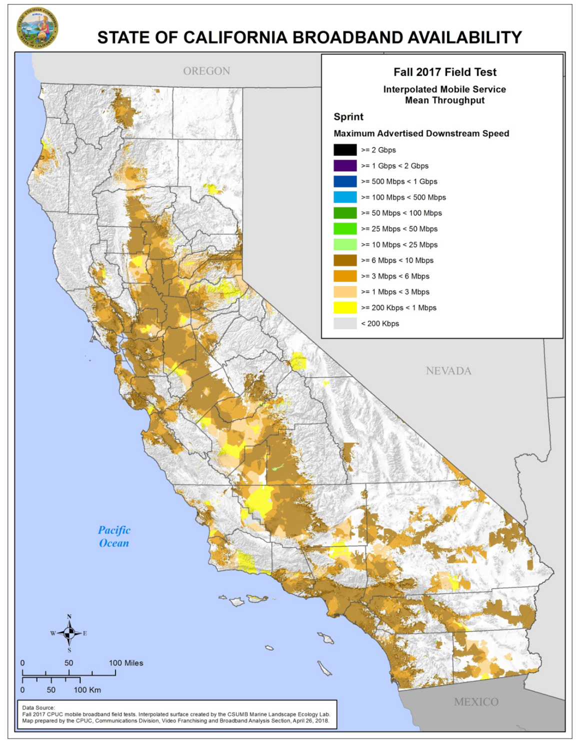
STATE OF CALIFORNIA BROADBAND AVAILABILITY
Fall 2017 Field Test
Interpolated Mobile Service Mean Throughput
Sprint
Maximum Advertised Downstream Speed
>= 2 Gbps
>= 1 Gbps < 2 Gbps
>= 500 Mbps < 1 Gbps
>= 100 Mbps < 500 Mbps
>= 50 Mbps < 100 Mbps
>= 25 Mbps < 50 Mbps
>= 10 Mbps < 25 Mbps
>= 6 Mbps < 10 Mbps
>= 3 Mbps < 6 Mbps
>= 1 Mbps < 3 Mbps
>= 200 Kbps < 1 Mbps
< 200 Kbps
OREGON
NEVADA
MEXICO
Pacific Ocean
N
W
E
S
0 50 100 Miles
0 50 100 Km
Data Source:
Fall 2017 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 Analysis Section, April 26, 2018.

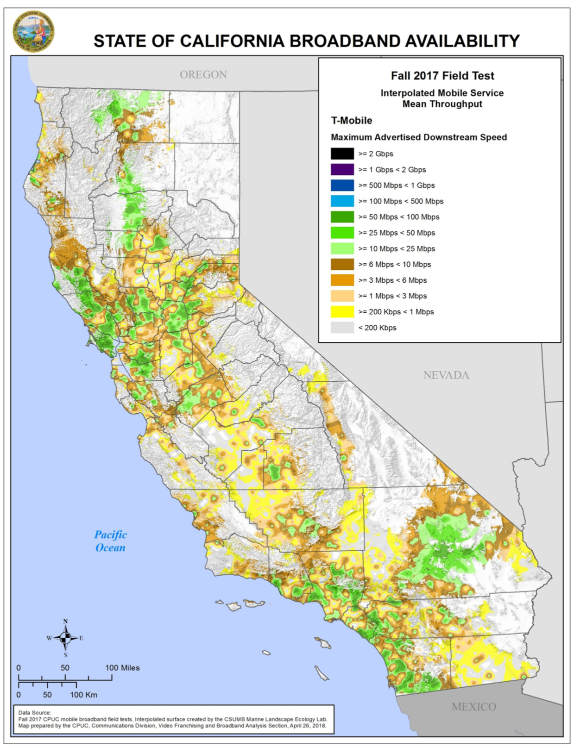
STATE OF CALIFORNIA BROADBAND AVAILABILITY
Fall 2017 Field Test
Interpolated Mobile Service Mean Throughput
T-Mobile
Maximum Advertised Downstream Speed
>= 2 Gbps
>= 1 Gbps < 2 Gbps
>= 500 Mbps < 1 Gbps
>= 100 Mbps < 500 Mbps
>= 50 Mbps < 100 Mbps
>= 25 Mbps < 50 Mbps
>= 10 Mbps < 25 Mbps
>= 6 Mbps < 10 Mbps
>= 3 Mbps < 6 Mbps
>= 1 Mbps < 3 Mbps
>= 200 Kbps < 1 Mbps
< 200 Kbps
OREGON
NEVADA
MEXICO
Pacific Ocean
N
W
E
S
0 50 100 Miles
0 50 100 Km
Data Source:
Fall 2017 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 Analysis Section, April 26, 2018.

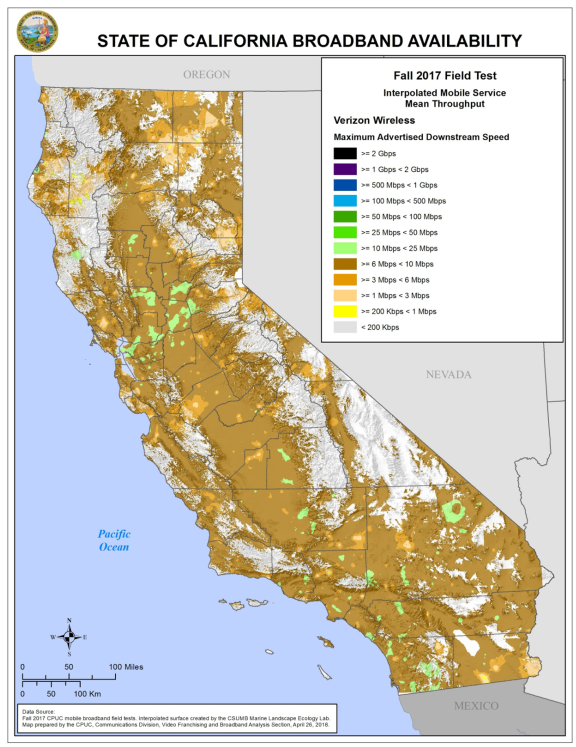
STATE OF CALIFORNIA BROADBAND AVAILABILITY
Fall 2017 Field Test
Interpolated Mobile Service Mean Throughput
Verizon Wireless
Maximum Advertised Downstream Speed
>= 2 Gbps
>= 1 Gbps < 2 Gbps
>= 500 Mbps < 1 Gbps
>= 100 Mbps < 500 Mbps
>= 50 Mbps < 100 Mbps
>= 25 Mbps < 50 Mbps
>= 10 Mbps < 25 Mbps
>= 6 Mbps < 10 Mbps
>= 3 Mbps < 6 Mbps
>= 1 Mbps < 3 Mbps
>= 200 Kbps < 1 Mbps
< 200 Kbps
OREGON
NEVADA
MEXICO
Pacific Ocean
N
W
E
S
0 50 100 Miles
0 50 100 Km
Data Source:
Fall 2017 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 Analysis Section, April 26, 2018.

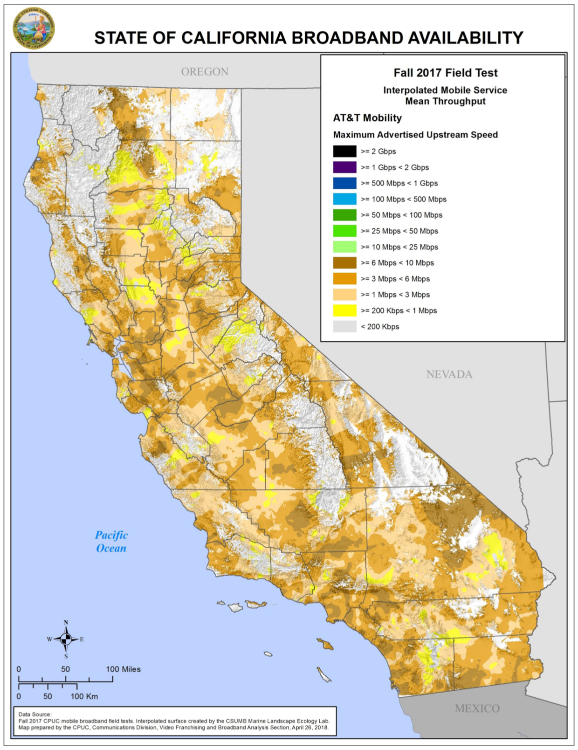
STATE OF CALIFORNIA BROADBAND AVAILABILITY
OREGON
Fall 2017 Field Test
Interpolated Mobile Service
Mean Throughput
AT&T Mobility
Maximum Advertised Upstream Speed
>= 2 Gbps
>= 1 Gbps < 2 Gbps
>= 500 Mbps < 1 Gbps
>= 100 Mbps < 500 Mbps
>= 50 Mbps < 100 Mbps
>= 25 Mbps < 50 Mbps
>= 10 Mbps < 25 Mbps
>= 6 Mbps < 10 Mbps
>= 3 Mbps < 6 Mbps
>= 1 Mbps < 3 Mbps
>= 200 Kbps < 1 Mbps
< 200 Kbps
NEVADA
Pacific Ocean
N
W
E
S
0 50 100 Miles
0 50 100 Km
MEXICO
Data Source:
Fall 2017 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 Analysis Section, April 26, 2018.

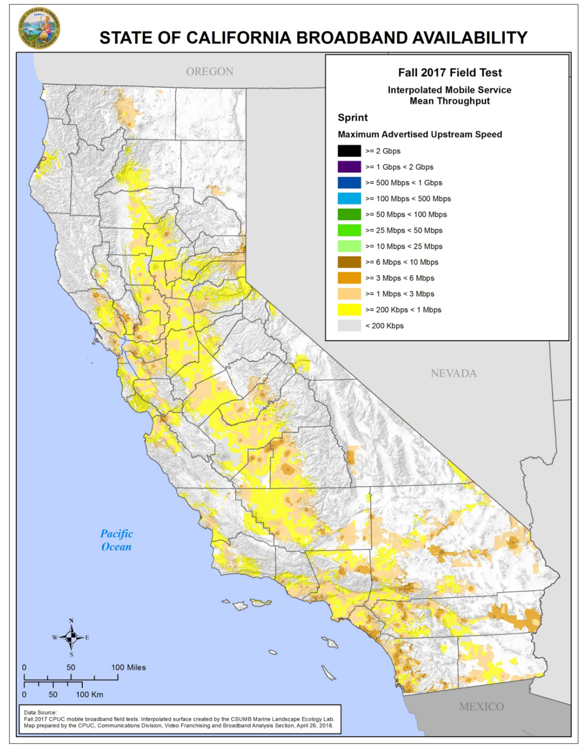
STATE OF CALIFORNIA BROADBAND AVAILABILITY
OREGON
Fall 2017 Field Test
Interpolated Mobile Service
Mean Throughput
Sprint
Maximum Advertised Upstream Speed
>= 2 Gbps
>= 1 Gbps < 2 Gbps
>= 500 Mbps < 1 Gbps
>= 100 Mbps < 500 Mbps
>= 50 Mbps < 100 Mbps
>= 25 Mbps < 50 Mbps
>= 10 Mbps < 25 Mbps
>= 6 Mbps < 10 Mbps
>= 3 Mbps < 6 Mbps
>= 1 Mbps < 3 Mbps
>= 200 Kbps < 1 Mbps
< 200 Kbps
NEVADA
Pacific Ocean
N
W
E
S
0 50 100 Miles
0 50 100 Km
MEXICO
Data Source:
Fall 2017 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 Analysis Section, April 26, 2018.

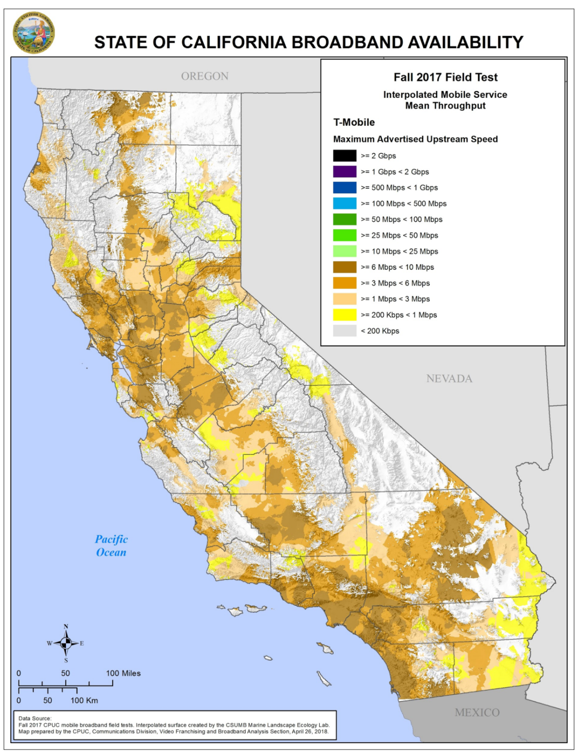
STATE OF CALIFORNIA BROADBAND AVAILABILITY
OREGON
Fall 2017 Field Test
Interpolated Mobile Service
Mean Throughput
T-Mobile
Maximum Advertised Upstream Speed
>= 2 Gbps
>= 1 Gbps < 2 Gbps
>= 500 Mbps < 1 Gbps
>= 100 Mbps < 500 Mbps
>= 50 Mbps < 100 Mbps
>= 25 Mbps < 50 Mbps
>= 10 Mbps < 25 Mbps
>= 6 Mbps < 10 Mbps
>= 3 Mbps < 6 Mbps
>= 1 Mbps < 3 Mbps
>= 200 Kbps < 1 Mbps
< 200 Kbps
NEVADA
Pacific Ocean
N
W
E
S
0 50 100 Miles
0 50 100 Km
MEXICO
Data Source:
Fall 2017 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 Analysis Section, April 26, 2018.

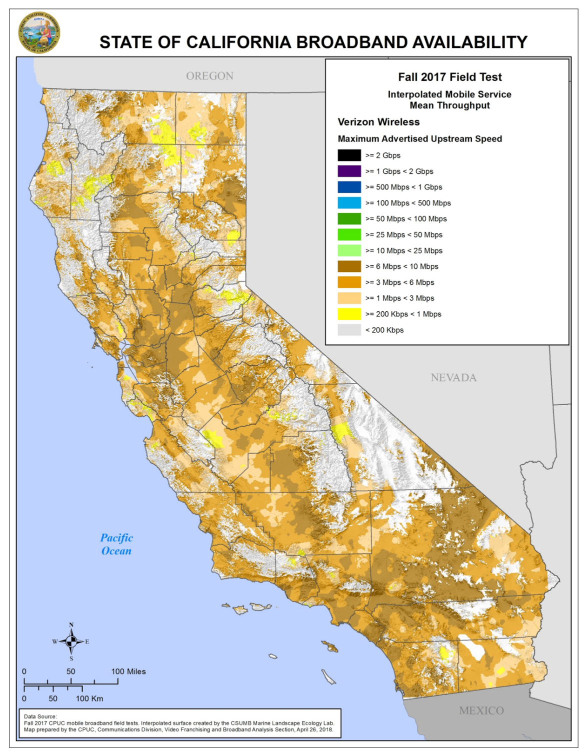
STATE OF CALIFORNIA BROADBAND AVAILABILITY
OREGON
Fall 2017 Field Test
Interpolated Mobile Service
Mean Throughput
Verizon Wireless
Maximum Advertised Upstream Speed
>= 2 Gbps
>= 1 Gbps < 2 Gbps
>= 500 Mbps < 1 Gbps
>= 100 Mbps < 500 Mbps
>= 50 Mbps < 100 Mbps
>= 25 Mbps < 50 Mbps
>= 10 Mbps < 25 Mbps
>= 6 Mbps < 10 Mbps
>= 3 Mbps < 6 Mbps
>= 1 Mbps < 3 Mbps
>= 200 Kbps < 1 Mbps
< 200 Kbps
NEVADA
Pacific Ocean
N
W
E
S
0 50 100 Miles
0 50 100 Km
MEXICO
Data Source:
Fall 2017 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 Analysis Section, April 26, 2018.

2.3 Throughput Variation Over Time

CalSPEED measures throughput performance every second during its measurements and can thus assess how throughput changes during a TCP communications session.   For each such session, CalSPEED computes the standard deviation of the throughput during the duration of the session as a metric of consistency of the throughput.

The following charts plot that standard deviation of TCP throughput, for each carrier, during the CalSPEED mobile measurement project. Note some conclusions:

- Normalized[11] mean TCP standard deviation, both upstream and downstream, improves from about 35% to about 25% from 2012 to 2017.
- AT&T and Verizon show much the same variation downstream and upstream.   T-mobile and particularly Sprint, show much more variation on the upstream service.
- Rural and urban show much the same downstream variation for all carriers.   But upstream service shows much more variation, for all carriers, for rural users.   A poorer, less predictable service by a factor of 50%.

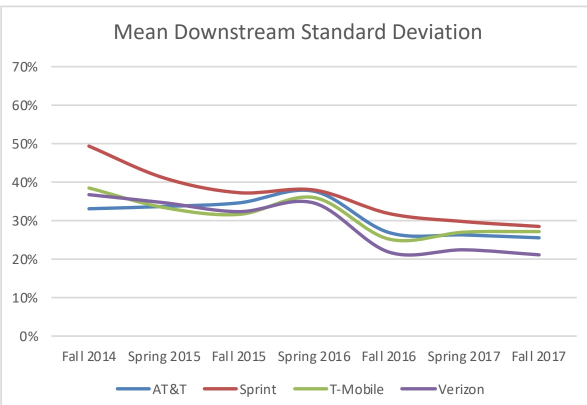

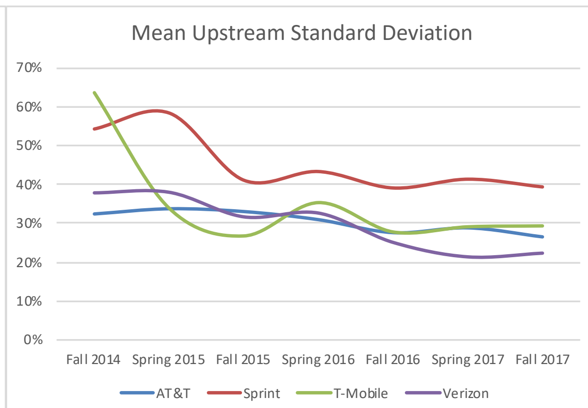

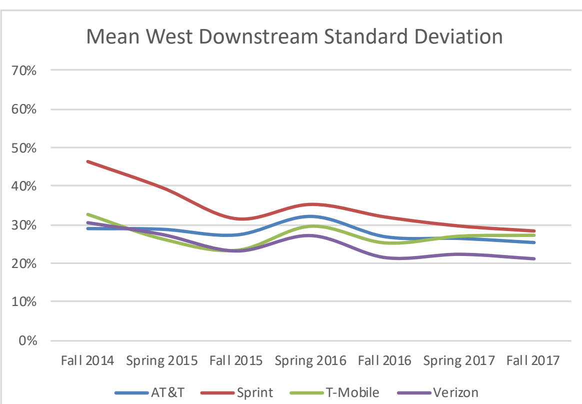

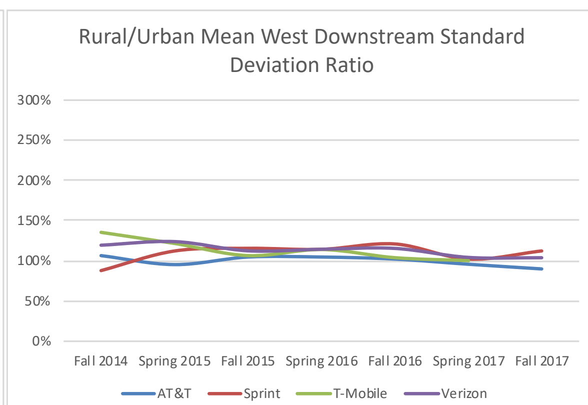

[11] Normalized as a percentage of the TCP throughput.   A measure of how much variation a user can see in TCP throughput.

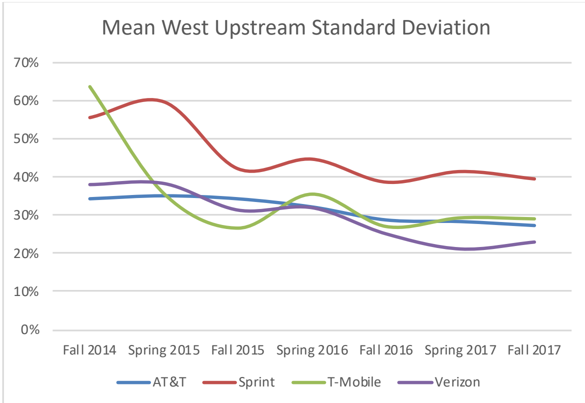
Mean West Upstream Standard Deviation
70%
60%
50%
40%
30%
20%
10%
0%
Fall 2014
Spring 2015
Fall 2015
Spring 2016
Fall 2016
Spring 2017
Fall 2017
AT&T
Sprint
T-Mobile
Verizon

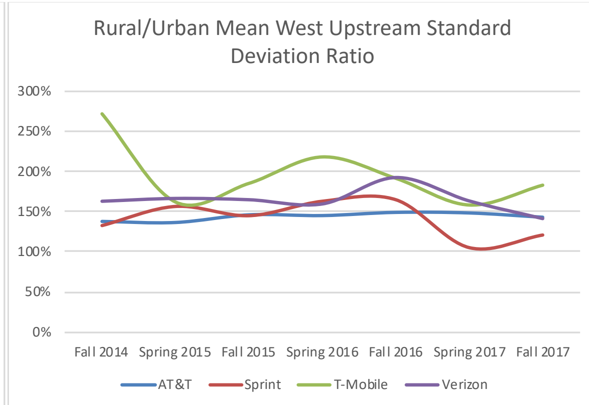
Rural/Urban Mean West Upstream Standard Deviation Ratio
300%
250%
200%
150%
100%
50%
0%
Fall 2014
Spring 2015
Fall 2015
Spring 2016
Fall 2016
Spring 2017
Fall 2017
AT&T
Sprint
T-Mobile
Verizon

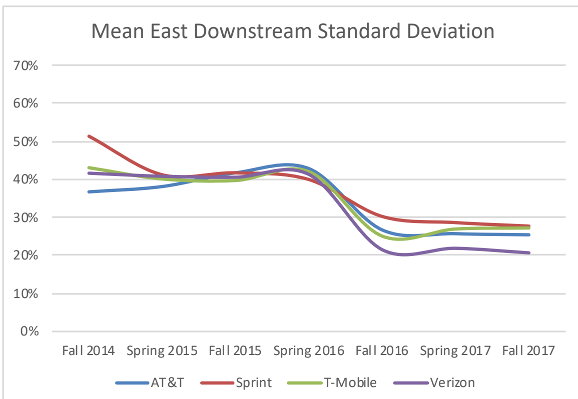
Mean East Downstream Standard Deviation
70%
60%
50%
40%
30%
20%
10%
0%
Fall 2014
Spring 2015
Fall 2015
Spring 2016
Fall 2016
Spring 2017
Fall 2017
AT&T
Sprint
T-Mobile
Verizon

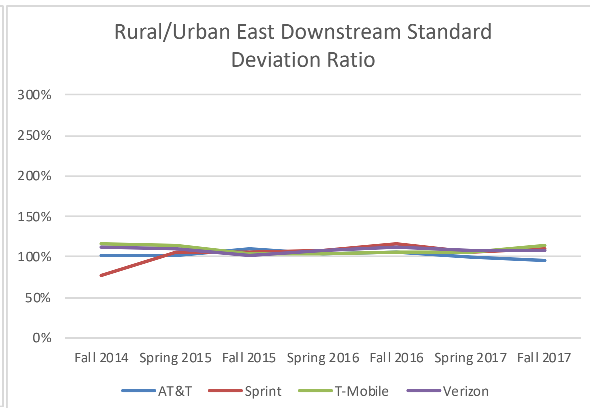
Rural/Urban East Downstream Standard Deviation Ratio
300%
250%
200%
150%
100%
50%
0%
Fall 2014
Spring 2015
Fall 2015
Spring 2016
Fall 2016
Spring 2017
Fall 2017
AT&T
Sprint
T-Mobile
Verizon

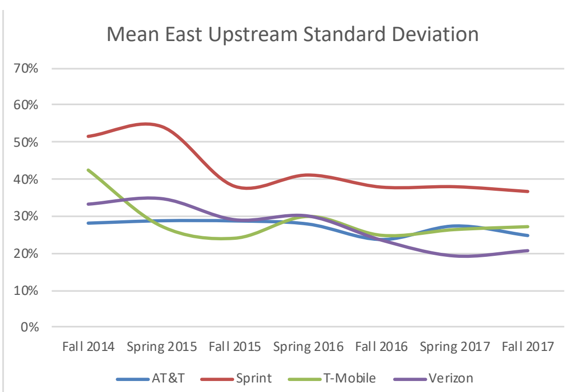
Mean East Upstream Standard Deviation
70%
60%
50%
40%
30%
20%
10%
0%
Fall 2014
Spring 2015
Fall 2015
Spring 2016
Fall 2016
Spring 2017
Fall 2017
AT&T
Sprint
T-Mobile
Verizon

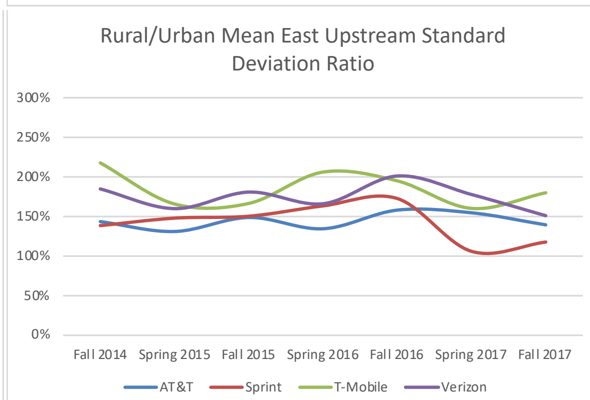
Rural/Urban Mean East Upstream Standard Deviation Ratio
300%
250%
200%
150%
100%
50%
0%
Fall 2014
Spring 2015
Fall 2015
Spring 2016
Fall 2016
Spring 2017
Fall 2017
AT&T
Sprint
T-Mobile
Verizon

2.4 Latency

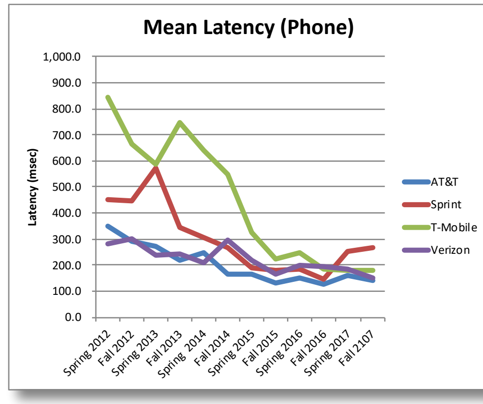

Latency is round trip packet delay experienced between the user and the destination server. Lower latency improves performance, particularly at very high speeds and for streaming media services such as over-the-top voice and video.

Mean latency has improved for both AT&T and Verizon from 300 msec in 2012 to under 150 msec at the end of 2017. T-Mobile and Sprint have a history of longer latency with dramatic improvements in the last three years but remained with consistency higher latency than the carriers.

Charted to the right, is the mean latency difference between access to the East server compared to the West server. At the end of 2017, there is about 50 msec of latency difference between West and East from a baseline of about 65 msec for urban AT&T and Verizon users. Rural users can see an additional 20-30% latency penalty.

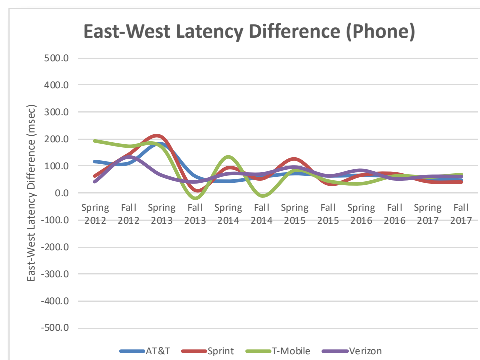

Charted above, each carrier consistently delivers lower latency for urban users than for rural or tribal users. This is particularly true for T-Mobile. Tribal and rural users have a similar pattern.

Latency maps for each carrier identify the spatial variation across the state. Of particular note are the dark blue areas of each map, identifying extraordinary long mean latency - approaching a full half a second. These same areas will have poor voice and video streaming service.

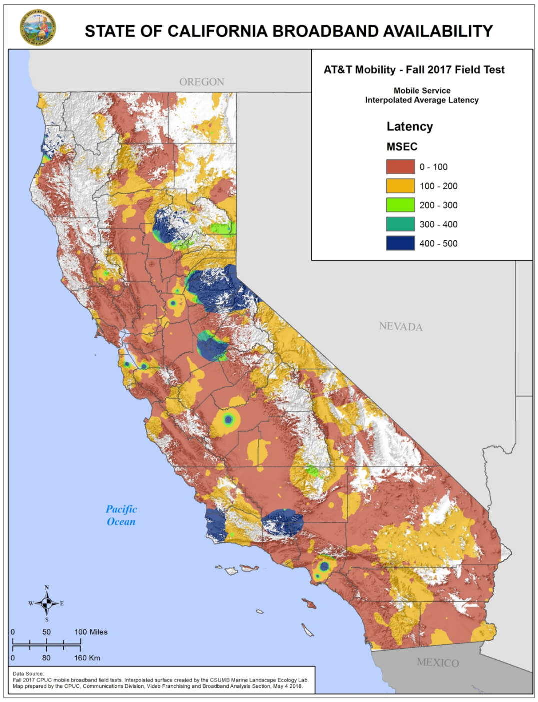

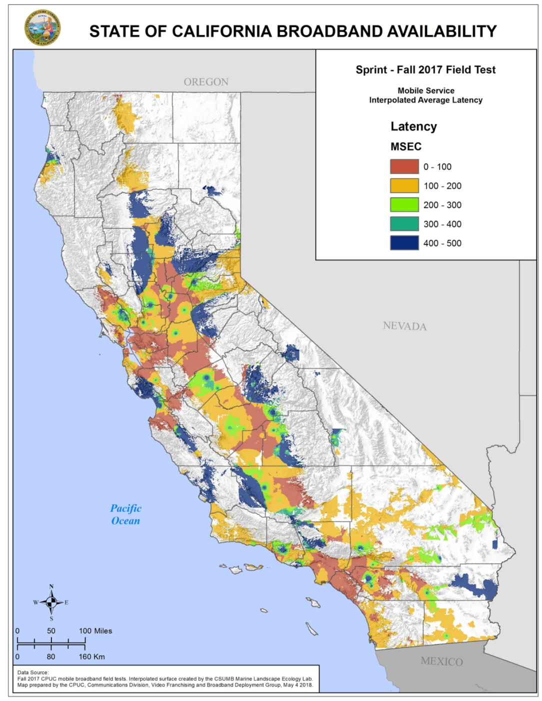

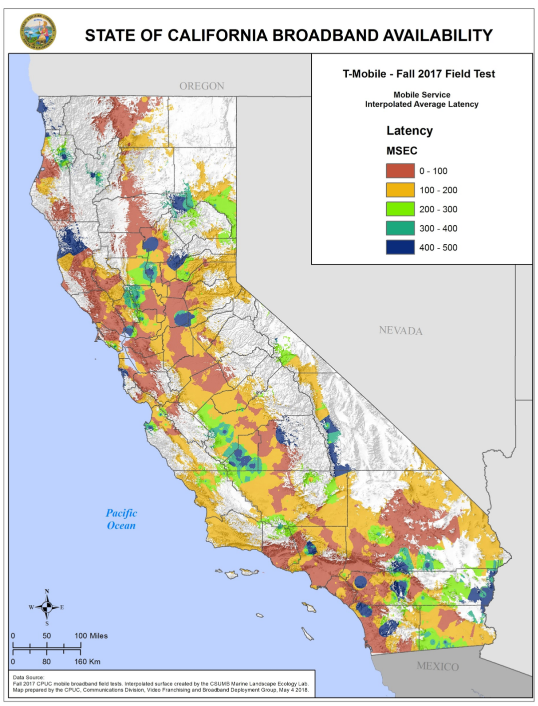

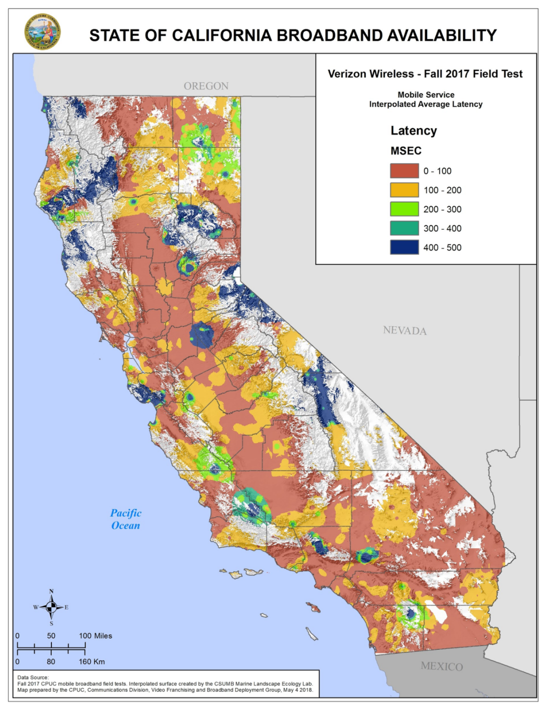

2.5 Jitter

Jitter is the variation in latency as packet networks struggle for consistent latency. Mean jitter[12] improved

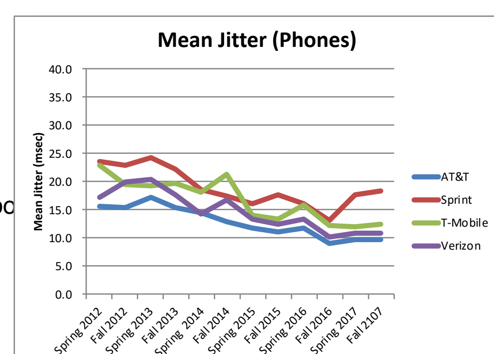

[12] Jitter is the measure of variation in latency. It is important compo audio and video.

for all carriers from 2012 thru 2017.   AT&T and Verizon lead with both the lowest jitter and well as the most consistent improvement.   Both T-Mobile and Sprint improved, but with far less consistency and remained with absolutely more jitter than the leading two carriers.

Jitter was similar between East and West for all carriers but with spikes of change in both 2014 and 2016. Dramatic 20% increase in East jitter in 2014 and a 10% decrease in East jitter, relative to West, in 2014.

Jitter of urban users is always less than rural and tribal users for all carriers.

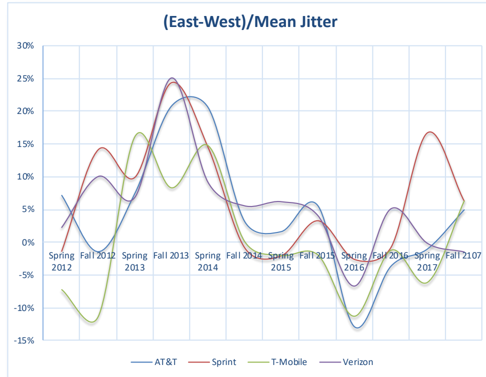

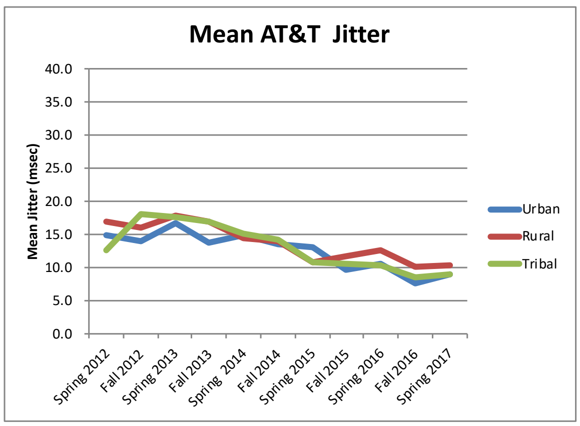

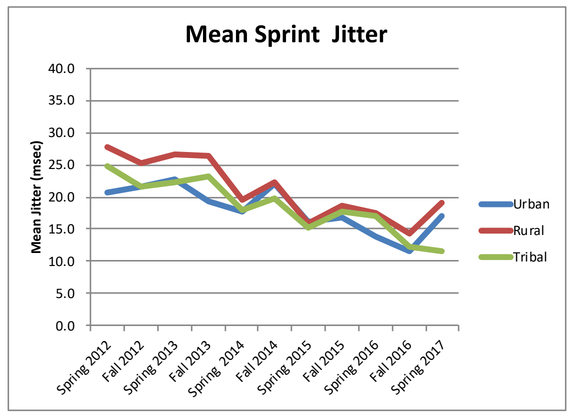

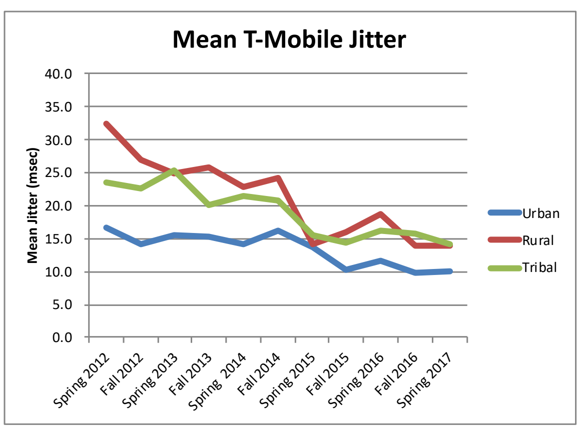

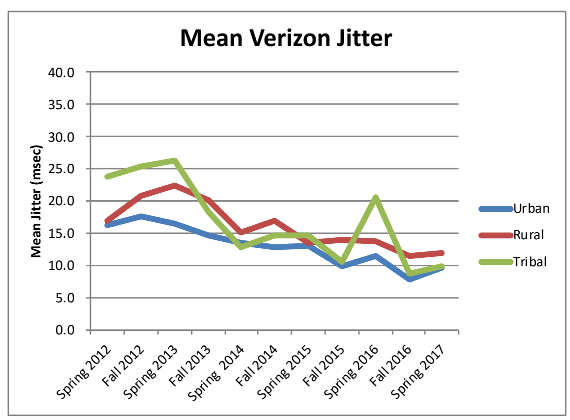

2.6 Packet Loss

Mean packet loss[13] rates, for all carriers, show high volatility from year to year with a trend to increase from 2012 thru the end of 2017

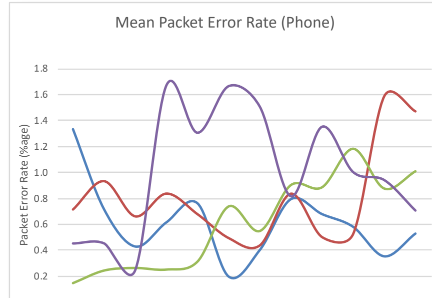

On average, rural users have about 50% more packet loss than urban users for all carriers.

Tribal packet loss has much more volatility than either urban or rural, particularly since 2015.

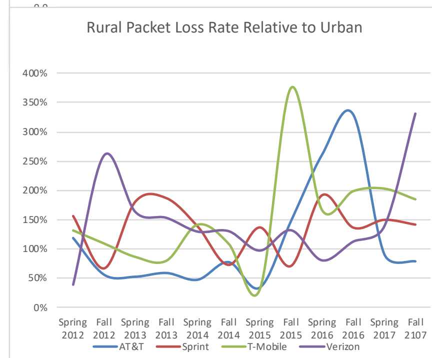

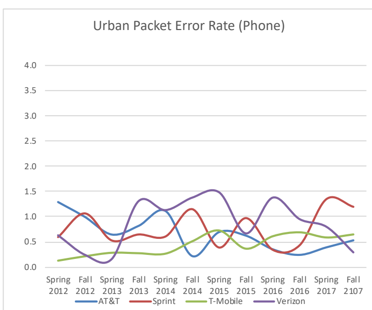

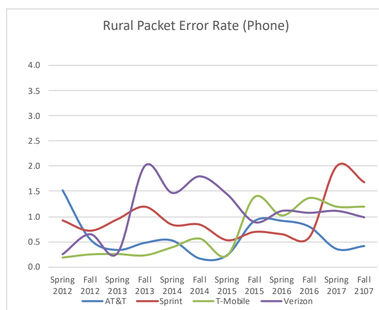

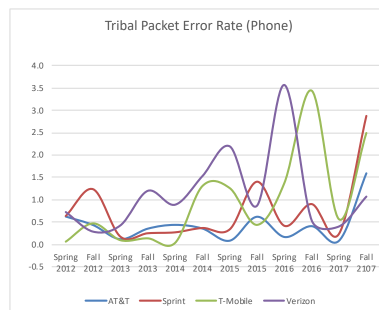

[13] Mean packet loss is the average percentage of packets that are lost during transmission. Small increases in packet loss are particularly degrading for streaming media services.

2.7 TCP Connection Quality

TCP connection quality[14], as measured by the ratio of failed to attempted TCP connection attempts[15], is a metric that assesses the reliability of the broadband service. While overall, about 15% of all mobile broadband TCP connection attempts fail, this metric is very different for urban versus rural and tribal users. Urban users consistently see less than 10% of connection attempts fail and the leading carriers, AT&T and Verizon, offer a less than 5% failure rate. Rural (and tribal users) see this failure rate increase more than 5 times than for urban users.

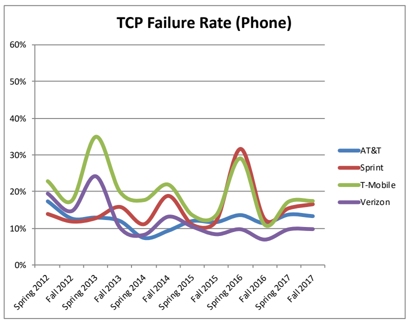

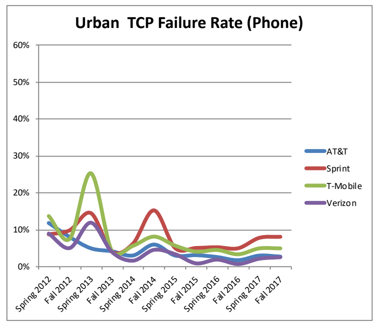

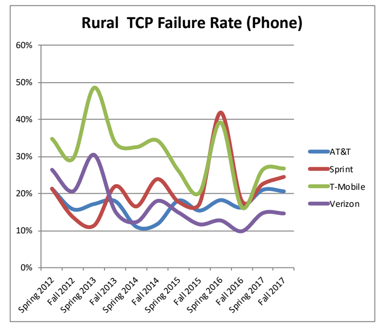

Charting the relative TCP failure rate between rural and urban users reveals that not only do rural users see a more TCP failures than urban users but also that that failure rate is differentially increasing.

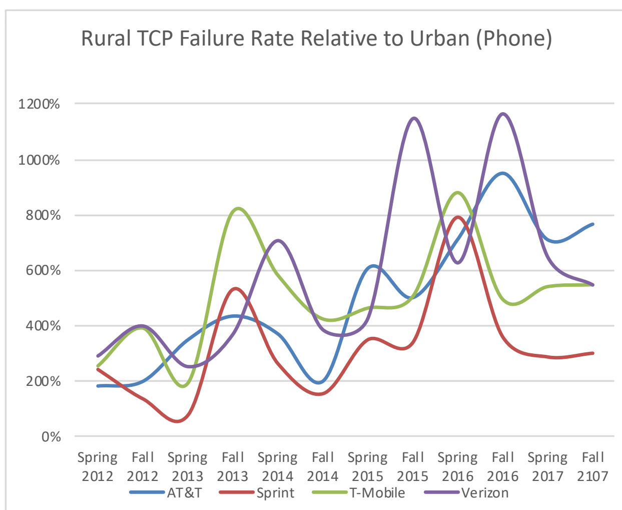

This is a dramatic result. For many users, this TCP connection failure is masked by network applications automatically retrying a failed connection, but it is an underlying reliability failure of the networks.

[14] 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.

[15] 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.

2.8 Over The Top Voice Quality

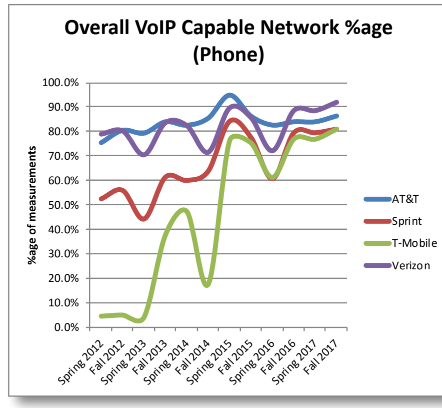

Over the top digital voice over IP is increasingly the standard for audio conversations in the modern Internet. CalSPEED calculates a simulated Mean Opinion Score to assess OTTP Voice over IP (VoIP) service. This calculation is heavily dependent on measured latency, jitter and packet loss in each location for each carrier. A computed MOS of greater than or equal to 4.0 suggests a quality voice service.

AT&T and Verizon have been consistent leaders with over 85% of CalSPEED measurement locations offering a quality voice service. Both Sprint and T-Mobile have offered a voice service of lower quality than the leaders and have only begun to catch up in breadth of voice service since 2016.

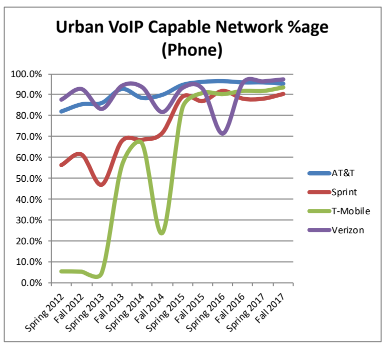

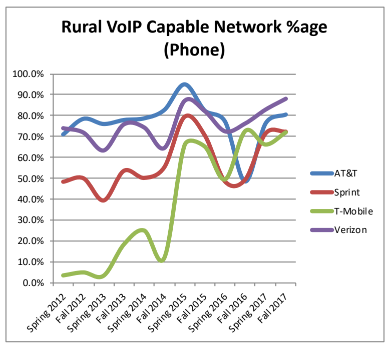

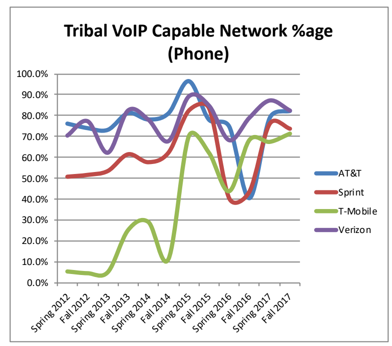

All carriers are challenged to deliver broad VoIP coverage for rural and tribal users. The charts above illustrate that all carriers offer over 90% VoIP coverage for urban users while Spring and T-Mobile struggle to deliver 70% coverage for rural and tribal users. In general, rural users get about 85% of the VoIP coverage of urban users for all carriers.

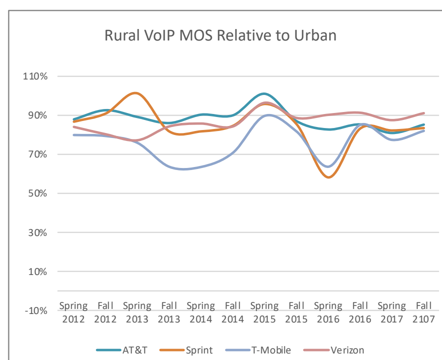

CalSPEED maps VoIP MOS across the state for each carrier. AT&T and Verizon both have extensive VoIP service across the state, with Sprint and T-Mobile are limited to dense urban areas and major transportation routes. In the following maps, note the state locations not red for locations with poor or non-existent VoIP service.

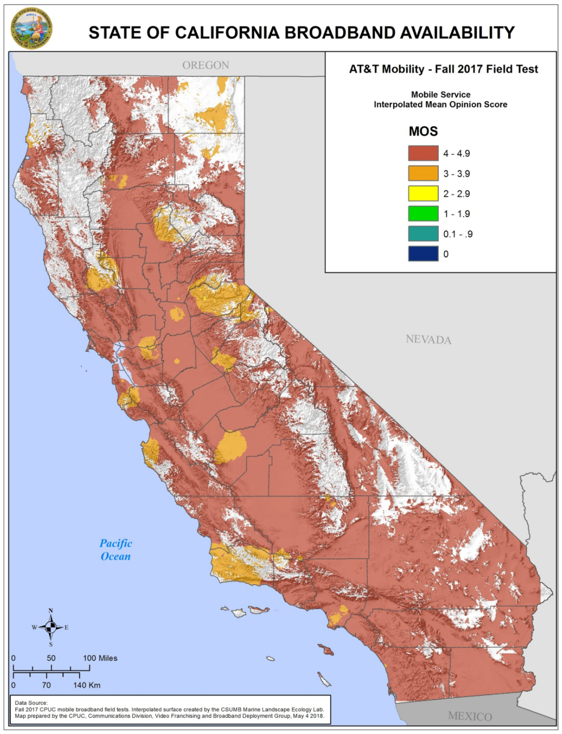

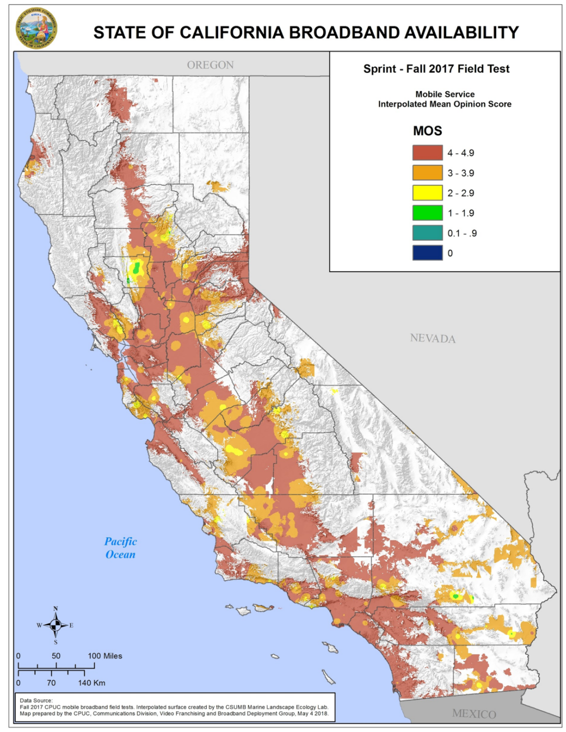

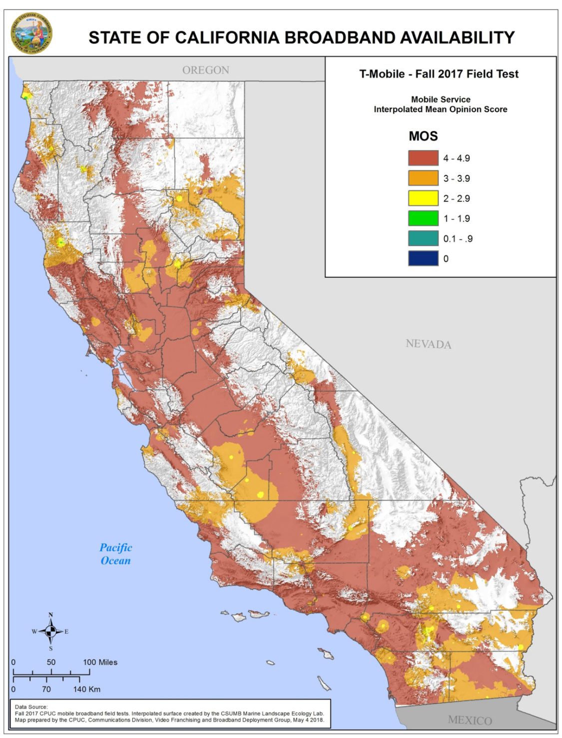

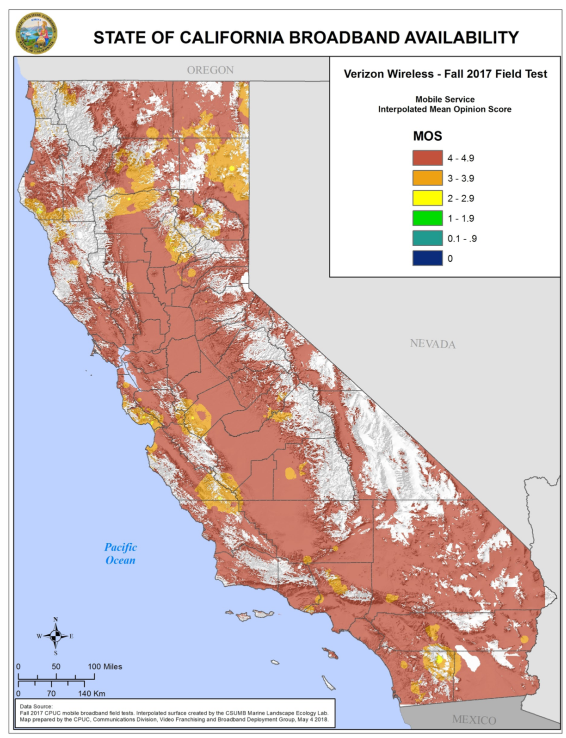

Section 9 examines two ways we can use this mapping information, integrated with other mapping data, to make predictions about emergency management performance.

2.9 Over The Top Video

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

2.9.1 Video Streaming

Video streaming is the service of viewing a video, in real-time, stored on an Internet server that is streamed in real-time to the user. 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 Fall 2014 measurement and this report concludes 3.5 years of measurement.

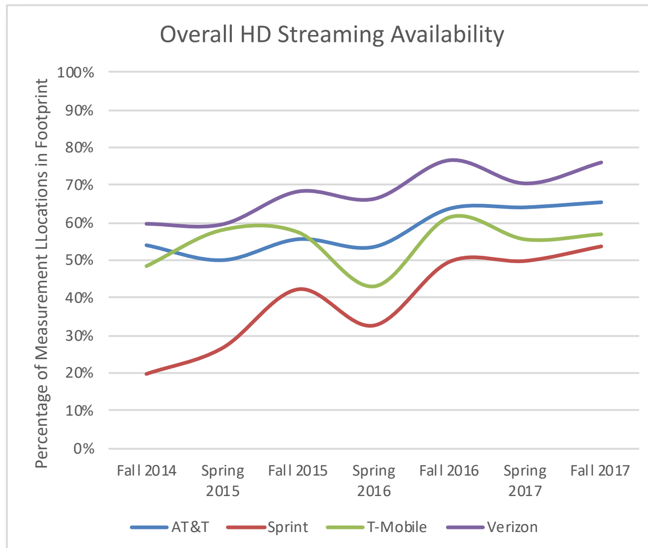

Assessing coverage by the percentage of measurement locations with a standard of service, some trends can be easily seen.

- All carriers exhibited a trend to increasing coverage of HD quality streaming mobile video. This was observed for all demographics.
- Verizon has the most consistent streaming video service, followed by AT&T. T-Mobile has a strong stream service but limited to only urban users.
- Rural and tribal HD streaming video lags urban for all carriers during the entire measurement period. Verizon offers the best availability for rural video streaming.

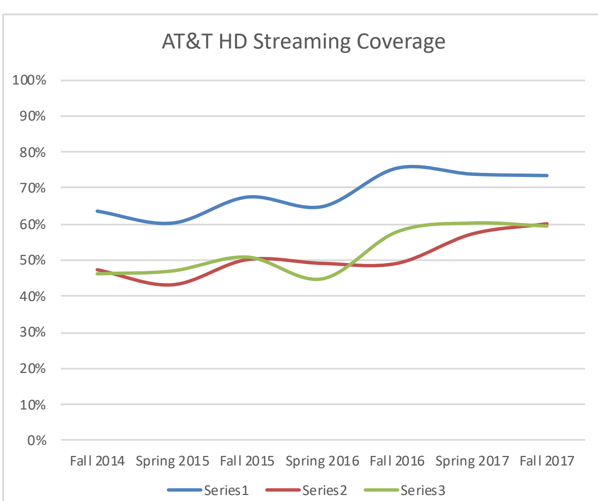

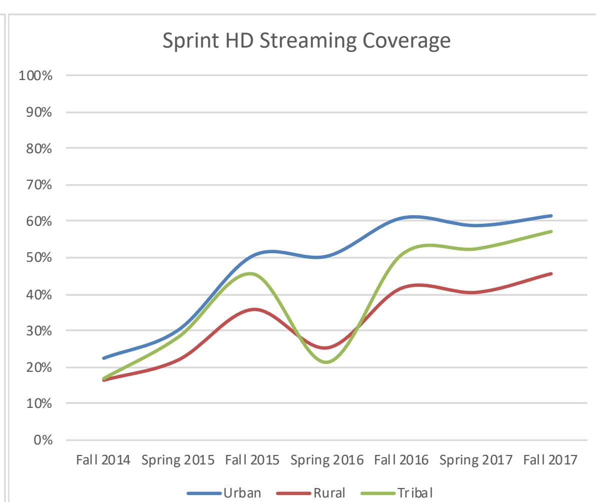

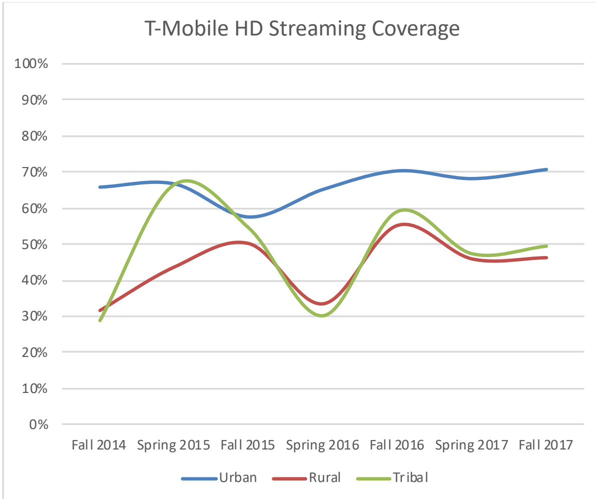

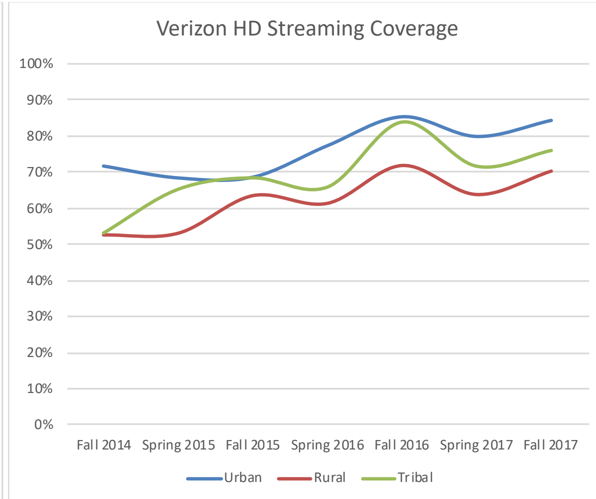

Rural and tribal HD coverage, for all carriers, is consistently degraded from urban users. For AT&T and Verizon, rural users see about 80% of the availability of this quality of service. For Sprint and Verizon, rural users see about 70% of the availability.

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 measurement rounds

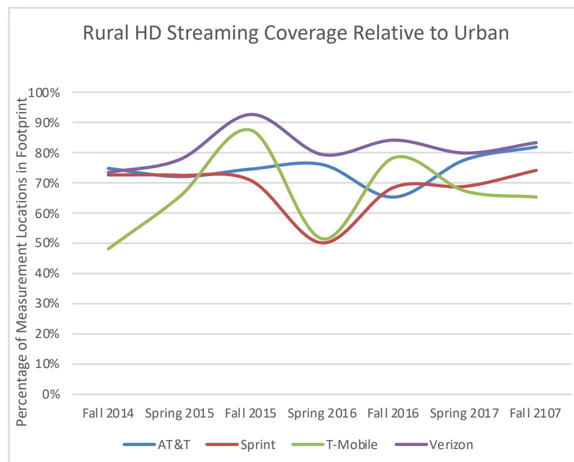

Mobile video streaming coverage is illustrated in the following coverage maps estimating mobile video service across the state for the four carriers. HD streaming coverage for all carriers is localized with Verizon giving the best overall state-wide coverage.

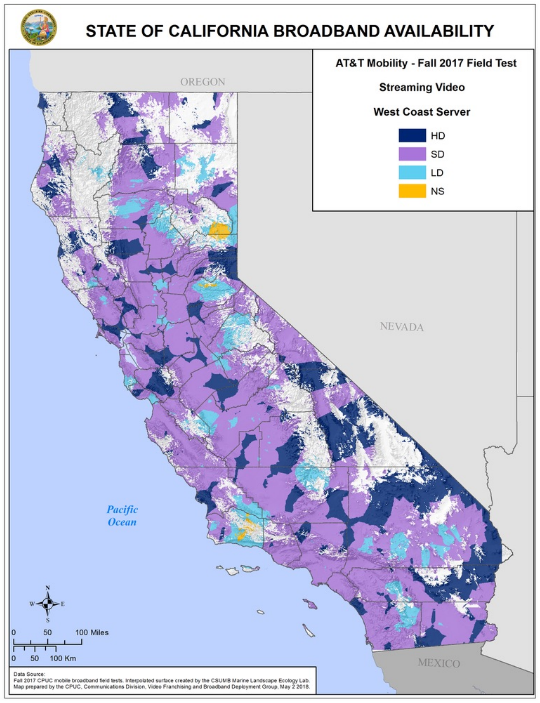
STATE OF CALIFORNIA BROADBAND AVAILABILITY
AT&T Mobility - Fall 2017 Field Test
Streaming Video
West Coast Server
HD
SD
LD
NS
OREGON
NEVADA
Pacific Ocean
MEXICO
N
W
E
S
0 50 100 Miles
0 50 100 Km
Data Source:
Fall 2017 CPUC mobile broadband field tests. Interpolated surface created by the CSUMB Marine Landscape Ecology Lab.
Map prepared by the CPUC, Communications Division, Video Franchising and Broadband Deployment Group, May 2 2018.

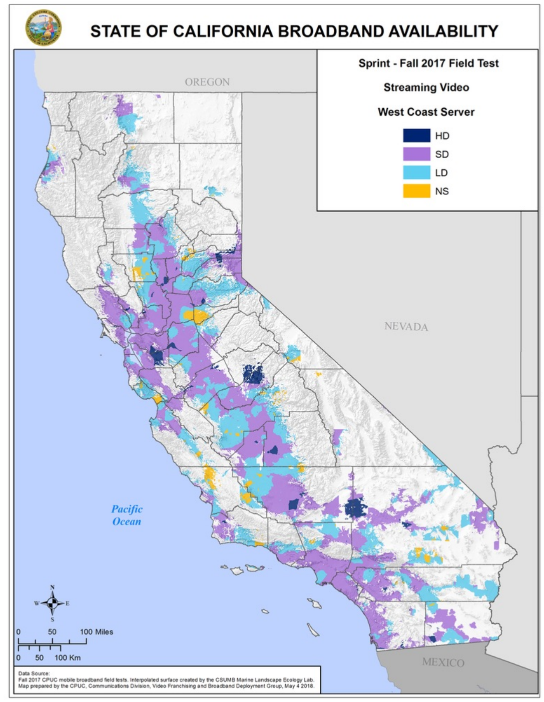
STATE OF CALIFORNIA BROADBAND AVAILABILITY
Sprint - Fall 2017 Field Test
Streaming Video
West Coast Server
HD
SD
LD
NS
OREGON
NEVADA
Pacific Ocean
MEXICO
0 50 100 Miles
0 50 100 Km
Data Source:
Fall 2017 CPUC mobile broadband field tests. Interpolated surface created by the CSUMB Marine Landscape Ecology Lab.
Map prepared by the CPUC, Communications Division, Video Franchising and Broadband Deployment Group, May 4 2018.

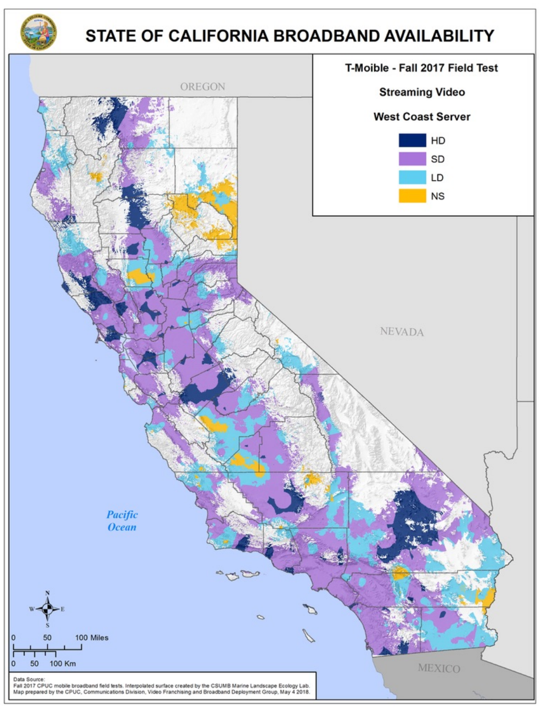
STATE OF CALIFORNIA BROADBAND AVAILABILITY
T-Moible - Fall 2017 Field Test
Streaming Video
West Coast Server
HD
SD
LD
NS
OREGON
NEVADA
Pacific Ocean
MEXICO
0 50 100 Miles
0 50 100 Km
Data Source:
Fall 2017 CPUC mobile broadband field tests. Interpolated surface created by the CSUMB Marine Landscape Ecology Lab.
Map prepared by the CPUC, Communications Division, Video Franchising and Broadband Deployment Group, May 4 2018.

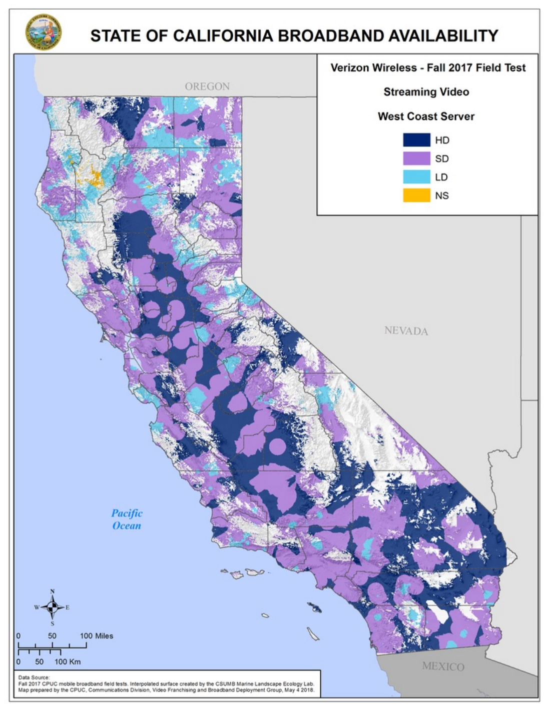
STATE OF CALIFORNIA BROADBAND AVAILABILITY
Verizon Wireless - Fall 2017 Field Test
Streaming Video
West Coast Server
HD
SD
LD
NS
OREGON
NEVADA
Pacific Ocean
MEXICO
0 50 100 Miles
0 50 100 Km
Data Source:
Fall 2017 CPUC mobile broadband field tests. Interpolated surface created by the CSUMB Marine Landscape Ecology Lab.
Map prepared by the CPUC, Communications Division, Video Franchising and Broadband Deployment Group, May 4 2018.

2.9.2 Conference Video

Conference video is a two-way audio/video stream between two users. CalSPEED uses two-way MOS streaming audio quality and 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.

For all carriers, HD conference video coverage has improved from 2012-2017. Modest improvements for conferencing to the West with dramatic improvements beginning in 2013 for conferencing to the East.

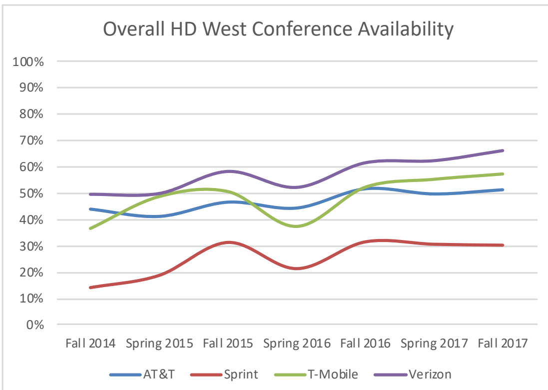

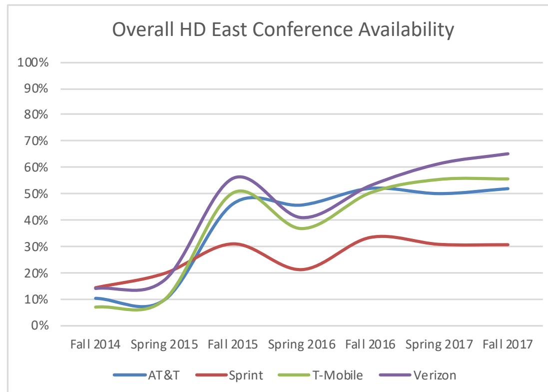

The video quality chart on the right illustrates that for all carriers and all geographies, conference video has modestly improved to the East server since Spring 2016. Particular improvements have been for Sprint and T-Mobile for rural users.

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).

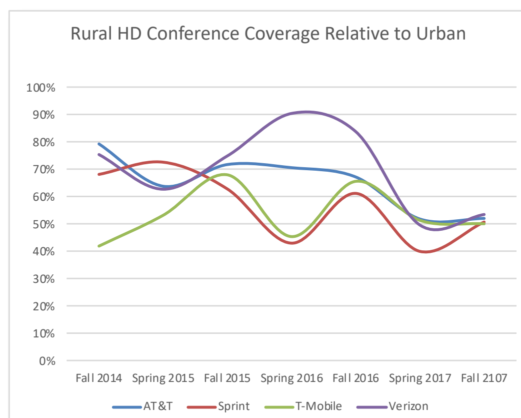

While video conference quality overall showed improvement, conference quality for rural users degraded relative to urban users in the same period. Overall, CalSPEED assessed a rural users had 70% of the conference availability of urban users in 2012, but that degraded to 50% at the end 2017.

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)

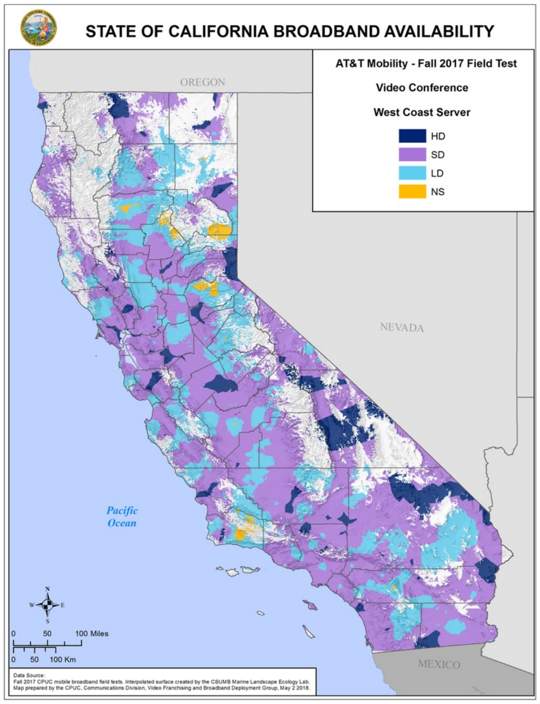
STATE OF CALIFORNIA BROADBAND AVAILABILITY
AT&T Mobility - Fall 2017 Field Test
Video Conference
West Coast Server
HD
SD
LD
NS
OREGON
NEVADA
Pacific Ocean
MEXICO
N
W
E
S
0 50 100 Miles
0 50 100 Km
Data Source:
Fall 2017 CPUC mobile broadband field tests. Interpolated surface created by the CSUMB Marine Landscape Ecology Lab.
Map prepared by the CPUC, Communications Division, Video Franchising and Broadband Deployment Group, May 2 2018.

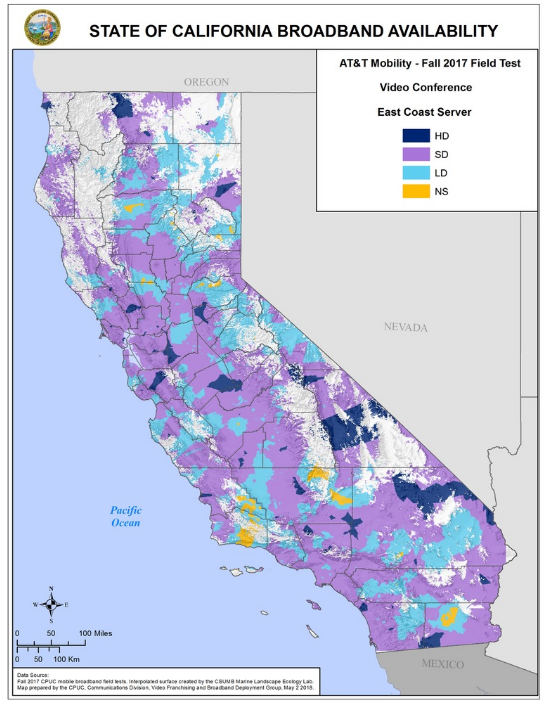
STATE OF CALIFORNIA BROADBAND AVAILABILITY
AT&T Mobility - Fall 2017 Field Test
Video Conference
East Coast Server
HD
SD
LD
NS
OREGON
NEVADA
Pacific Ocean
MEXICO
N
W
E
S
0 50 100 Miles
0 50 100 Km
Data Source:
Fall 2017 CPUC mobile broadband field tests. Interpolated surface created by the CSUMB Marine Landscape Ecology Lab.
Map prepared by the CPUC, Communications Division, Video Franchising and Broadband Deployment Group, May 2 2018.

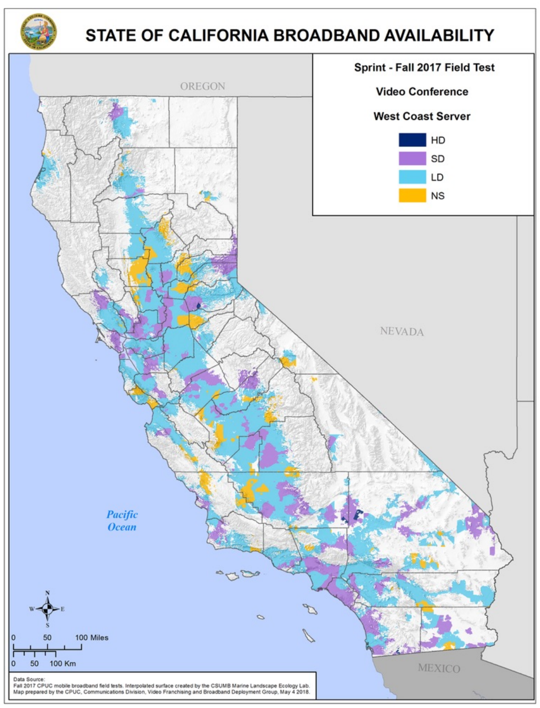
STATE OF CALIFORNIA BROADBAND AVAILABILITY
Sprint - Fall 2017 Field Test
Video Conference
West Coast Server
HD
SD
LD
NS
OREGON
NEVADA
Pacific Ocean
MEXICO
N
W
E
S
0 50 100 Miles
0 50 100 Km
Data Source:
Fall 2017 CPUC mobile broadband field tests. Interpolated surface created by the CSUMB Marine Landscape Ecology Lab.
Map prepared by the CPUC, Communications Division, Video Franchising and Broadband Deployment Group, May 4 2018.

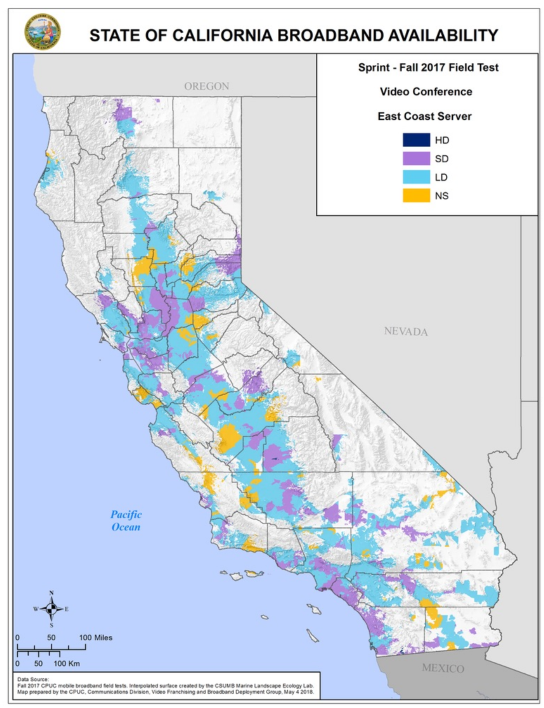
STATE OF CALIFORNIA BROADBAND AVAILABILITY
Sprint - Fall 2017 Field Test
Video Conference
East Coast Server
HD
SD
LD
NS
OREGON
NEVADA
Pacific Ocean
MEXICO
N
W
E
S
0 50 100 Miles
0 50 100 Km
Data Source:
Fall 2017 CPUC mobile broadband field tests. Interpolated surface created by the CSUMB Marine Landscape Ecology Lab.
Map prepared by the CPUC, Communications Division, Video Franchising and Broadband Deployment Group, May 4 2018.

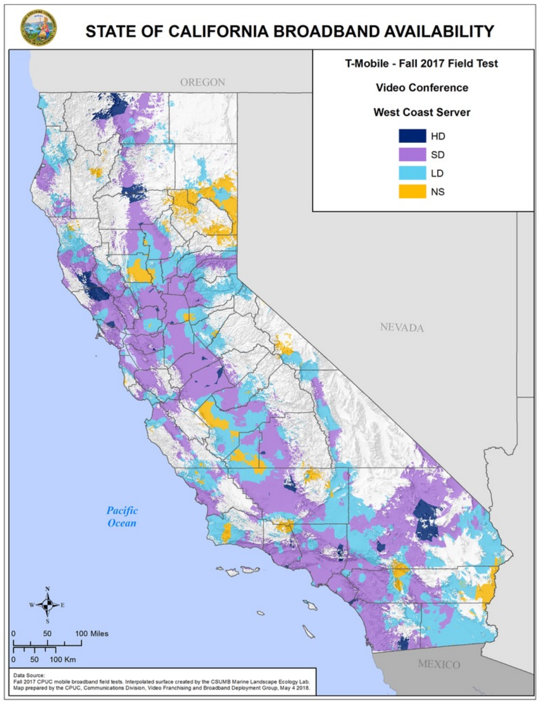

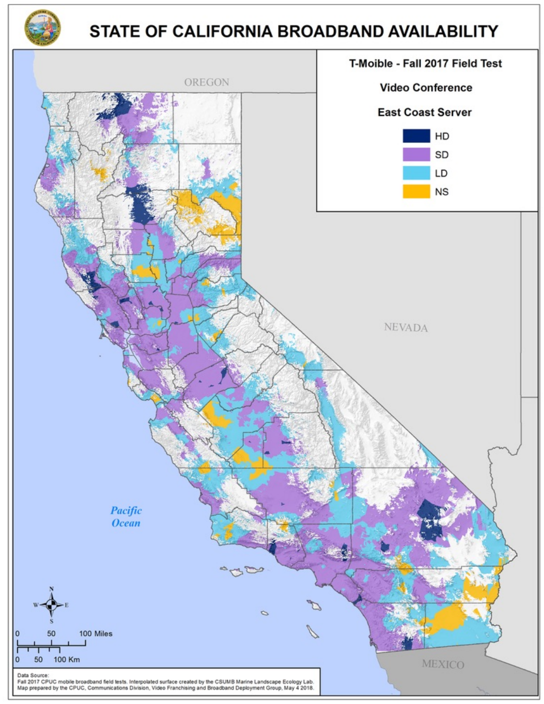

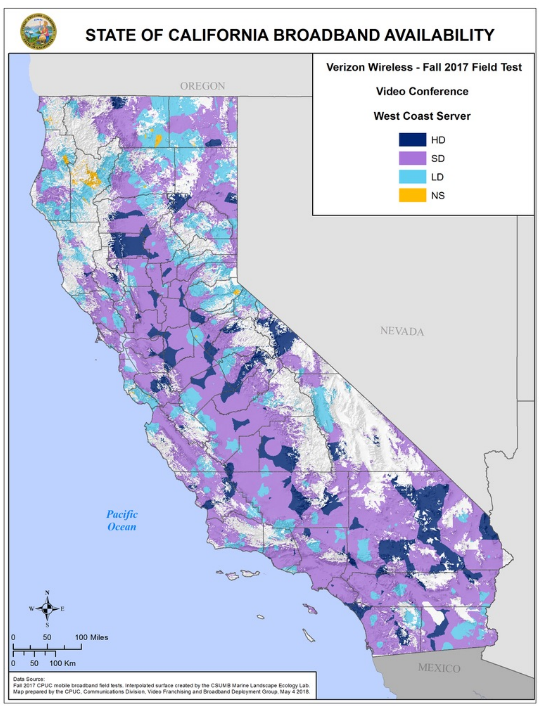

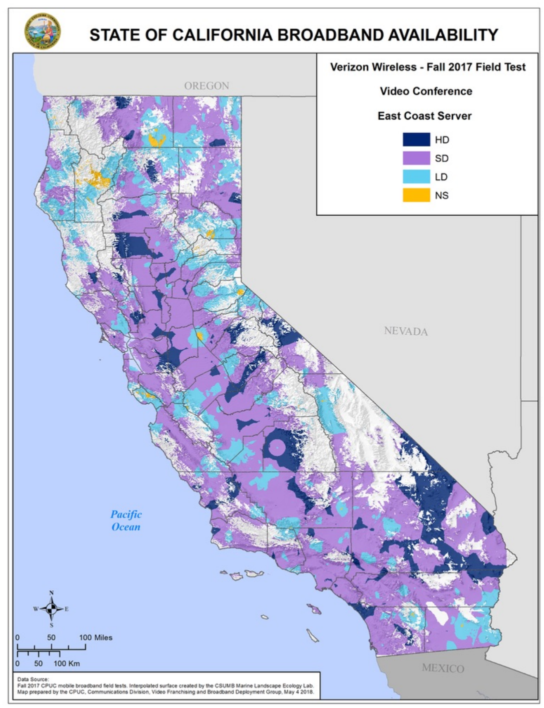

The Internet backbone, thru increased latency, can decrease the availability of HD quality conference video for all carriers. In measurements since Fall 2015, HD conference availability is about equal between conference destinations - West or East.

Interactive video service quality is mapped across the state for all four carriers in the above charts. With the increased East availability beginning in Fall 2015, service quality East vs West for all carriers is similar.

3. Persistent Rural/Urban Mobile Digital Divide

The six years of CalSPEED mobile broadband measurement documents 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. The coming technology change to 5G will likely only further degrade rural and tribal relative to urban since both 5G technology and deployment economics bias towards urban deployment.

Let's examine each of these metrics.

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

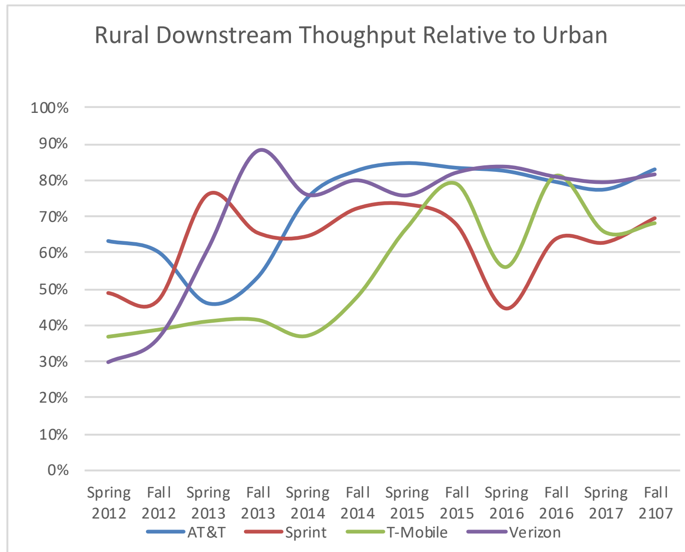

Rural mean upstream TCP throughput is consistently only 60-70% of urban.

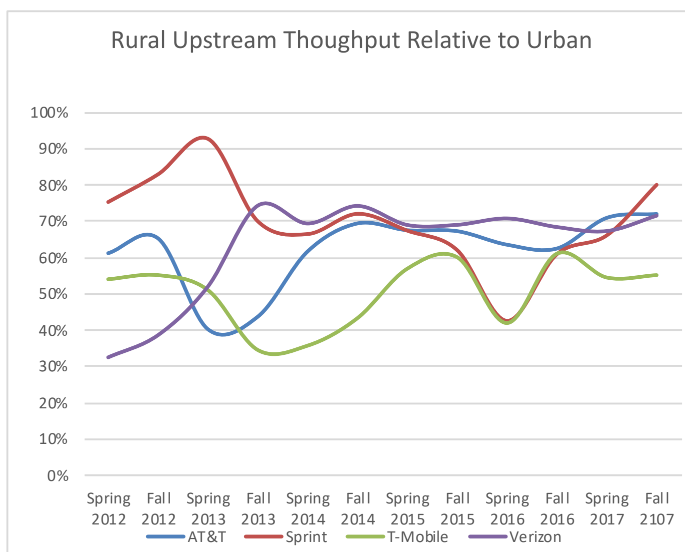

Rural mean latency is consistently at least 50% (T-Mobile and Sprint) to 100% (AT&T and Verizon) worse than urban.

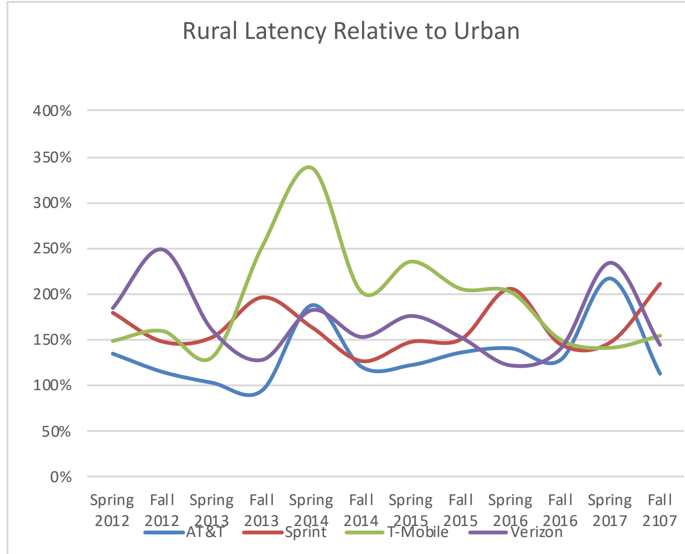

Rural mean jitter is consistently at least 20% worse than urban.

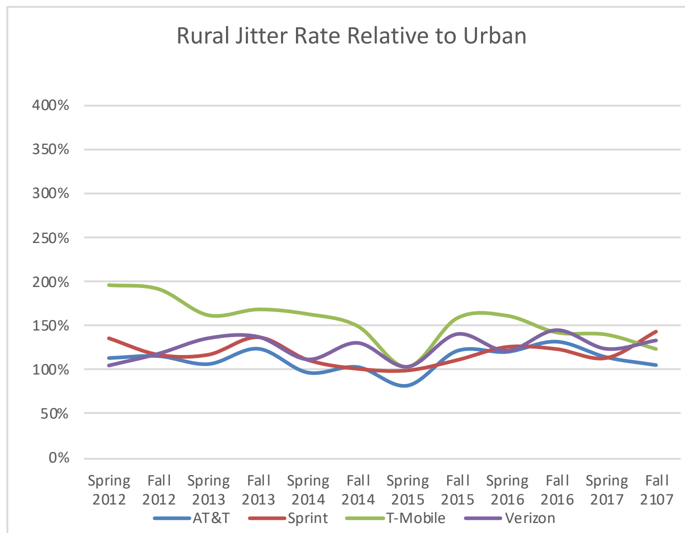

Rural packet loss is consistently 100-200% worse than urban. There is arguably a modest trend for worsening packet loss.

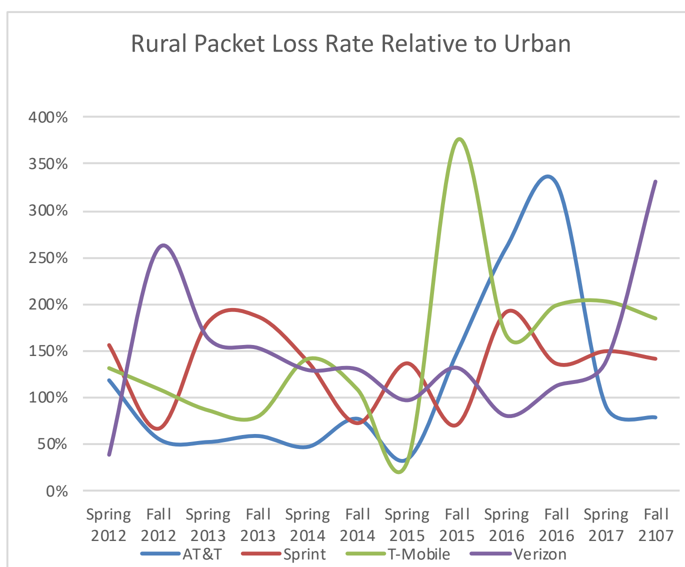

The poor performance of rural latency, jitter and packet loss strongly implies that real-time streaming voice and video services will perform materially worse in rural than urban areas.

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 of all TCP connection attempts fail.

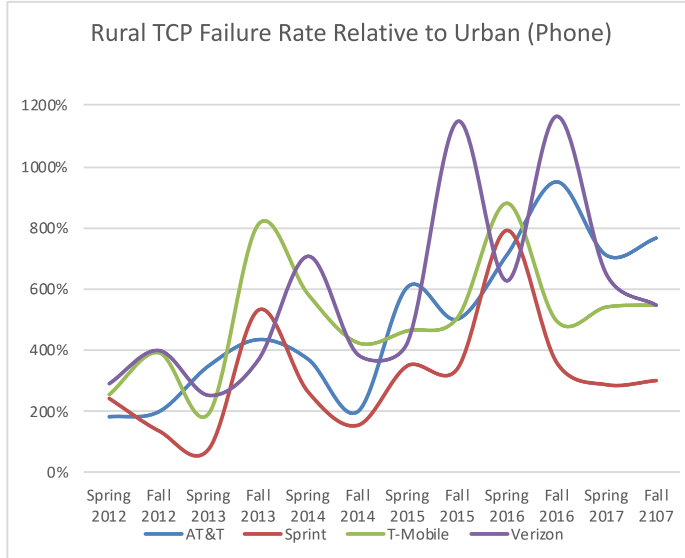

Rural OTT digital voice coverage is consistently only 85% of urban.

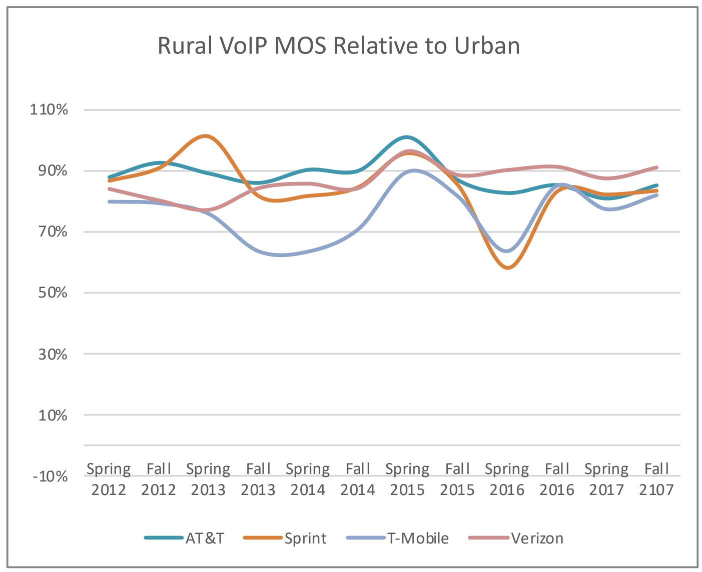

Rural HD streaming video coverage is consistently 60-80% of urban. Rural users will encounter no HD streaming availability 2.5x more frequently than urban users.

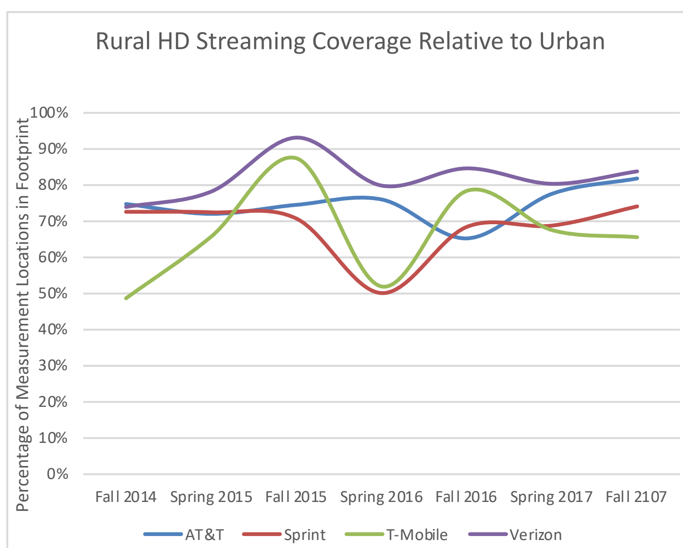

Rural HD interactive video coverage is consistently 60-80% of urban with a significant trend towards worsening. Rural users will encounter no HD conference availability 1.5x more frequently than urban users.

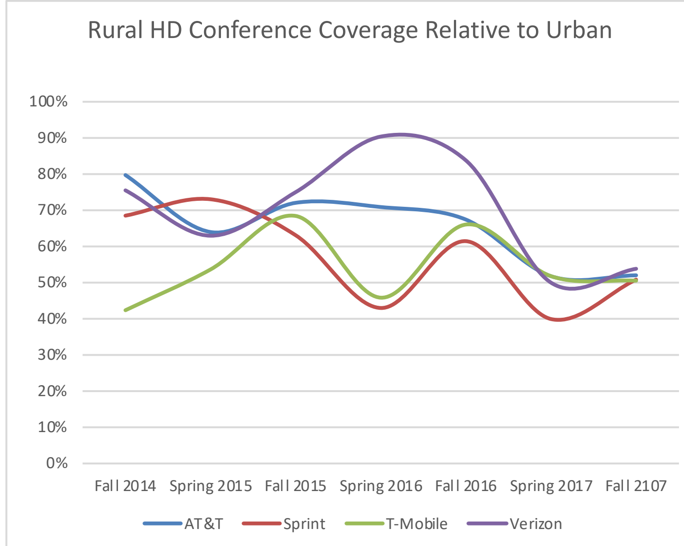

Rural mobile broadband coverage at the standard of 25 Mb/s down and 3 Mb/s up is consistently 60-80% of urban.

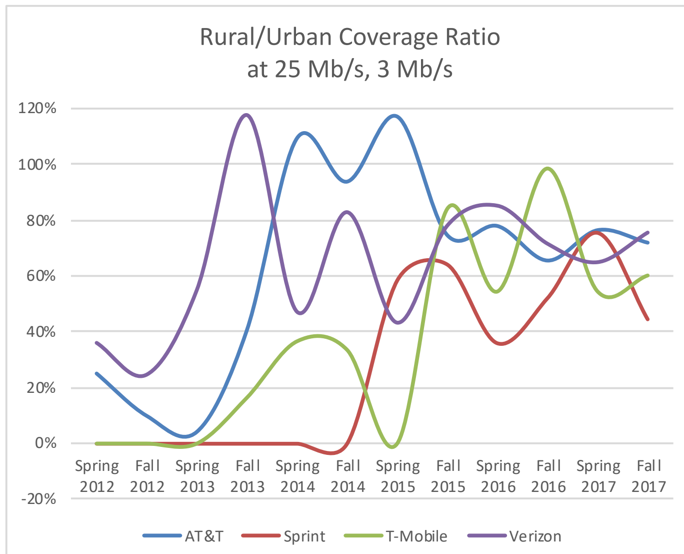

Rural LTE coverage is consistently less than 80% of urban (with the significant exception of Verizon). This has substantial implications for the use of mobile broadband as a replacement for wired broadband in rural areas and for the capability of public safety services.

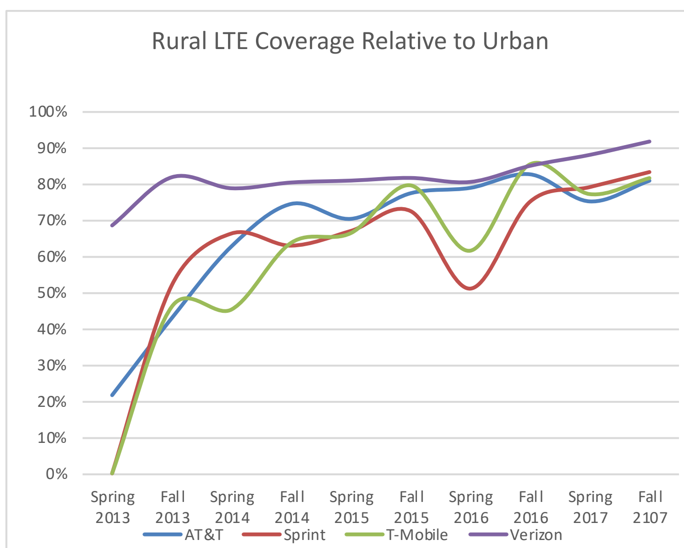

Rural mobile broadband consistently has more than a 5x use of legacy 1G and 2G mobile broadband technology than urban. It is much more likely for a rural user to encounter very poor legacy mobile broadband service than an urban user.

A notable change in Spring 2017 is the effective elimination of legacy rural coverage for AT&T. Sprint, T-Mobile and Verizon still have remnants of legacy technology deployed and in daily use.

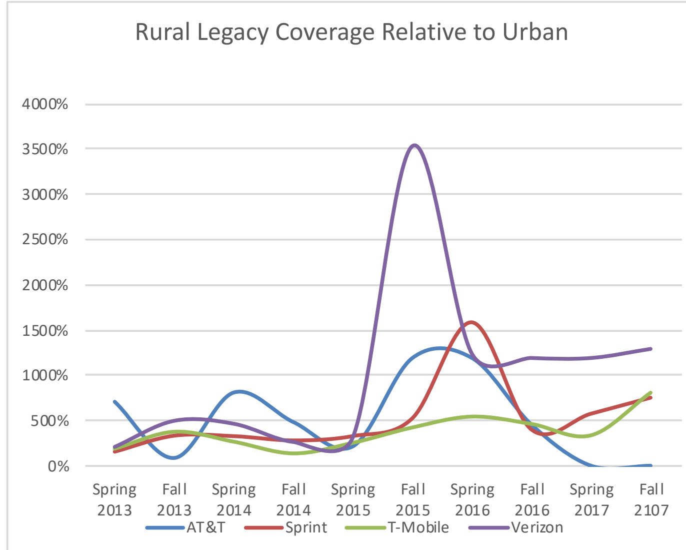

4. The Internet Degrades Mobile Broadband Service Quality

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"). Measurements to the latter gives CalSPEED an assessment to user experience to the rest of the Internet.

In the best case, the physics of data transmission[16] 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.

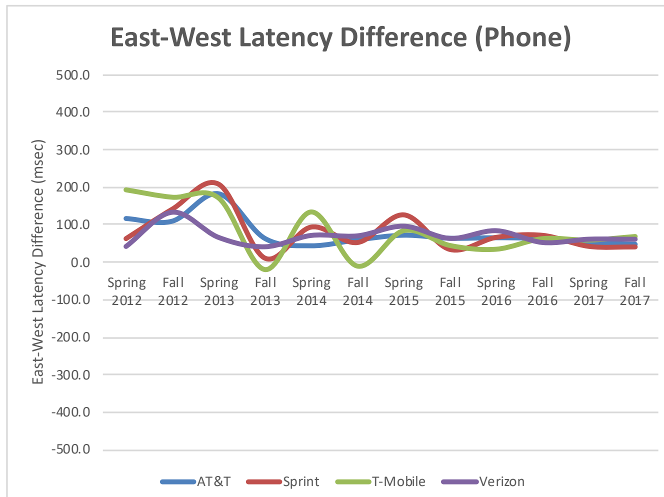

TCP throughput is related to latency … the longer the latency, the smaller the throughput[17]. Historically, we have seen that downstream throughput from the East server to California clients is 10-50% less than throughput from the West server. The chart to the right demonstrates this observation.

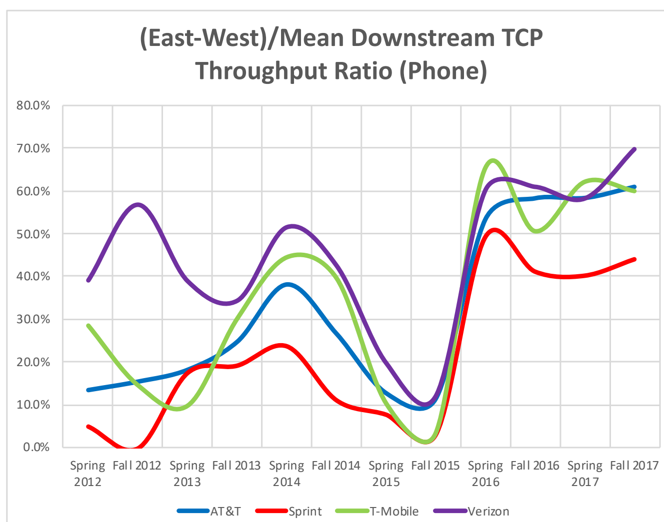

In Fall 2014 thru Spring 2015, note the substantial decrease in the TCP throughput difference and in the Spring 2016 the dramatic increase in the TCP difference for all carriers. This Spring 2016 dramatic decrease in TCP throughput to the East server has continued thru the end of 2017. At the end of 2017,

[16] Including the speed of light.

[17] A consequence of TCP's data reliability and congestion control mechanisms.

downstream TCP from the East to California users was 60% of the throughput from the West server - and only 40% for Sprint. In Fall 2015, there was essentially no difference in throughput between these servers. This behavior seems to correlate with changes in traffic shaping we can see in the histogram of throughput for each carrier - analyzed below.

CalSPEED servers are based in the Amazon AWS service and landline fiber and cable measurements have observed performance of over 500 Mb/s[18] downstream and upstream from the West server and over 50 Mb/s downstream and upstream from the East server. The CalSPEED Home measurement engine has been benchmarked at over 900 Mb/s to a local server using Ethernet. This strongly suggests neither AWS nor the CalSPEED measurement tool limit performance, but rather limitations of either the access network or the choice of Internet routing.

This becomes apparent as histograms of upstream and downstream measured throughput are examined. These charts plot the percentage of measurement resultsBelow are histogram charts of West and East, downstream and upstream throughput for all carriers. Some patterns are clear.

- Increasing average downstream and upstream throughput. Not easily seen on these charts but refer to the charts in Section 2.
- Increasing median throughput - downstream and upstream for all carriers. Easily seen at the 50% percentile in the histogram.
- Limits on maximum downstream West throughput. The high water mark for most carriers was the Spring 2016 measurement with maximum downstream West throughput occasionally exceeding 60 Mb/s. Since then West downstream throughput has been sharply limited for all carriers to below 40 Mb/s. This is despite wired measurements exceeding 500 Mb/s to the same server.
- Limits on maximum downstream East throughput. The high water mark, for all carriers, was Fall 2015 when maximum downstream East throughput exceeded 30 Mb/s and ofter 40 Mb/s. Since then, maximum throughput is rather abruptly curtailed to be below 20 Mb/s. This is consistent with the above analysis of the ratio of mean East and West throughput.
- Limits on maximum upstream West throughput. The high water mark for upstream West throughput was Spring 2016 or earlier. Subsequent upstream West throughput appears capped at 20-22 Mb/s for all carriers.
- Limits on maximum upstream East throughput. The high water mark for upstream throughput was Fall 2015 at close to 40 Mb/s. Subsequent to Fall 2015, maximum upstream East throughput seems capped at 10-12 Mb/s.

This behavior can be seen in histogram plots of all the throughput measurements for each carrier, upstream and downstream, East and West server.

[18] Ken Biba,"CalSPEED Home: Preliminary Measurements of Residential Wired and WiFi Broadband Quality", Novarum, January 2019.

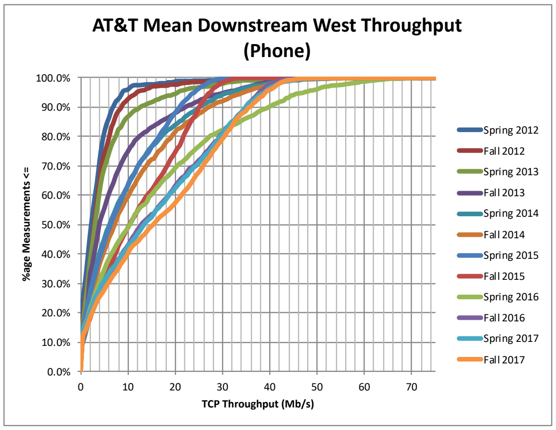
AT&T Mean Downstream West Throughput
(Phone)
%age Measurements <=
TCP Throughput (Mb/s)
Spring 2012
Fall 2012
Spring 2013
Fall 2013
Spring 2014
Fall 2014
Spring 2015
Fall 2015
Spring 2016
Fall 2016
Spring 2017
Fall 2017

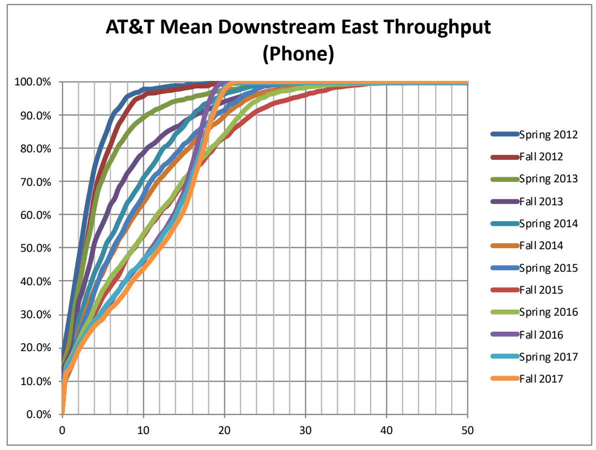
AT&T Mean Downstream East Throughput
(Phone)
Spring 2012
Fall 2012
Spring 2013
Fall 2013
Spring 2014
Fall 2014
Spring 2015
Fall 2015
Spring 2016
Fall 2016
Spring 2017
Fall 2017

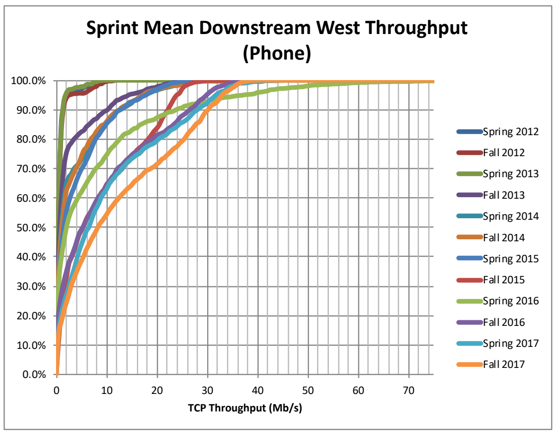
Sprint Mean Downstream West Throughput
(Phone)
TCP Throughput (Mb/s)
Spring 2012
Fall 2012
Spring 2013
Fall 2013
Spring 2014
Fall 2014
Spring 2015
Fall 2015
Spring 2016
Fall 2016
Spring 2017
Fall 2017

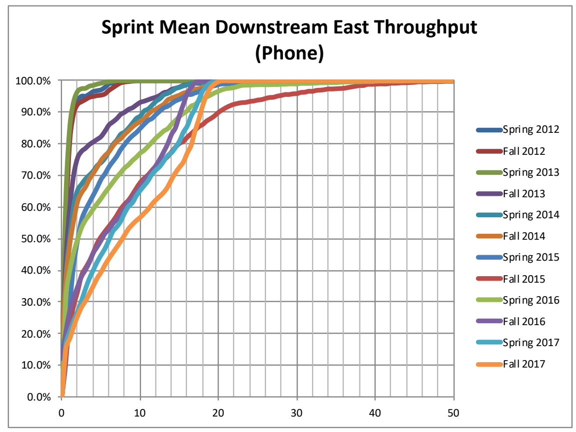
Sprint Mean Downstream East Throughput
(Phone)
Spring 2012
Fall 2012
Spring 2013
Fall 2013
Spring 2014
Fall 2014
Spring 2015
Fall 2015
Spring 2016
Fall 2016
Spring 2017
Fall 2017

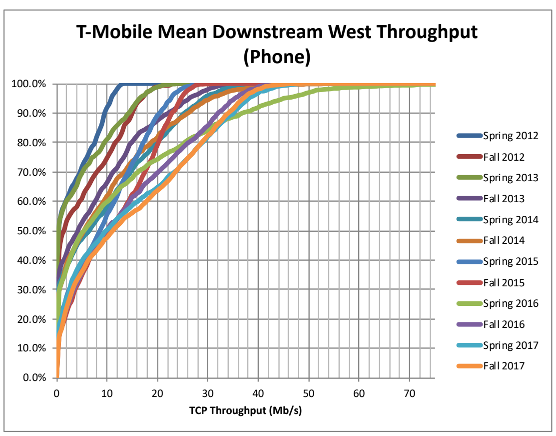
T-Mobile Mean Downstream West Throughput
(Phone)
TCP Throughput (Mb/s)
Spring 2012
Fall 2012
Spring 2013
Fall 2013
Spring 2014
Fall 2014
Spring 2015
Fall 2015
Spring 2016
Fall 2016
Spring 2017
Fall 2017

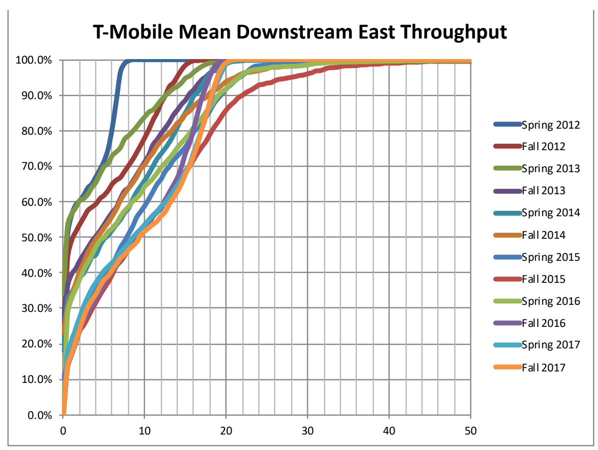
T-Mobile Mean Downstream East Throughput
Spring 2012
Fall 2012
Spring 2013
Fall 2013
Spring 2014
Fall 2014
Spring 2015
Fall 2015
Spring 2016
Fall 2016
Spring 2017
Fall 2017

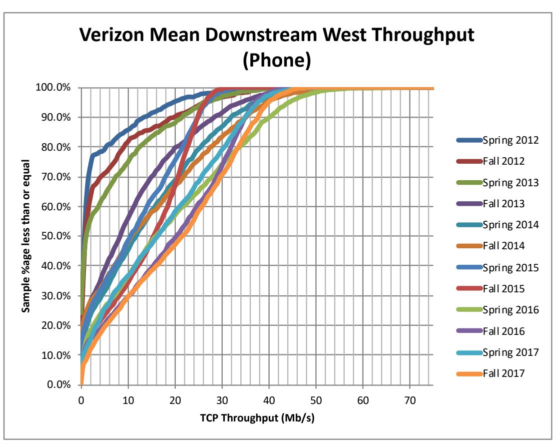
Verizon Mean Downstream West Throughput
(Phone)
Sample %age less than or equal
TCP Throughput (Mb/s)
Spring 2012
Fall 2012
Spring 2013
Fall 2013
Spring 2014
Fall 2014
Spring 2015
Fall 2015
Spring 2016
Fall 2016
Spring 2017
Fall 2017

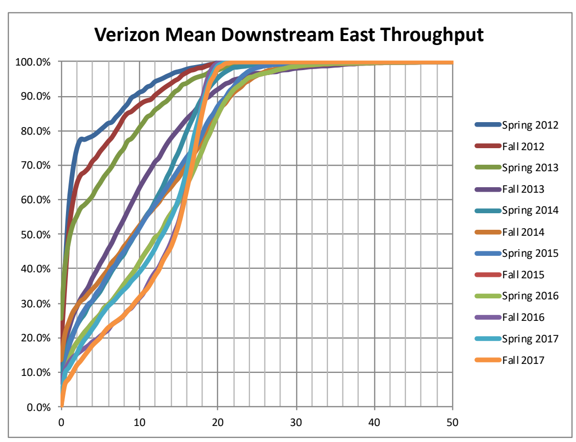
Verizon Mean Downstream East Throughput
Spring 2012
Fall 2012
Spring 2013
Fall 2013
Spring 2014
Fall 2014
Spring 2015
Fall 2015
Spring 2016
Fall 2016
Spring 2017
Fall 2017

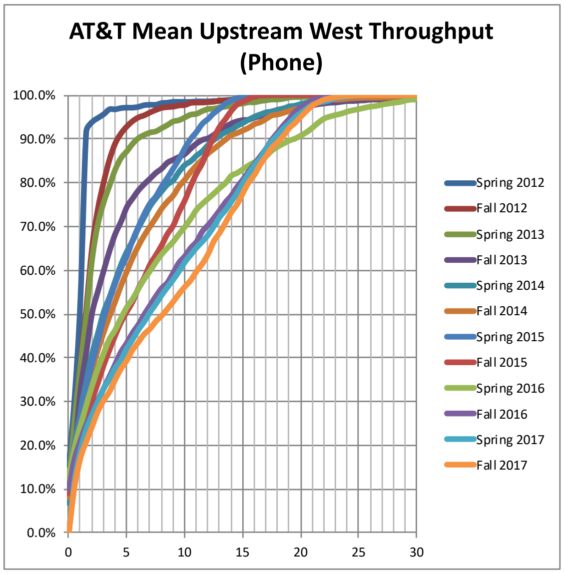
AT&T Mean Upstream West Throughput (Phone)
100.0%
90.0%
80.0%
70.0%
60.0%
50.0%
40.0%
30.0%
20.0%
10.0%
0.0%
0
5
10
15
20
25
30
Spring 2012
Fall 2012
Spring 2013
Fall 2013
Spring 2014
Fall 2014
Spring 2015
Fall 2015
Spring 2016
Fall 2016
Spring 2017
Fall 2017

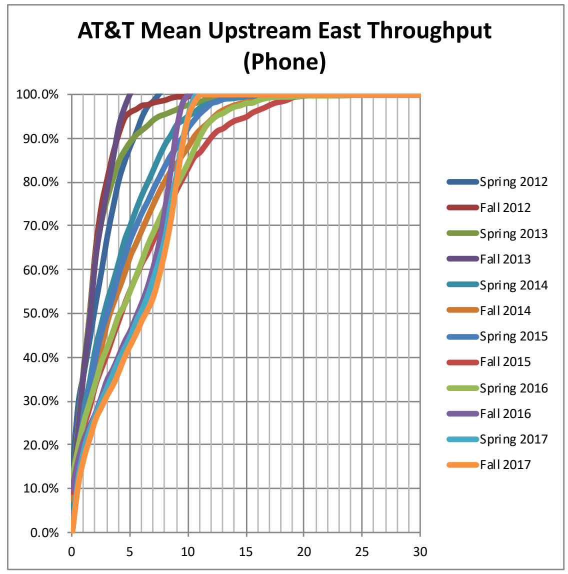
AT&T Mean Upstream East Throughput (Phone)
100.0%
90.0%
80.0%
70.0%
60.0%
50.0%
40.0%
30.0%
20.0%
10.0%
0.0%
0
5
10
15
20
25
30
Spring 2012
Fall 2012
Spring 2013
Fall 2013
Spring 2014
Fall 2014
Spring 2015
Fall 2015
Spring 2016
Fall 2016
Spring 2017
Fall 2017

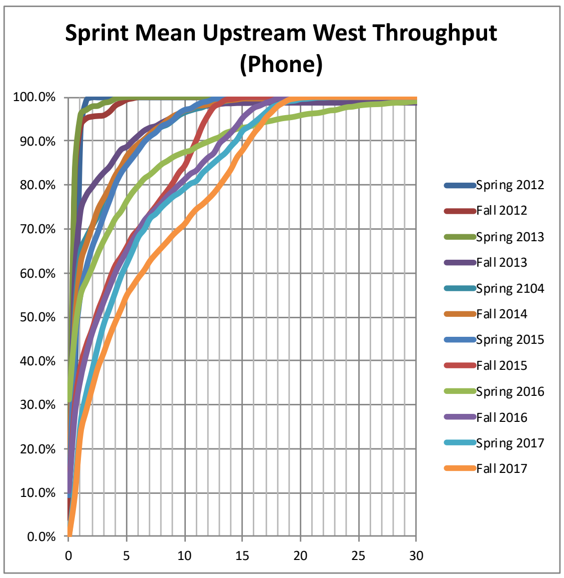
Sprint Mean Upstream West Throughput (Phone)
100.0%
90.0%
80.0%
70.0%
60.0%
50.0%
40.0%
30.0%
20.0%
10.0%
0.0%
0
5
10
15
20
25
30
Spring 2012
Fall 2012
Spring 2013
Fall 2013
Spring 2104
Fall 2014
Spring 2015
Fall 2015
Spring 2016
Fall 2016
Spring 2017
Fall 2017

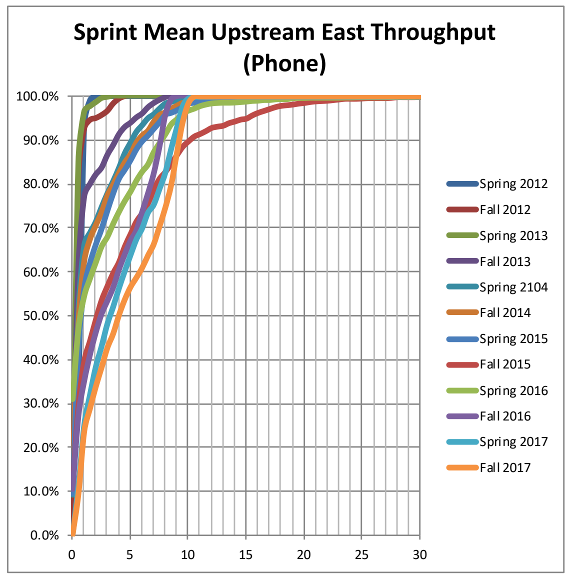
Sprint Mean Upstream East Throughput (Phone)
100.0%
90.0%
80.0%
70.0%
60.0%
50.0%
40.0%
30.0%
20.0%
10.0%
0.0%
0
5
10
15
20
25
30
Spring 2012
Fall 2012
Spring 2013
Fall 2013
Spring 2104
Fall 2014
Spring 2015
Fall 2015
Spring 2016
Fall 2016
Spring 2017
Fall 2017

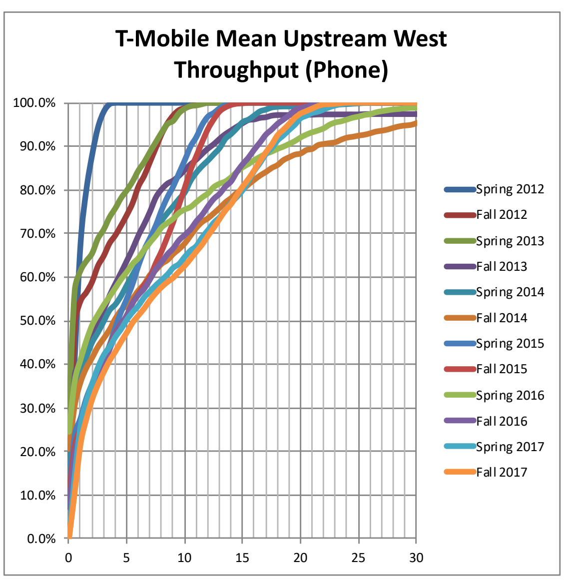
T-Mobile Mean Upstream West Throughput (Phone)
100.0%
90.0%
80.0%
70.0%
60.0%
50.0%
40.0%
30.0%
20.0%
10.0%
0.0%
0
5
10
15
20
25
30
Spring 2012
Fall 2012
Spring 2013
Fall 2013
Spring 2014
Fall 2014
Spring 2015
Fall 2015
Spring 2016
Fall 2016
Spring 2017
Fall 2017

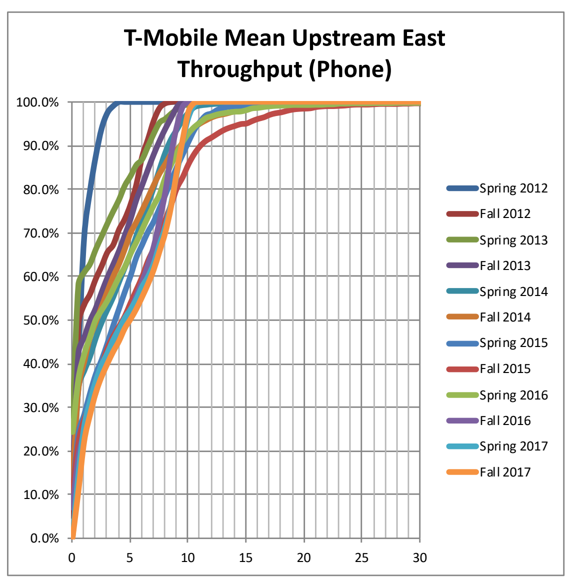
T-Mobile Mean Upstream East Throughput (Phone)
100.0%
90.0%
80.0%
70.0%
60.0%
50.0%
40.0%
30.0%
20.0%
10.0%
0.0%
0
5
10
15
20
25
30
Spring 2012
Fall 2012
Spring 2013
Fall 2013
Spring 2014
Fall 2014
Spring 2015
Fall 2015
Spring 2016
Fall 2016
Spring 2017
Fall 2017

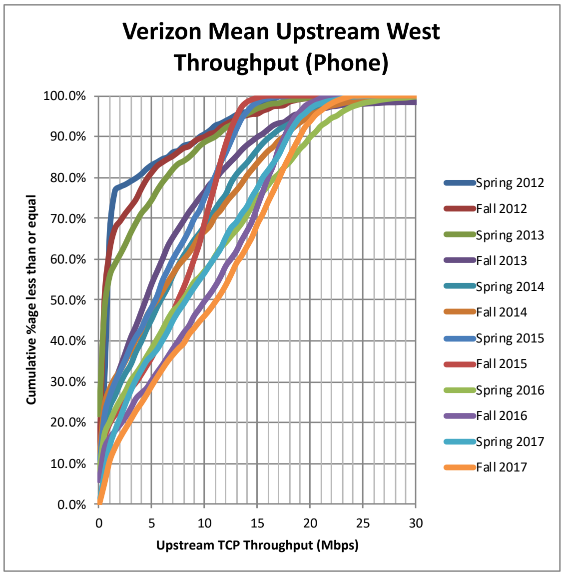
Verizon Mean Upstream West Throughput (Phone)
Cumulative %age less than or equal
Upstream TCP Throughput (Mbps)
0.0%
10.0%
20.0%
30.0%
40.0%
50.0%
60.0%
70.0%
80.0%
90.0%
100.0%
0
5
10
15
20
25
30
Spring 2012
Fall 2012
Spring 2013
Fall 2013
Spring 2014
Fall 2014
Spring 2015
Fall 2015
Spring 2016
Fall 2016
Spring 2017
Fall 2017

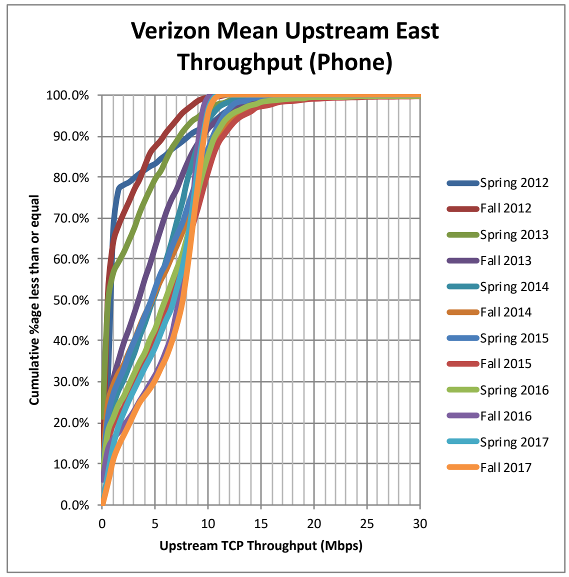
Verizon Mean Upstream East Throughput (Phone)
Cumulative %age less than or equal
Upstream TCP Throughput (Mbps)
0.0%
10.0%
20.0%
30.0%
40.0%
50.0%
60.0%
70.0%
80.0%
90.0%
100.0%
0
5
10
15
20
25
30
Spring 2012
Fall 2012
Spring 2013
Fall 2013
Spring 2014
Fall 2014
Spring 2015
Fall 2015
Spring 2016
Fall 2016
Spring 2017
Fall 2017

5. Mobile Broadband Coverage

Let's first look at mobile broadband coverage at the 25 Mb/s down/3 Mb/s up broadband standard - then at the new FCC proposed 10/1 mobile broadband standard. This analysis estimates that coverage by the percentage of measurement locations meeting the standard.

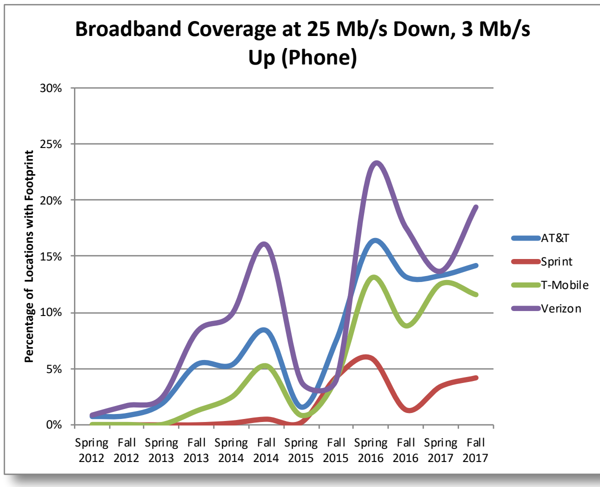

In the Spring of 2015, there was a dramatic (~50%) decrease for AT&T, T-Mobile and Verizon in the number of measurement locations meeting the 25/3 standard. For the Spring 2016 measurement round a similar dramatic increase in the number of locations meeting standard is seen. The Fall 2016 measurement saw a material decrease in the number of locations meeting the standard. Spring 2017 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 essentially stable with a ratio of ~75% 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.

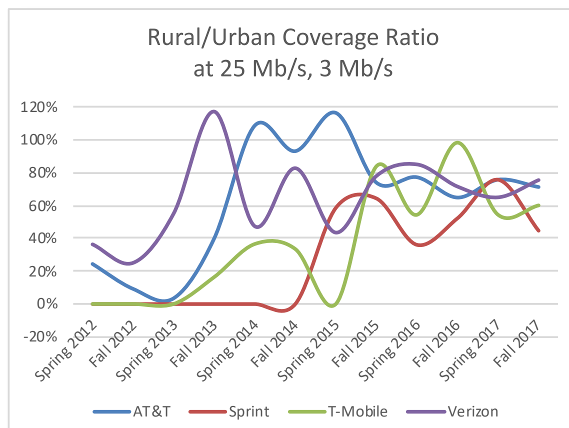

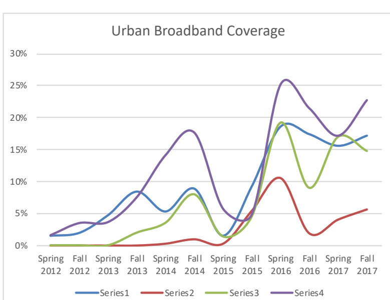

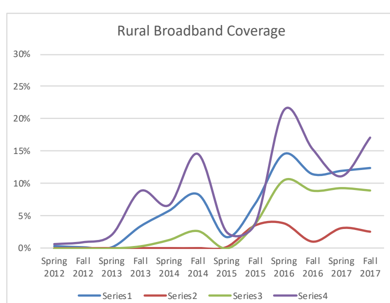

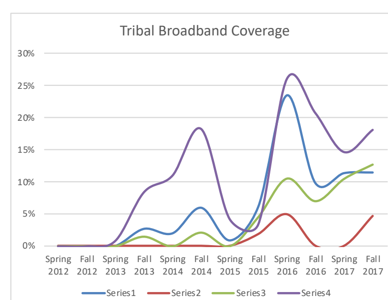

When looking at the 10/1 standard, we see substantive differences from the 25/3 standard.

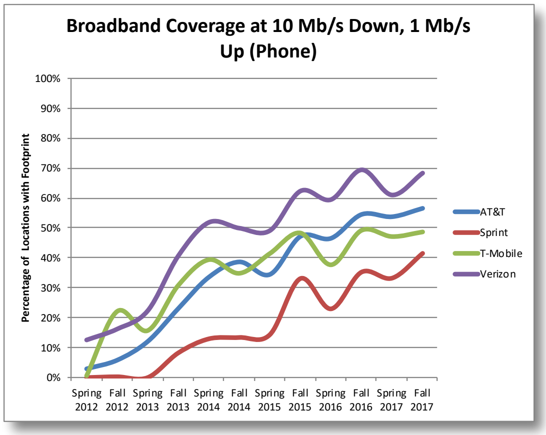

First, the average broadband coverage at the end of 2017 for the 10/1 standard is substantially higher that at 25/3 standard, at ~50% rather than the ~13% for 25/3. Second, the coverage at 10/1 has grown over the last 5 years at a steady rate rather than the dramatic swings at the 25/3 standard.

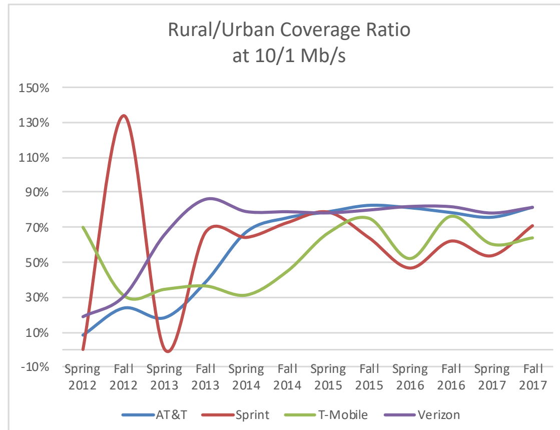

Yet, while the patterns of coverage are quite different from the 25/3 standard, the rural/urban ratio is surprisingly similar for 10/1 (~75%) and 25/3 (75%). Note that at 10/1, this coverage ratio was apparent as early as Spring 2014, while a similar ratio appeared for 25/3 eighteen months later in Fall 2015. So regardless of the standard, rural and tribal users are getting about 50-75% of the service availability of urban users.

Demographic differences mirror the overall coverage. Whichever way mobile broadband coverage is define, urban users continues to dominate rural and tribal users.

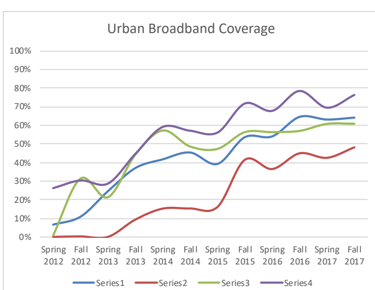

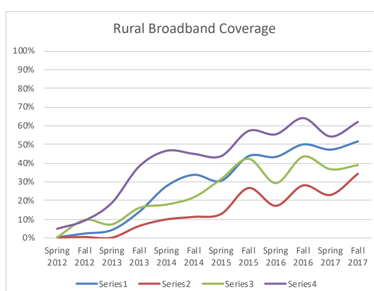

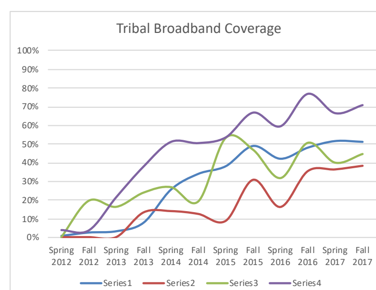

6. Mobile Technology Deployment Is About to Transition (Again)

The CalSPEED Mobile Broadband measurement covered the time of the change of mobile broadband from largely 3G technology thru the deployment of 4G LTE among all four major carriers. At the beginning of this project, legacy 1G and 2G connections were used for between 10-20% of all measurements while LTE connections were used for between 0-25% of all connections. There were major differences between carriers with Sprint and T-Mobile, in particular, being slow to deploy LTE technology while Verizon clearly led in deployment.

By the end of 2017, all four carriers are above 80% connections of 4G LTE and 1/2G legacy connections are under 5%. We can expect that 5G service will begin for some urban users by the end of 2019. It is not unreasonable to expect that there will be a long tail of legacy 1/2G service, particularly in rural or tribal areas with small user populations, and a capping of 4G service availability as carriers concentrate on 5G deployment.

These trends will be especially focussed on rural and tribal users as deployment economics will bias towards infrastructure investment covering more users per cellular tower. The CalSPEED data suggests about a 3-4 year deployment window as rural and tribal high performance infrastructure begins to catch up with urban deployment.

CalSPEED data documents that rural and tribal users had slower removal of 1/2G legacy service and slower deployment of 4G LTE service than urban users. At the end of 2017, rural users have a 5-10x greater likelihood of legacy 1/2G service than urban users as well as only an 80% likelihood of 4G LTE service.

Advanced network services (such as VoIP and streaming video) are really only practical on LTE (or on 5G) and the extent to which LTE is not available, or legacy services are experienced instead, will dramatically degrade mobile broadband service for rural and tribal users.

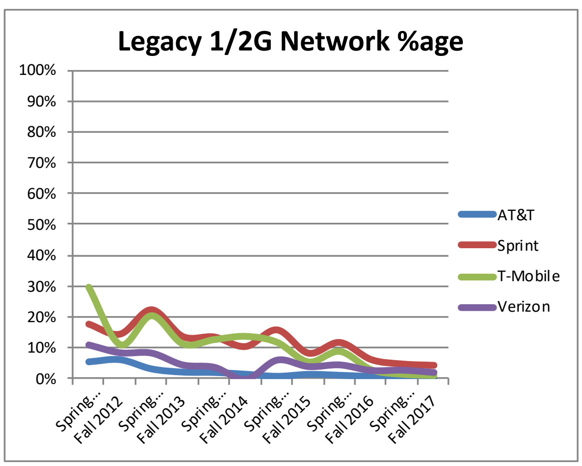

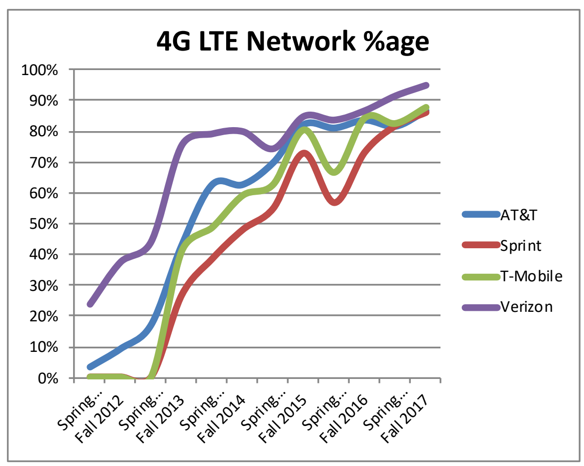

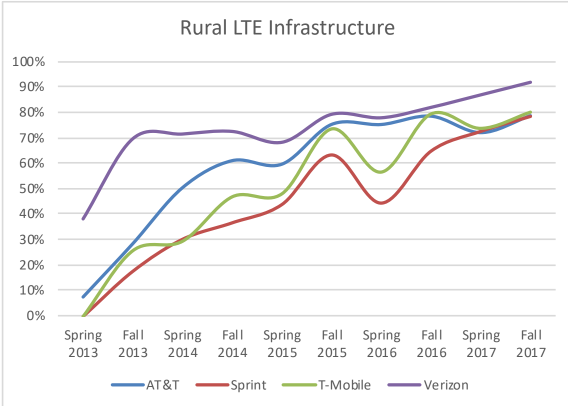

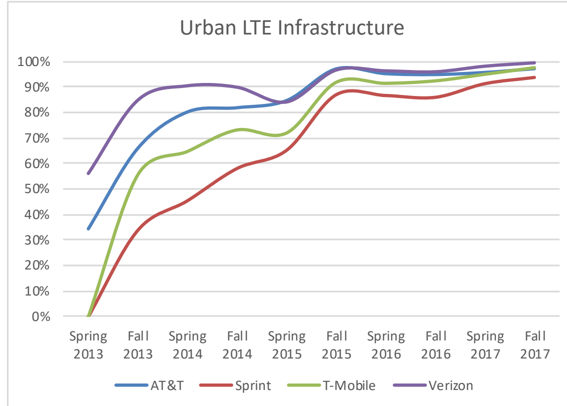

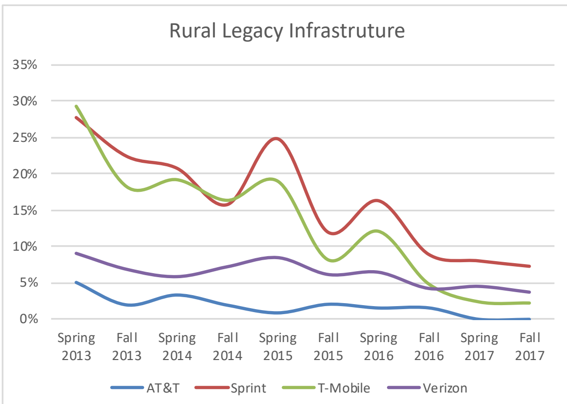

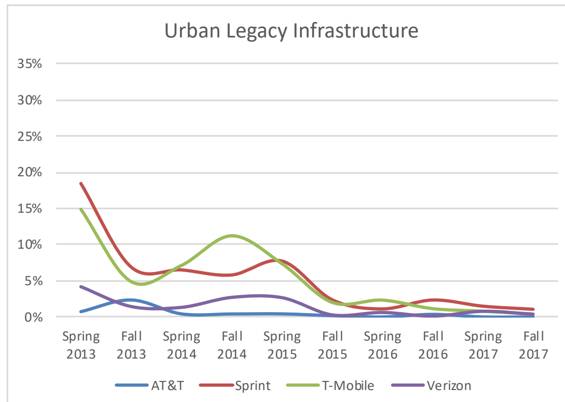

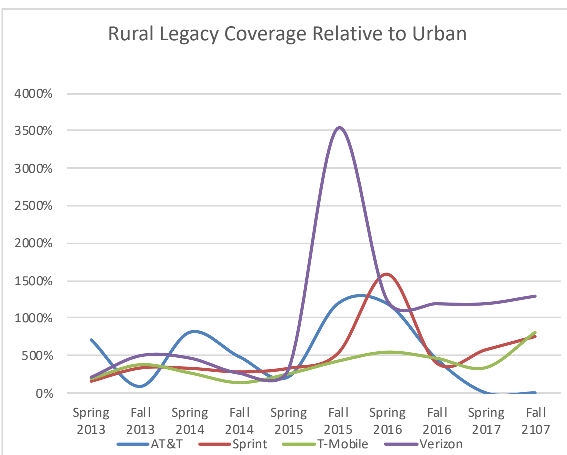

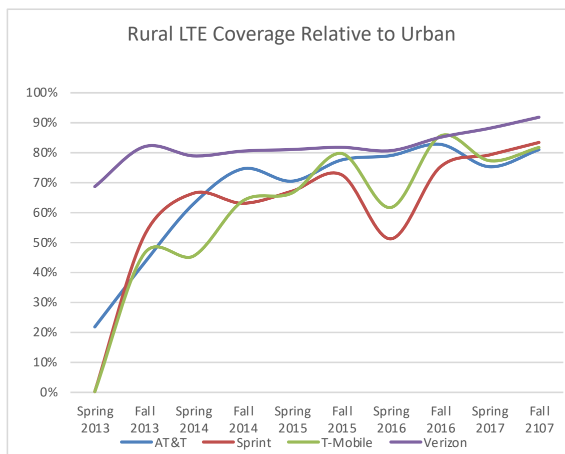

A Caution About the Coming Transition to 5G

CalSPEED Mobile documented the California mobile broadband transition from 3G to 4G LTE. One of its primary observations is the lag in service between urban users and rural/tribal users. It took 3-4 years for rural and tribal users to begin to get the same level of advanced LTE service and urban users. And as documented in Section 3, today's rural/tribal mobile broadband quality is still degraded from urban users by a factor of at least 1/3 - particularly in advanced streaming service for voice and video.

The is much carrier provided excitement about the potential of 5G services. It is important to recognize that there are really two, quite different, 5G services and their impact on rural and tribal users will be quite different.

The excitement about 5G is mostly about new services in new millimeter wave frequency bands[19]. These bands have substantial new capacity for very high performance (up to gigabit/second speeds) but have some substantial not well articulated limitations.

- The physics of radio propagation at these frequencies is challenging. Limited propagation around corners (of buildings), through foliage, almost zero propagation through walls and glass.
- Some mitigation of these limitations can be accomplished with extensive use of advanced MIMO antennas - but these antennas are large, expensive and in the near term will not be available in affordable user devices. But even with these antennas, there are severe deployment challenges both technically and economically for carriers.
- The consequences of the physics mandate that transmissions at these frequencies will require a base station within every room, at every street corner. This deployment density is likely prohibitive for rural and tribal infrastructure, certainly in the medium term.

In addition to excitement for mmwave bands, 5G offers incremental improvements to existing 4G LTE services, in the largely the same frequency bands and assets currently deployed. These bands will offer similar coverage to existing LTE and similar performance and will likely be the only 5G service to rural and tribal users except in special cases. Only modest, if any, performance and quality gains are anticipated for this version of 5G. AT&T is already marketing this version of 5G under the 5G E(volution) branding.

Urban users, where mmwave infrastructure is deployed and users purchase new phones including this technology, may see substantial increase in performance (5-10x?) and decrease in latency that will make interactive and streaming services even more effective. These are unlikely to be available to rural and tribal users where population density and geography make 5G mmwave deployments uneconomical or physically impossible.

It is likely that the current mobile broadband digital divide between urban and rural/tribal users will not only widen, but widen dramatically.

[19] In California, these bands will be between 30-60 GHz.

7. Old Devices Give Old Performance

CalSPEED does it measurements with the highest performance mobile user device, for that carrier, available at the time of measurement. It thus assesses the best quality of service that deployed network infrastructure of each carrier can deliver to a user with the best quality equipment. But - users do not always upgrade to the latest technology. What is the penalty in performance and quality for not upgrading?

CalSPEED has always measured each carrier, in each location, with two devices. One a smartphone and the other has varied. Originally it was a USB modem stick with a Windows laptop, and in the Spring of 2015 was changed to a then current technology tablet. The device type was changed because while the technology of the smartphones improved (essentially on an annual basis), the USB stick did not. Likely because of diminishing demand not motivating the investment.

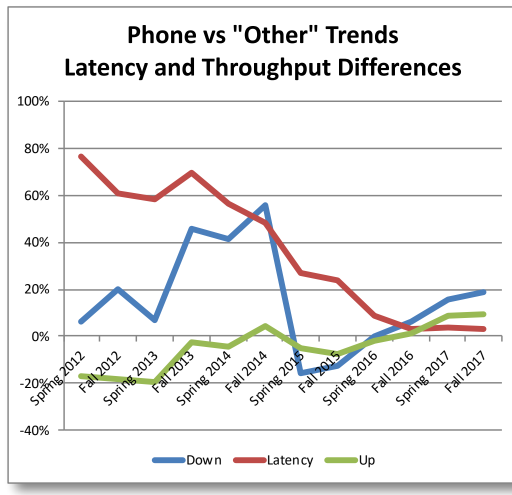

And the lack of change in technology was obvious in the results. The chart on the right documents the relative difference between the smartphone and the "other" device (USB stick prior to Spring 2015, tablet subsequently). For simplicity, this chart aggregates all carriers together despite having different models of device.

There was always a big latency difference between a smartphone and a USB stick, but the difference in throughput was modest. But clear differences in throughput began to quickly appear, particularly downstream. The change to tablet in 2015 was abrupt, with the differences between devices largely eliminated.

While CalSPEED continued measurement with two devices thru 2017, the core analysis is based solely on continuously updated smartphone user devices that paced carrier deployment.

What can be learned from this data?

- Different devices can have dramatically different performance.
- This difference could easily be ~25%/year which means a 3 year old user device (a phone with 3 year old technology) can easily be a factor of 2 lower in performance than a state of the art device.

While not studied in CalSPEED Mobile, the real user experience is a mixture of device types, old and new. Many users choose, or cannot afford, the latest technology. These users will not have the performance and quality of service documented by CalSPEED - but rather something substantially less.

Pew Research[20], in 2018, found that 26% of rural users used a legacy cellphone rather than a smartphone (much less a modern smartphone) compared to only 13% of urban users. While Pew did not ask the question about the generation of smartphone, a reasonable inference is that the age of rural user smartphones trends older than urban users. CalSPEED suggests that a 3 year old smartphone has half the performance of a current smartphone. This compounds the 3/5ths quality deficit of the mobile infrastructure.

[20] "Mobile Fact Sheet", Pew Research Center, February 2018, http://www.pewinternet.org/fact-sheet/mobile/

8. The Signal is Not the Internet Experience

Carriers would like users to believe that the amount of signal "bars" means better service. And while in the days of analog voice service (in the previous century) this was certainly true, for modern digital mobile networks this is much less true. While some signal is certainly required, there is small correlation in modern networks between signal strength and network performance and a more significant correlation with a more esoteric measurement - SNR[21] - or signal to noise ratio. But not much of a predictive relationship.

Wireless network performance is coupled to two, different, underlying mechanisms. First, the difference between radio signal strength and the always present underlying noise couples to maximum modulation rate - the bits/second - that the network can carry. And second, each cell tower has multiple frequency separated channels - each of which is shared between multiple simultaneous users. For network cells with more users, at a given modulation rate, each user will get lower bandwidth as the underlying channel bandwidth is shared among those users. Both of these mechanisms are updated with new network technology.

As networks become more loaded with traffic, and more towers deployed, in modern networks there is almost always enough signal for service (except at network edges) but there is not always enough network capacity to share.

In 2017, each CalSPEED Mobile measurement in 2017 added a measurement of both raw radio signal strength - RSSI[22] - and SNR. Different mobile technologies represent these values differently with sadly no common standard - but the preponderance of LTE service in the CalSPEED dataset allows analysis by just relying on reported LTE RSSI and SNR.

The following scatter plots show the relationship between both signal strength and SNR for both urban and rural demographics for each carrier. There is only poor correlation between RSSI and Throughput and only modest correlation between SNR and throughput.

R-value Correlation				
	RSSI		SNR	
	Rural	Urban	Rural	Urban
AT&T	-0.10	-0.06	0.24	0.41
Sprint	0.26	0.38	0.67	0.61
T-Mobile	0.25	0.42	0.49	0.60
Verizon	0.33	0.45	0.66	0.62

At best, RSSI explains about 20% of the variation in throughput (Verizon urban), and at worst, explains 0% of the variation in throughput (AT&T urban). At best, SNR explains about 45% of throughput variation (Sprint rural) and at worst, about 6% of throughput variation (AT&T rural).

R^2			
RSSI		SNR	
Rural	Urban	Rural	Urban
1%	0%	6%	16%
7%	14%	45%	37%
6%	18%	24%	36%
11%	20%	43%	38%

[21] SNR measures the ratio between raw signal and the underlying noise. As more networks are deployed, there is more noise. SNR is almost directly coupled to the raw modulation speed of the wireless network. More SNR, more energy above the noise is available to carry modulated information.

[22] RSSI is the underlying measurement of signal strength that devices and carriers choose to represent raw radio signal strength. Bars are a phone unique graphic mapping to the underlying signal strength – Received Signal Strength Indicator. Sadly, each phone has a unique, almost always non-linear mapping from signal strength to graphic bars.

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, particularly for SNR.
- There is only a poor correlation between signal strength (and perhaps bars of service) and TCP throughput. There is a statistical hint that greater signal might give higher throughput.
- There is a moderate positive correlation between SNR and TCP throughput. Better SNR, on average, generally results in better TCP throughput.
- **But**. 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 (or RSSI) value and note the large range of possible TCP throughputs that were observed in the Fall 2017 data - for all carriers.

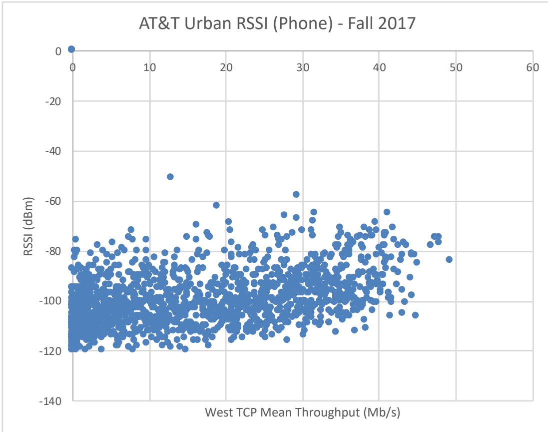
AT&T Urban RSSI (Phone) - Fall 2017
RSSI (dBm)
West TCP Mean Throughput (Mb/s)

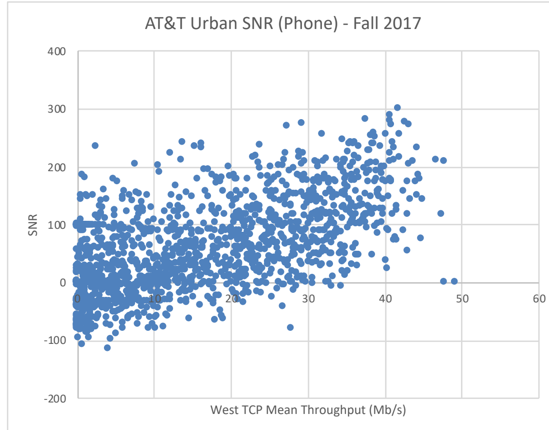
AT&T Urban SNR (Phone) - Fall 2017
SNR
West TCP Mean Throughput (Mb/s)

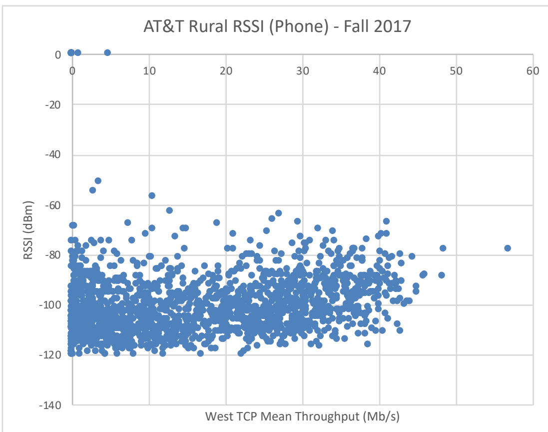
AT&T Rural RSSI (Phone) - Fall 2017
RSSI (dBm)
West TCP Mean Throughput (Mb/s)

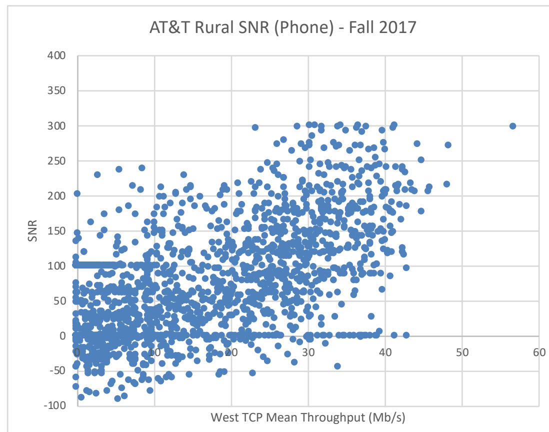
AT&T Rural SNR (Phone) - Fall 2017
SNR
West TCP Mean Throughput (Mb/s)

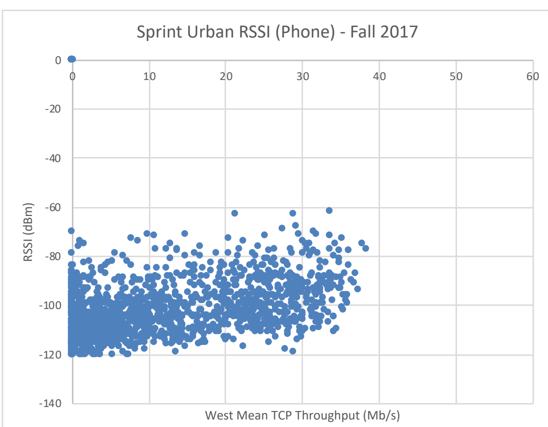
Sprint Urban RSSI (Phone) - Fall 2017
RSSI (dBm)
West Mean TCP Throughput (Mb/s)

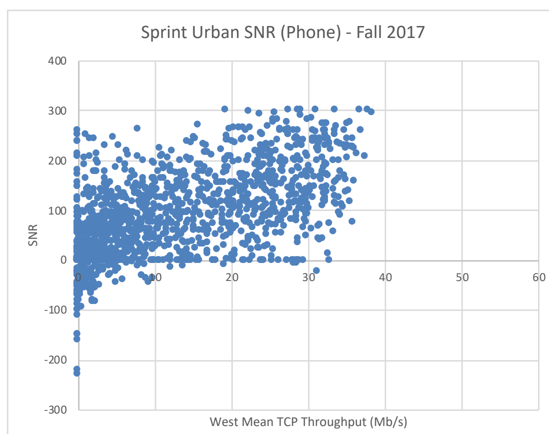
Sprint Urban SNR (Phone) - Fall 2017
SNR
West Mean TCP Throughput (Mb/s)

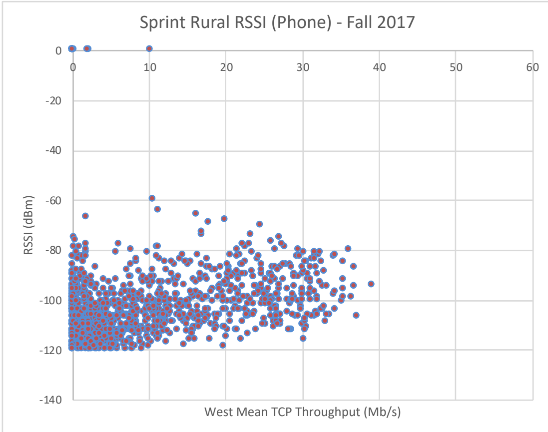
Sprint Rural RSSI (Phone) - Fall 2017
RSSI (dBm)
West Mean TCP Throughput (Mb/s)

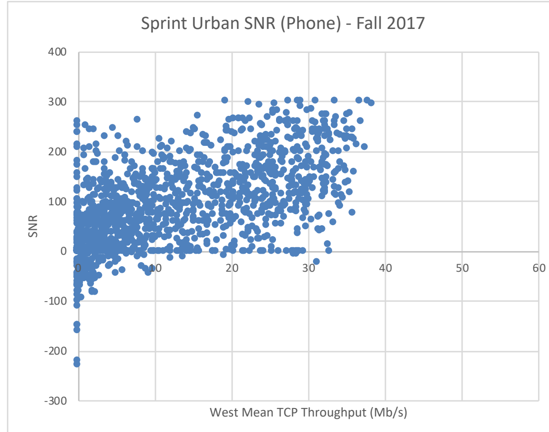
Sprint Urban SNR (Phone) - Fall 2017
SNR
West Mean TCP Throughput (Mb/s)

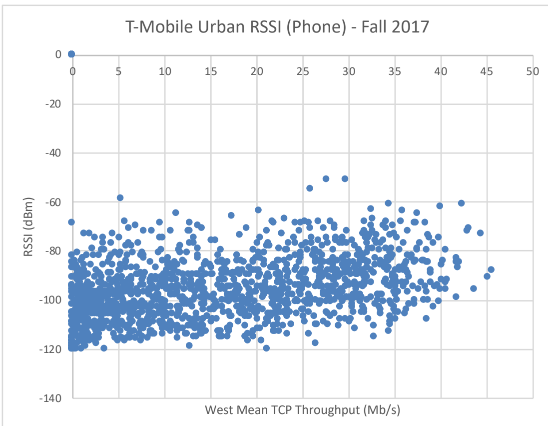
T-Mobile Urban RSSI (Phone) - Fall 2017
RSSI (dBm)
West Mean TCP Throughput (Mb/s)

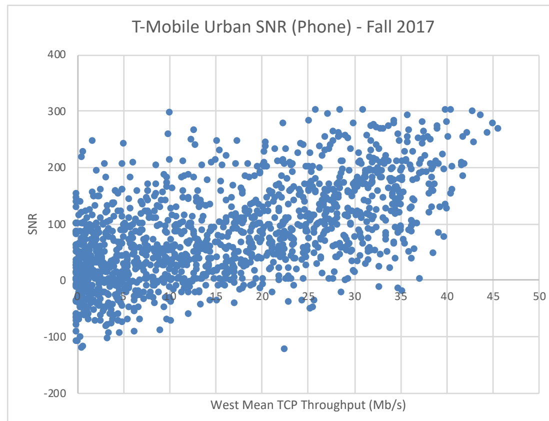
T-Mobile Urban SNR (Phone) - Fall 2017
SNR
West Mean TCP Throughput (Mb/s)

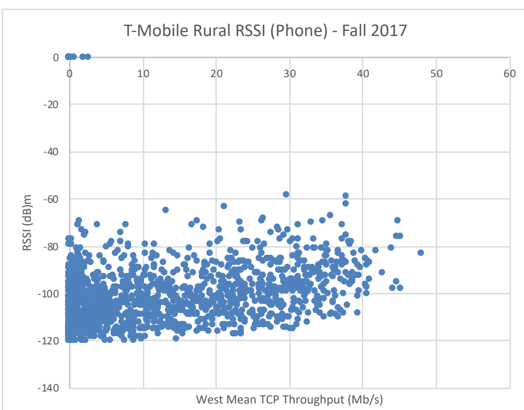
T-Mobile Rural RSSI (Phone) - Fall 2017
RSSI (dB)m
West Mean TCP Throughput (Mb/s)

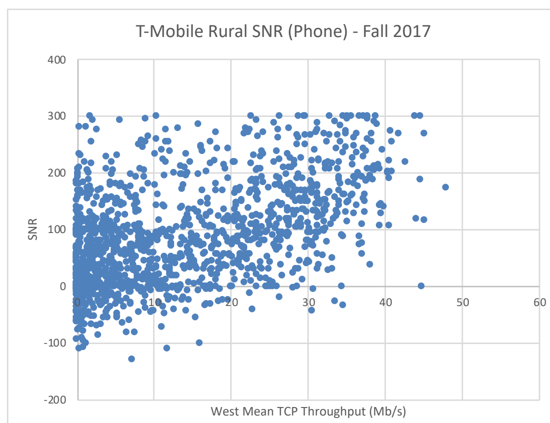
T-Mobile Rural SNR (Phone) - Fall 2017
SNR
West Mean TCP Throughput (Mb/s)

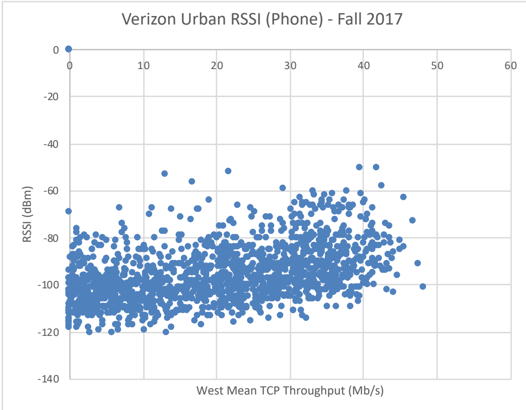
Verizon Urban RSSI (Phone) - Fall 2017
RSSI (dBm)
West Mean TCP Throughput (Mb/s)

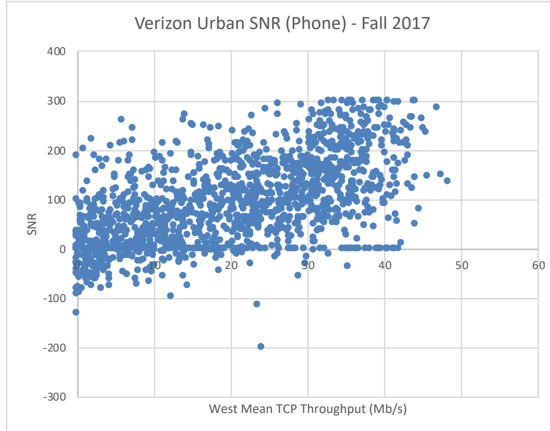
Verizon Urban SNR (Phone) - Fall 2017
SNR
West Mean TCP Throughput (Mb/s)

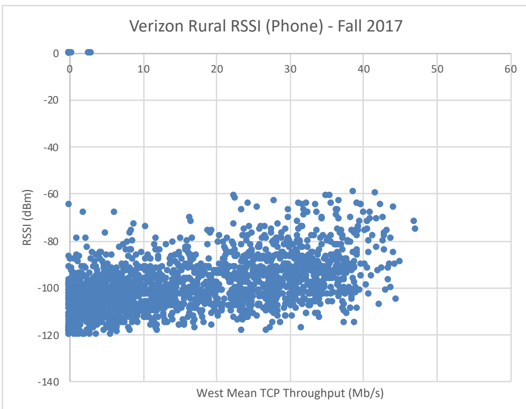
Verizon Rural RSSI (Phone) - Fall 2017
RSSI (dBm)
West Mean TCP Throughput (Mb/s)

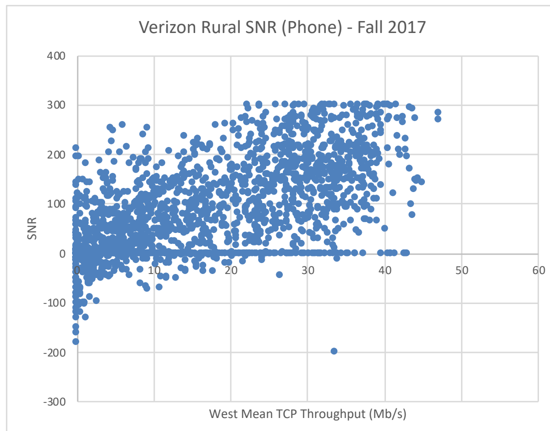
Verizon Rural SNR (Phone) - Fall 2017
SNR
West Mean TCP Throughput (Mb/s)

9. Applications

The team at the California Public Utility Commission has done a number of analyses of how the information from CalSPEED informs public policy and actions.   Two of these are particularly informative that make use of CalSPEED's unique method of creating broadband quality maps. Both try to assess the ability of mobile users to call for help in the case of emergency.   First with 911 and the second with fire.   Both critical to ascertaining the effectiveness of such services as FirstNet.

The first uses CalSPEED's estimate of mobile over-the-top voice quality (and the VoIP maps) to correlate with existing California mapping datasets for the location of 911 emergency calls.   This analysis estimates the likelihood of success of high quality 911 quality using mobile broadband.

The second uses the same VoIP maps to correlate with other state mapping datasets for the location of high fire danger.   This analysis estimates the likelihood of success of high quality voice reporting of a fire using mobile broadband in areas of high fire danger … which may have poor quality mobile service.

Both these analyses were done in 2014 with VoIP service at that time.   The analysis was restricted to AT&T and Verizon, which had the best VoIP services at that time.   As seen in Section 2, the estimated VoIP service for these two carriers has not much changed since then. Service for Sprint and T-Mobile would be of lower quality.

9.1. San Francisco Bay 911 Calls vs Mobile VoIP Quality

Let's start with AT&T and Verizon estimated VoIP service coverage as correlated with the locations of 911 calls in the San Francisco Bay Area.

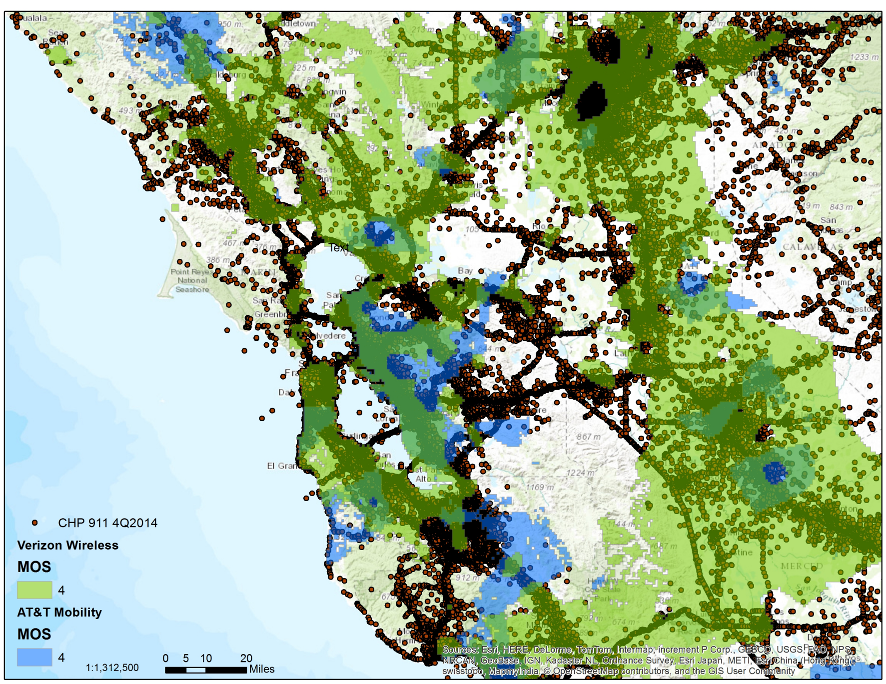

The CalSPEED estimate of the percentage of attempted VoIP calls in 911 calls without high quality over-the-top voice service as 48% for AT&T and 35% for Verizon.

AT&T Mobility "Good" VoIP Coverage			
Number of Calls Without	Percent of Calls Without	Number of Calls With	Percent of Calls With
3,297,260	48.09%	3,559,394	51.91%

Verizon Wireless "Good" VoIP Coverage			
Number of Calls Without	Percent of Calls Without	Number of Calls With	Percent of Calls With
2,399,405	34.99%	4,457,249	65.01%

9.2 Good Mobile VoIP Coverage of Fire Threat

Correlating maps of fire threat with maps of VoIP coverage gives an estimate of the likelihood of making a digital voice call in a high fire threat area of the state. The following composite map overlays good VoIP quality for AT&T and Verizon (MOS >= 4.0) with fire threat.

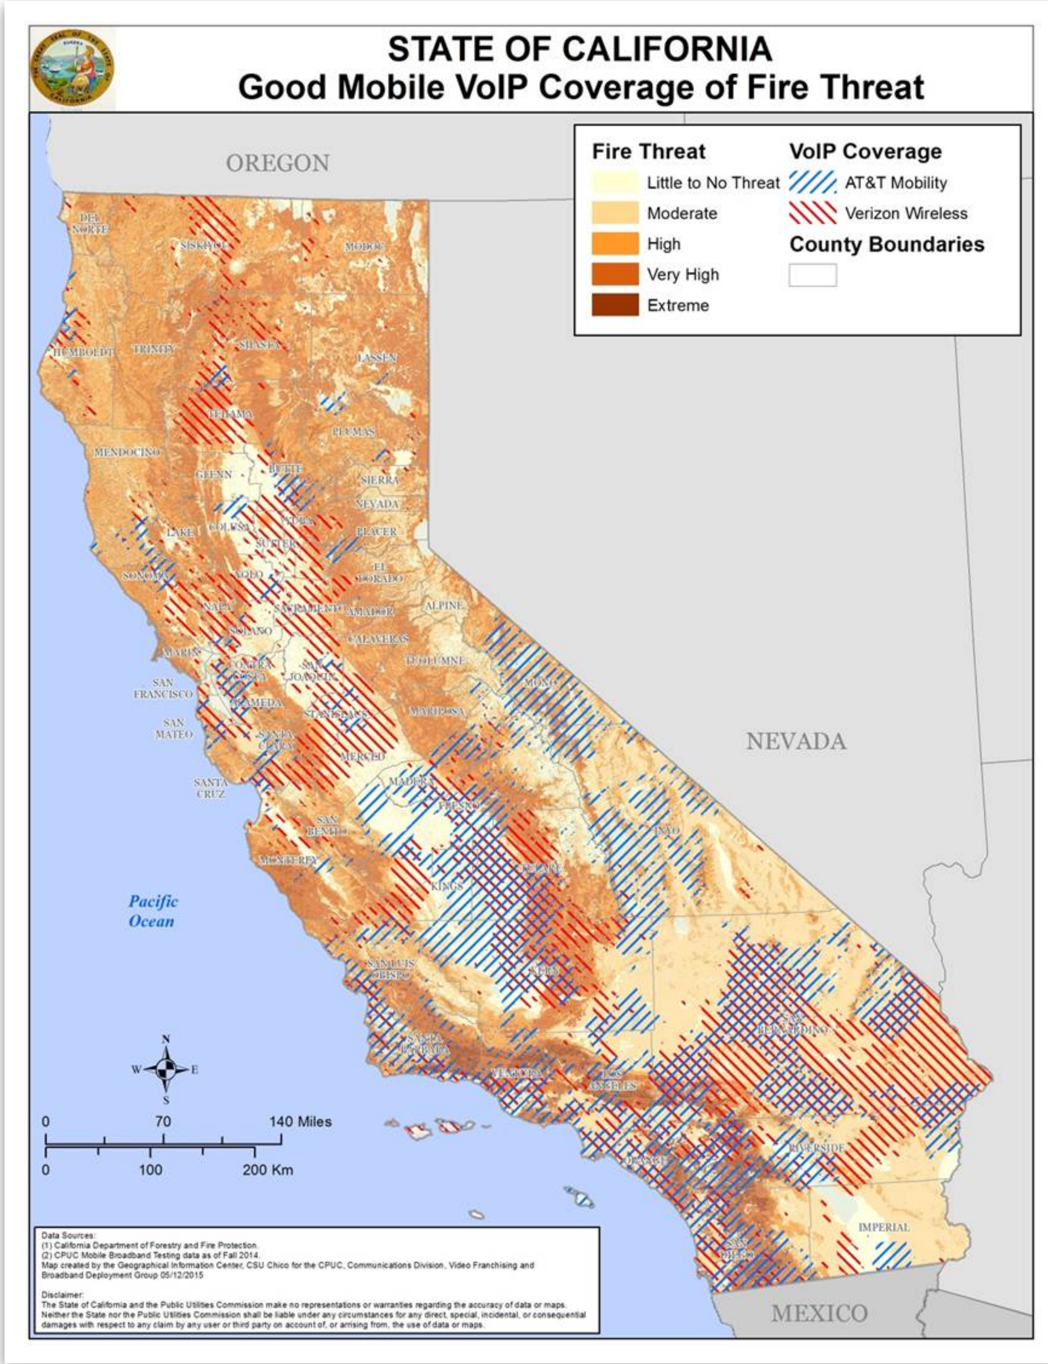

Now let's look at what this correlation reveals.

FIRE THREAT		Verizon Wireless Good VoIP Coverage			AT&T Mobility Good VoIP Coverage		
Threat Rank	Total Area (Sq. Miles)	Area Without	Area With	Percent Area Without	Area Without	Area With	Percent Area Without
Extreme	3,059	2,351	709	77%	1,397	1,663	46%
Very High	40,090	33,242	6,848	83%	32,706	7,384	82%
High	33,313	27,763	5,549	83%	27,613	5,699	83%
Moderate	56,911	37,956	18,955	67%	38,945	17,965	68%
Little to No Threat	24,780	16,014	8,766	65%	16,491	8,289	67%

Note that above 80% of the areas of very high fire risk are without high quality VoIP coverage for both AT&T and Verizon.

10. Conclusions

10.1 Accomplishments

CalSPEED Mobile has accomplished dramatic successes during its six year run.

- Insight in the Internet mobile broadband experience for millions of California users.
- The construction of great, open source broadband measurement tools.
- A measurement methodology that looks not only at local access, but the effect of the Internet on the broadband experience.
- Distillation of mobile broadband measurement into maps, that only visually describe the mobile Internet, but also provide a means to integrate broadband measurement with other geographic data - such as economics, public safety, fire fighting, and social trends.

10.2 Next Steps

CalSPEED Mobile now moves into hiatus and CalSPEED Home begins. CalSPEED Home uses the same measurement software, project structure and server infrastructure to measure residential broadband delivered to user devices either by a wire (Ethernet) or by WiFi. WiFi is important since almost all consumer Internet access is via WiFi, yet most regulatory broadband discussion is about an interface that is increasingly rarely used - Ethernet.

10.3. What We Learned

Mobile broadband quality improves overall	Mobile Internet quality has seen dramatic improvement from 2012 thru 2017: • A 4x improvement in average throughput • A 2x improvement in average latency • A 1.6x degradation in packet loss rates • A 1.7x improvement in packet jitter • A 1.5x improvement in TCP connection reliability. These changes have enabled the widespread deployment, across the state, of advanced communication services including voice over IP and streaming video.
Mobile broadband highly variable	Mobile broadband varies substantially by carrier (2x), infrastructure technology (10x), age of the user device (2x/3 years), location within state (10x) and destination information server (2x). But not much by time of day nor during a communication session (25%).

Rural/tribal mobile digital divide is substantial, persistent and likely to worsen

Rural and tribal users consistently experience about 3/5ths of the mobile broadband quality of their urban peers.

This appears for all broadband quality metrics of throughput, latency, packet loss, digital voice quality, digital video quality, jitter, connection reliability, broadband coverage and deployed modern digital communications technology.

Notably, 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 Internet application) while ~20% of rural users of those same carriers can expect TCP connection failures.

The deployment of 5G technology will likely improve urban users experience much more than rural users, increasing the broadband divide.

Mobile broadband is traffic shaped

Mobile broadband quality appears shaped, for all carriers, to have the following qualities:

- Capped maximum throughput upstream and downstream
- Improved median and average throughput
- Degraded performance outside of the regional Internet.

Mobile broadband coverage is modest

Broadband coverage at the 25 Mb/s down, 3 Mb/s is available at ~13% of all CalSPEED measurement locations at the end of 2017. At the degraded "broadband" standard of 10 Mb/s down and 1 Mb/s up, "coverage" improves to ~50% for all carriers.

But whichever way broadband is defined, rural users see about 50-75% of the service availability, at that standard as do urban users.

The End of 4G, the Hype of 5G, but 2G Lingers On

4G LTE has achieved high penetration for both urban and rural users. Some legacy 1G and 2G service areas remain particularly for rural users of Sprint, T-Mobile and Verizon. Notably there was no legacy 1/2G service detected for AT&T in California.

Urban users may see substantial increase in performance (5-10x?) and decrease in latency from the deployment of 5G that will make interactive and streaming services even more effective. These are unlikely to be available to rural and tribal users where population density and geography make 5G mmwave deployments uneconomical or physically impossible.

It is likely that the current mobile broadband digital divide between urban and rural/tribal users will not only widen, but widen dramatically.

Older Devices Mean Slower and Lower Quality Internet

CalSPEED measurements were largely made with the latest generation user devices in order to assess the deployed infrastructure. But CalSPEED has also measured with a variety of other devices in parallel.

What can be learned from this data?

- Different devices can have dramatically different performance.
- This difference could easily be ~25%/year which means a 3 year old user device (a phone with 3 year old technology) can easily be a factor of 2 lower in performance than a state of the art device.

Many users do not choose, or cannot afford, the latest technology. These users will not have the performance and quality of service documented by CalSPEED - but rather something substantially less. These user device choices amplify other differences of carrier and location.

The Signal is Not the Message

The signal bars on a smartphone are not a good predictor of mobile performance. There is little ability of signal strength to predict Internet performance.

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. A throughput value is captured for each second of each flow.
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. A throughput value is captured for each second of each flow.
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[23]. 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

23 http://en.wikipedia.org/wiki/Iperf

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.

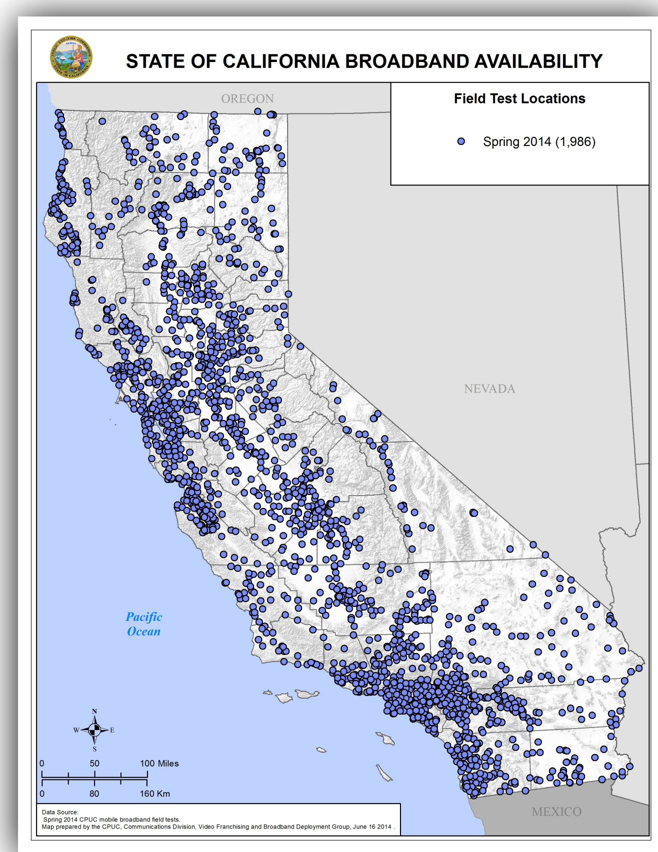

While it is impossible to measure all Internet servers, CalSPEED emulates this user experience with two fixed servers - one physically located in Northern California and the other in Northern Virginia - both in the Amazon AWS cloud. CalSPEED reports performance both to each individual server and the 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[24] 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.

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.

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 Googles' video quality metric[25] 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.

Maps for decision-makers not just for information. We then take the measurement data and create geospatial kriging[26] maps interpolating CalSPEED measurements of (but not limited to) latency, downstream and upstream throughput, jitter and packet loss over the entire state.

These maps can be overlaid with other geostatistical data on population, income, ethnicity, education, and census areas to provide more informed choices for consumers, businesses and governments. The CPUC web site uses this data to suggest what mobile service is available and at what performance at locations of the consumer's choice.

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

25 https://www.google.com/get/videoqualityreport/#methodology

26 http://en.wikipedia.org/wiki/Kriging

Massive Dataset. CalSPEED has now had eleven rounds of sampling California (Spring 2012, Fall 2012, Spring 2013, Fall 2013, Spring 2014, Fall of 2014, Spring 2015, Fall 2015, Spring 2016, Fall 2016, and Spring 2017) and is shortly to finish a twelfth round (Fall 2017). 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 the Northern Virginia Amazon AWS
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 in the Amazon AWS