CalSPEED: California Mobile Broadband - An Assessment - Fall 2014

Ken Biba
Managing Director and CTO
Novarum, Inc.

This paper examines the key findings of the first six measurement rounds of CalSPEED covering three years of measurement between Spring 2012 and Fall 2014. CalSPEED is the open source, mobile broadband measurement tool and methodology of the California Public Utilities Commission. This is an update thru Fall 2104 collected. The data reinforce the findings of that previous report and extend the foundation for five key incremental findings.

- Mobile broadband's overall performance and quality has stopped improving and shows signs of degradation.
- Mobile broadband continues trends of wide variation across California among carriers, locations of services, the growing digital divide between urban and rural,
- Quality degradation is particularly noticeable in rural areas - in which quality metrics can be 2x worse than in urban.
- Penetration of rural LTE shows signs of stalling.
- There is substantial variation between user devices on the performance and quality of service.

1. Calibrating the Mobile Internet Experience

Each of us relies on the Internet to research school papers, find a job, find and buy new products, read the news and increasingly to entertain ourselves. The Internet is not only becoming our newspaper, but also our phone, radio and television. How we do our jobs, raise our families, educate ourselves and our children, interact as responsible citizens, and entertain ourselves are all influenced by the quality of the Internet service we obtain. And ever increasingly, that service is not on our desk, but in our hand wherever we go.

Knowing the quality of this service is a vital piece of our modern ecosystem much in the same way as we research the brand of car we drive or the type of house we own. With multiple mobile Internet providers, an independent third party assessment of this quality allows consumers and policy makers to make informed choices.

CalSPEED is an open source, non-proprietary, network performance measurement tool and methodology created for the California Public Utilities Commission with the original assistance via a grant from the National Telecommunications and Information Administration. CalSPEED is now funded by California. CalSPEED uses a methodology pioneered by Novarum. The software measurement system is created and maintained by a team at California State University Monterey Bay, led by Professors Sathya Narayanan and YoungJoon Byun. CalSPEED mapping and measurement field operations are managed by the Geographic Information Center at California State University at Chico. Statisticians at CSU Monterey Bay assist the team with detailed geographic and statistical analysis of the dataset.

CalSPEED has now been in use in California for three years with six rounds of measurement over the entire state collecting over 10,000,000 measurements across California of the four major mobile broadband carriers: AT&T Mobility, Sprint, T-Mobile and Verizon Wireless. This paper does a deep analysis of the first six rounds of measurement. A previous paper analyzed the first five rounds of measurement[1]. The methodology has been rigorously analyzed with respect to other available mobile measurement tools[2].

This paper examines the incremental changes from the previous report extending thru the Fall of 2014.

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

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

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

2. Wide Variation in Mobile Broadband Continues

The following graphs show a histogram of the measured TCP throughput across the sampled locations to both of the two geographic measurement servers. Much of the growth in average throughput has occurred by dramatic increases in the high performance tail of the distribution. A minority of locations get much better throughput, while the majority of locations have much more modest improvements in throughput. The wide variation in delivered throughput across the entire sample set is apparent. For example, it is possible (though uncommon) to get a downstream throughput for Verizon to a local (West) server that is 50 Mbps even though the Verizon state-wide mean is 17.5 and the median is 13.8 Mbps.

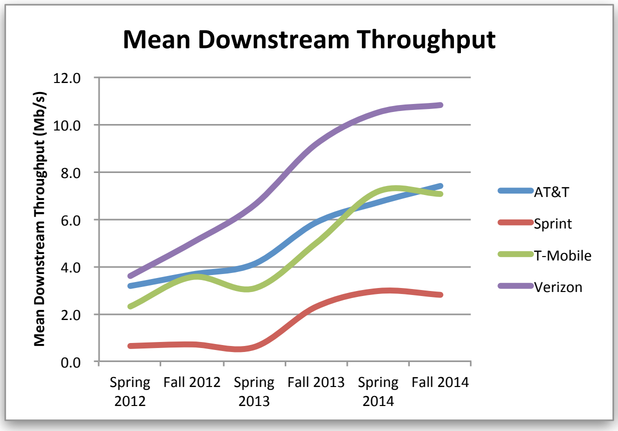

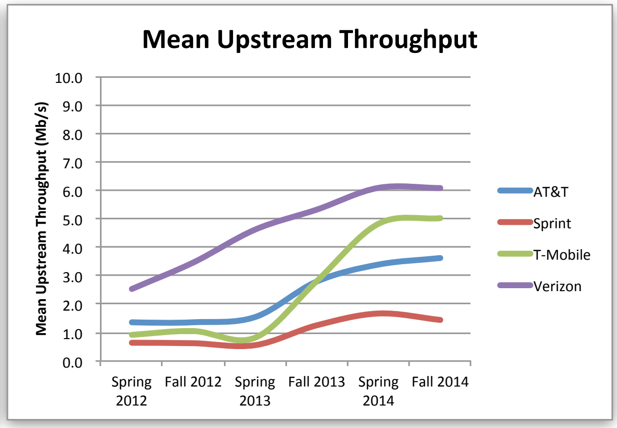

This variation in performance, echoed in other network metrics of upstream throughput, latency and jitter - is a composite of other, more fundamental variations. In order of importance these include:

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

2.1 Location of the User

The most important variation is location within the state of the user. The following interpolated kriging maps for downstream throughput for the four carriers illustrate this variation. Depending on the carrier there is almost a 25:1 variation between mean TCP downstream performance based on where in the state the user is at the time of Internet access.

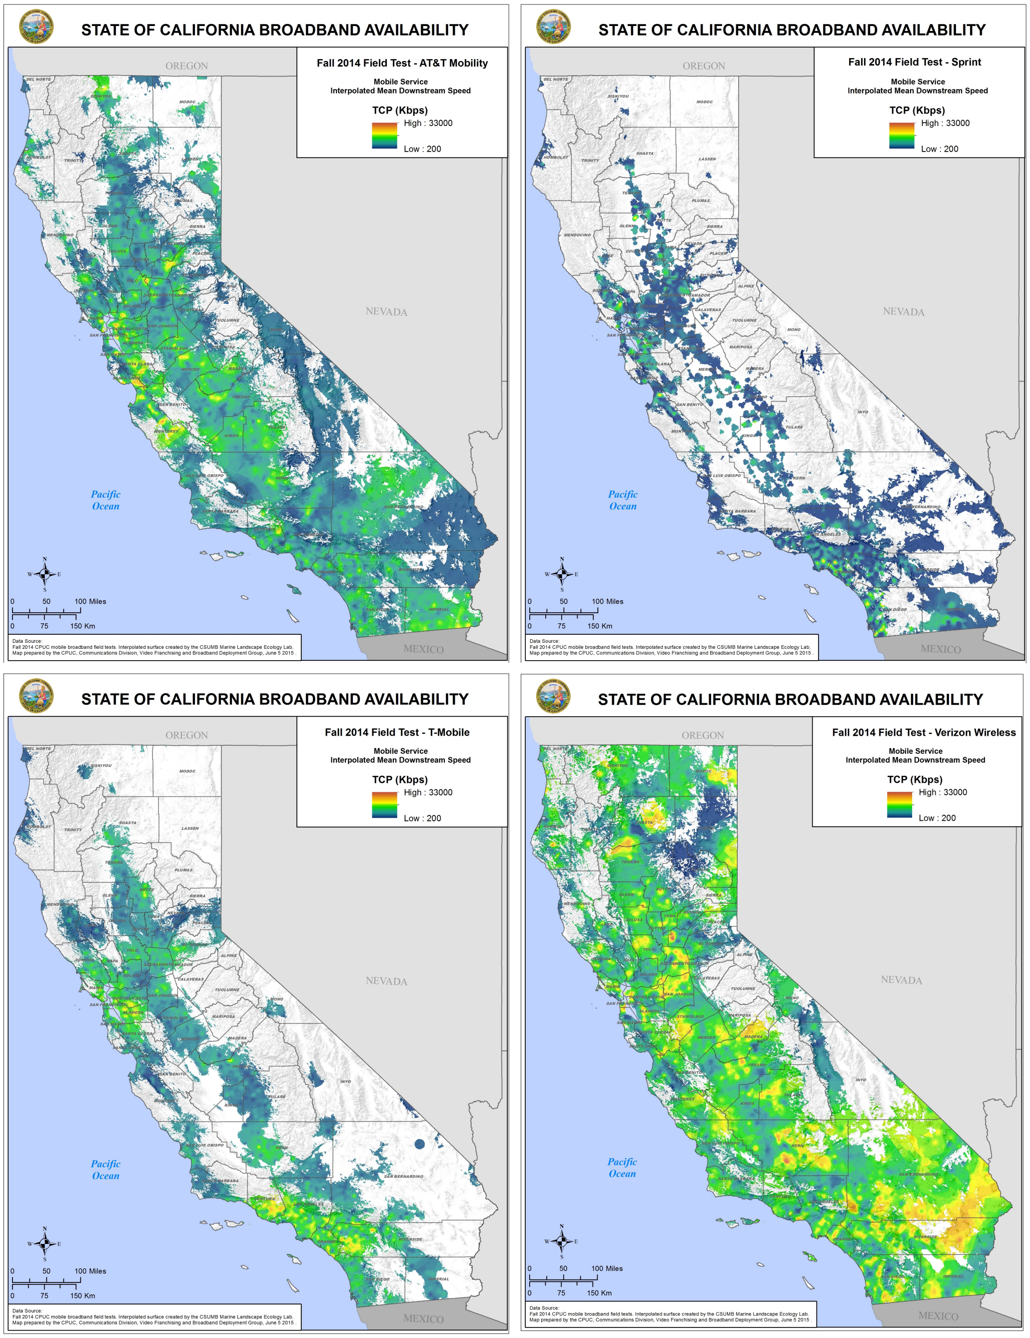

A similar variation exists for upstream TCP throughput.

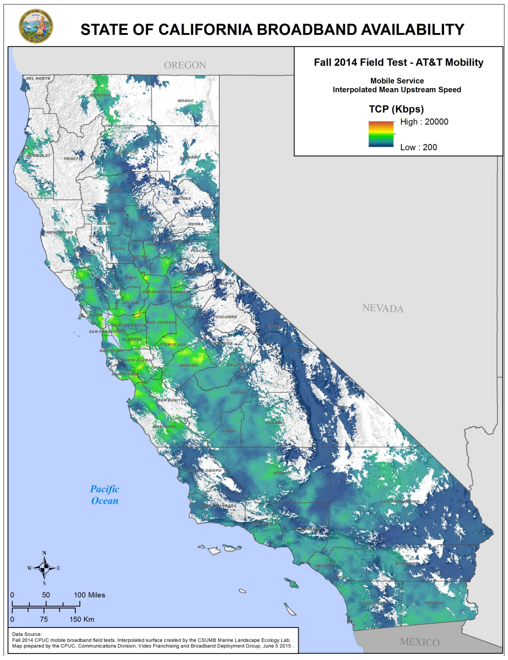

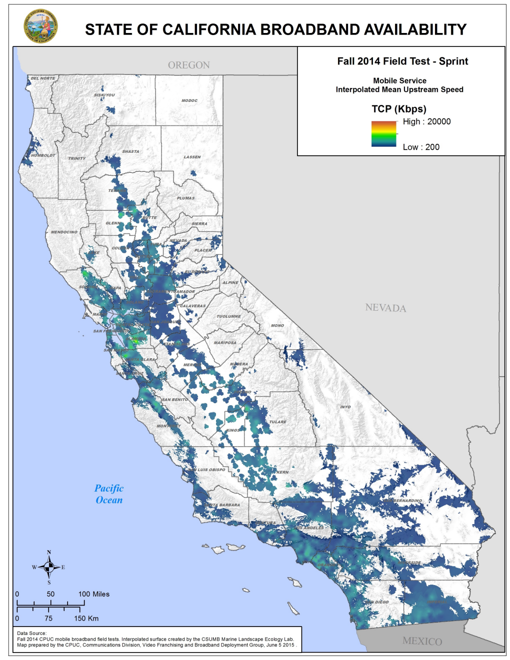

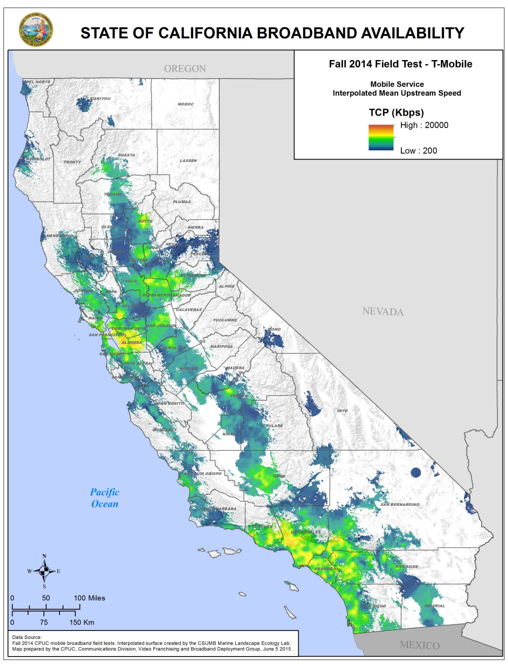

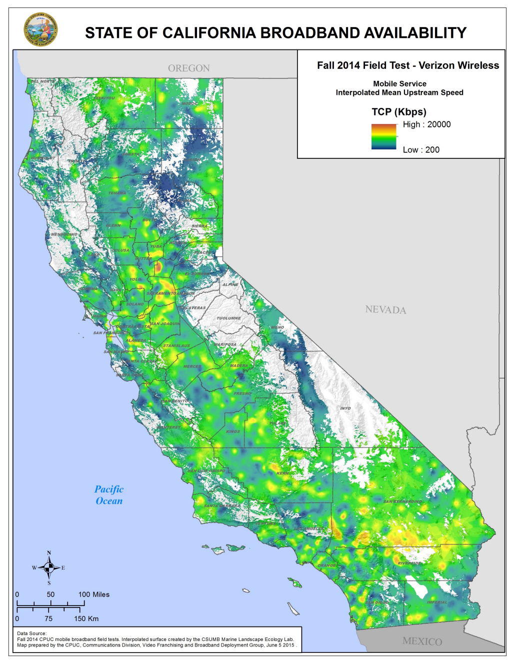

And for latency.

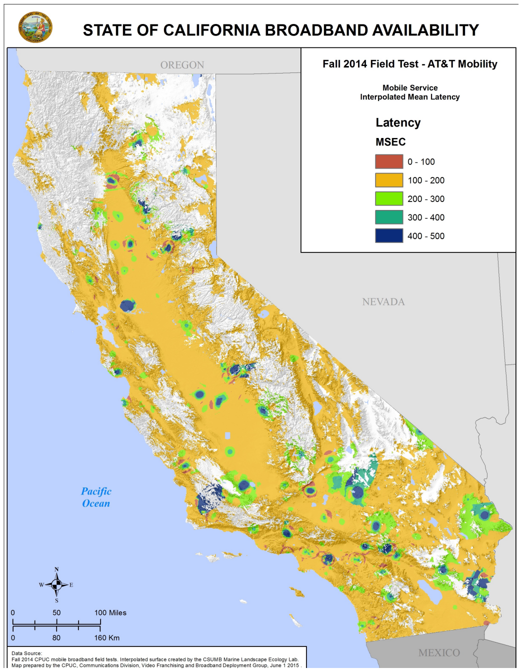

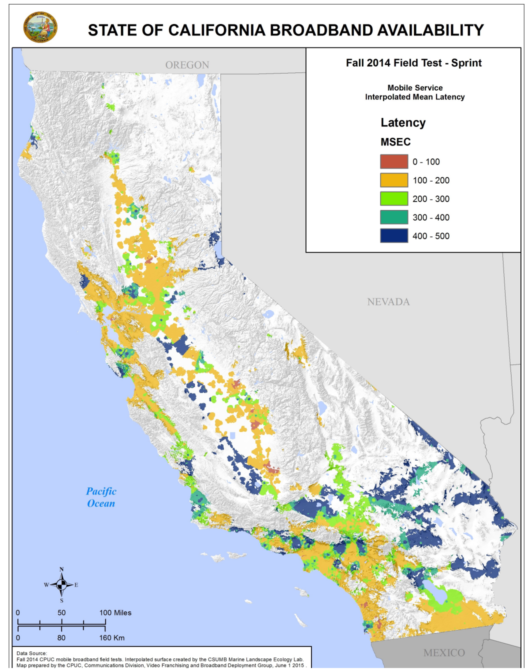

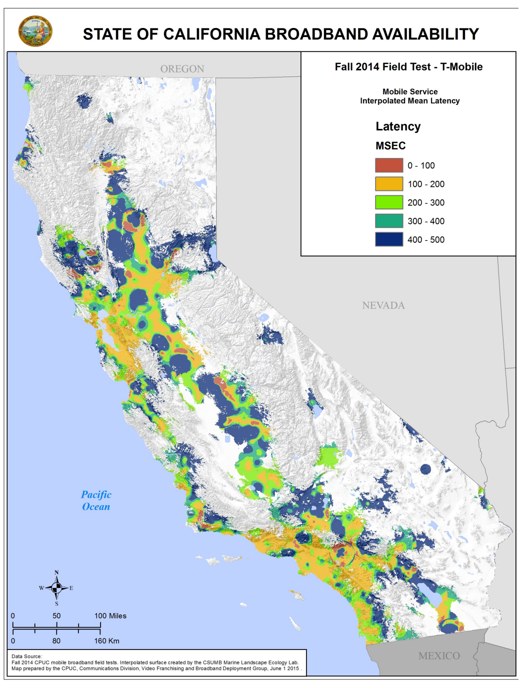

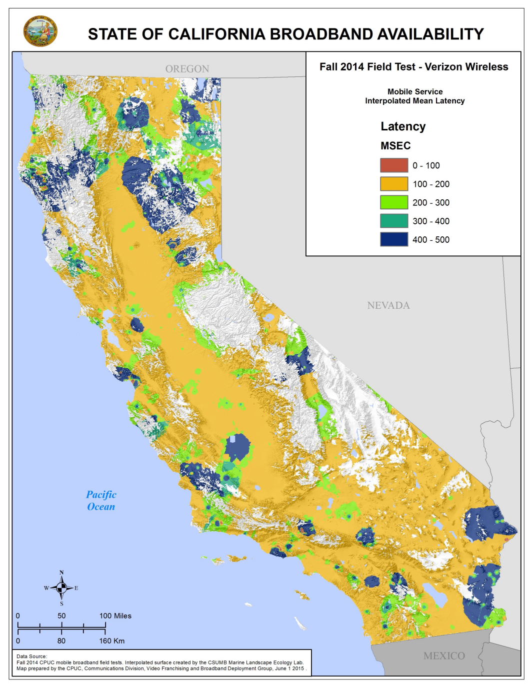

And for the integrated MOS metric, the wide variance across the state is easily seen.

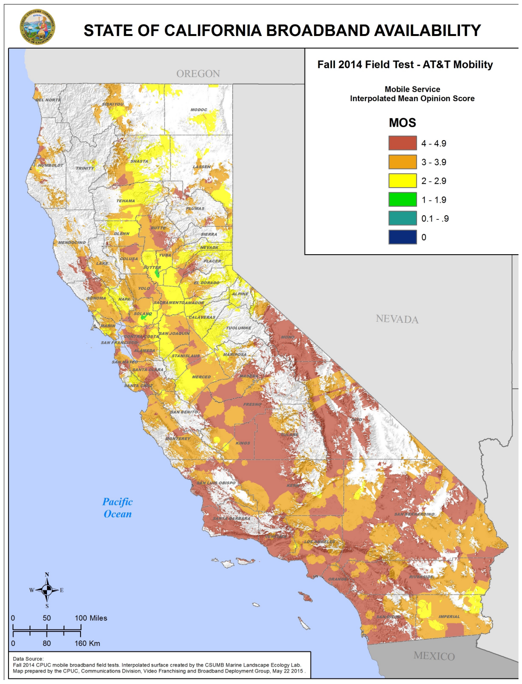

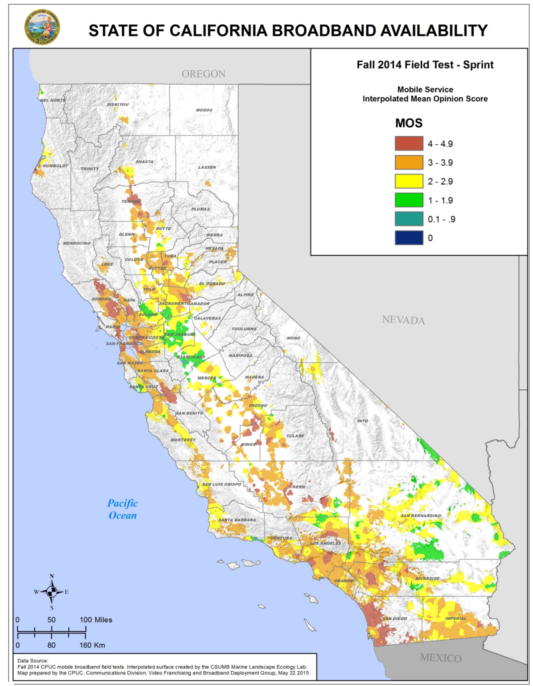

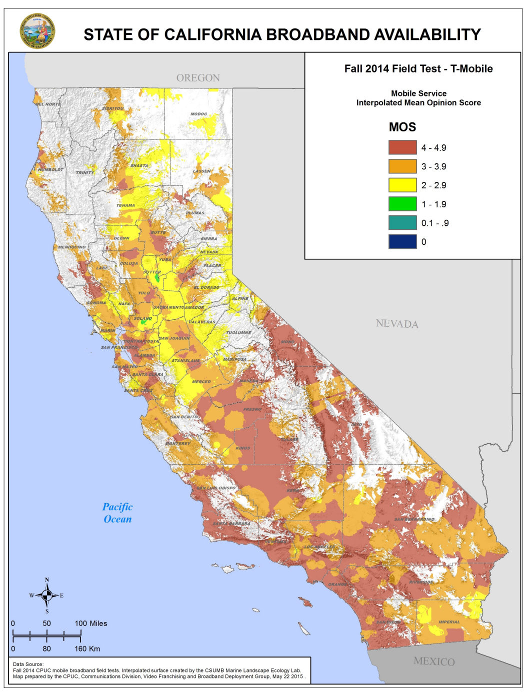

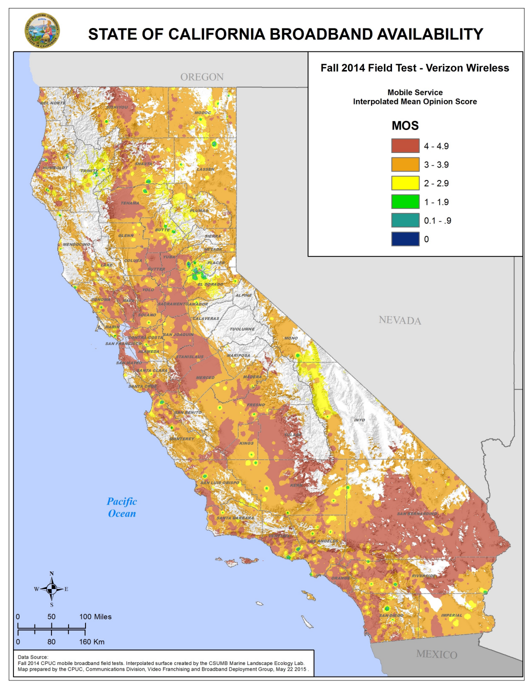

2.2 Choice of Carrier

There is a wide variation between the service delivered by each carrier. This variation is illustrated in the graphs below charting the overall mean downstream throughput for each carrier across the entire state. We can see a range of greater than 4:1 between the fastest and the slowest carriers in both upstream and downstream throughput.

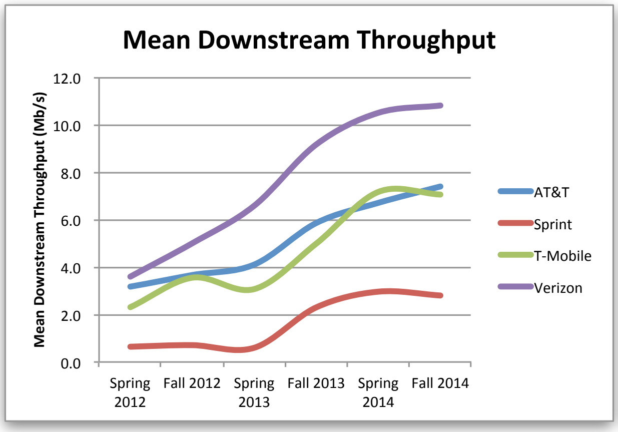

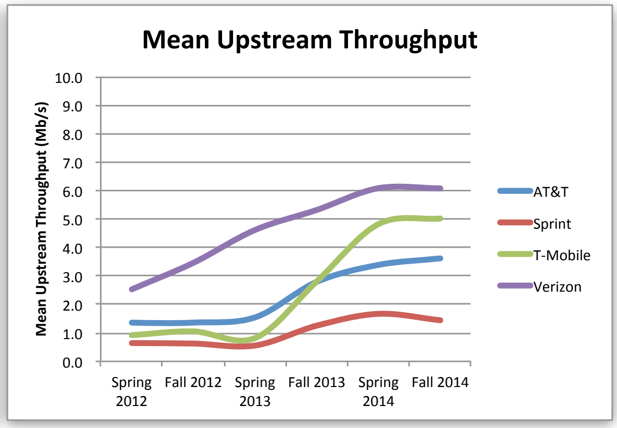

The coverage maps for throughput, latency and MOS in the previous section illustrate the wide variance between carriers in service in California.

2.3 Location of the Service

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

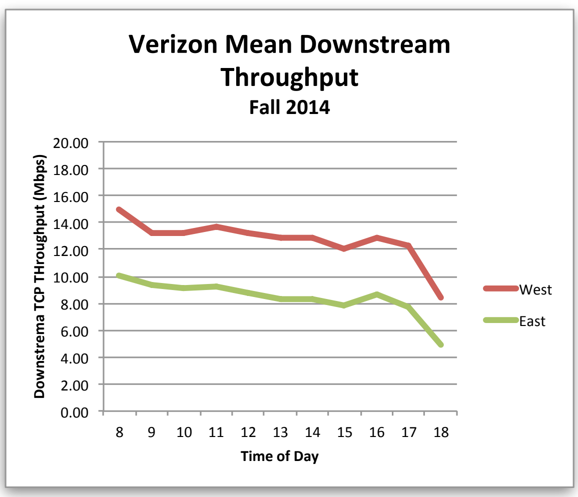

In Verizon's case in Fall 2014, this choice of backbone can result in about a 50% performance difference between a California user accessing a server on the East Coast vs a server in California.

The data suggest that the effects of server location get more pronounced as network performance increases as the data from the other three carriers suggest.

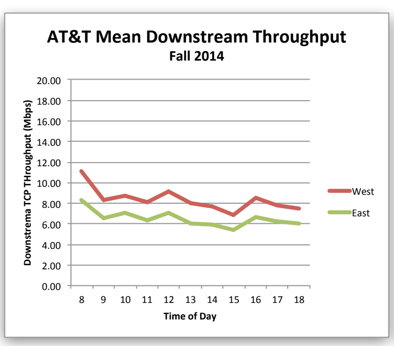

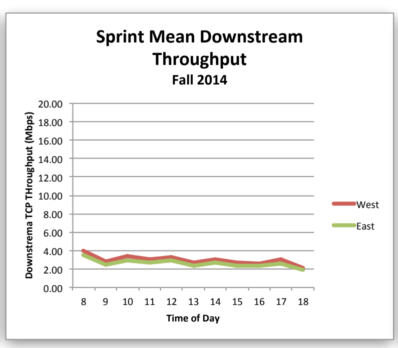

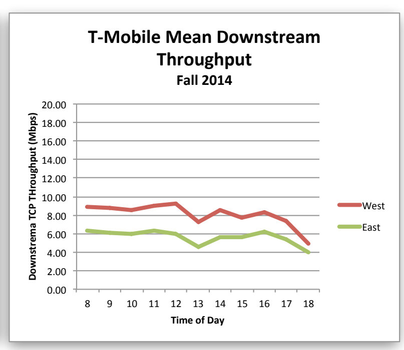

For Sprint, the lowest performing throughput carrier, there is almost no difference in performance due to server location.

2.4 Intra-Session Variation

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

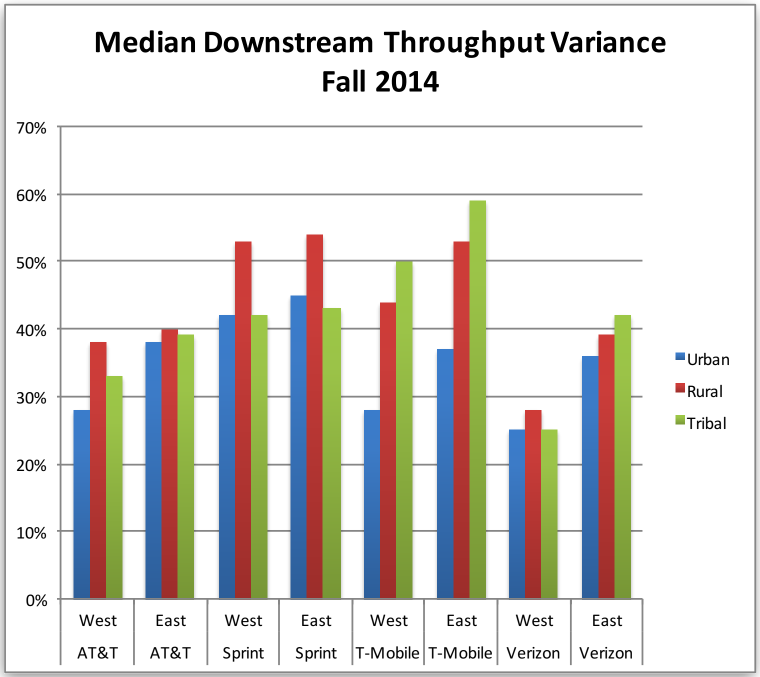

The CalSPEED analysis computes a standard deviation for the variation among these measurements for each test location - giving a metric for the variation in throughput during the duration of the measurement session. This local variation depends on carrier and location - as can be seen to the left charting median variance.

Some general trends can be noted:

- Rural and tribal see median variances higher than urban demographics; and
- backbone Internet contributes variance particularly for AT&T, T-Mobile and Verizon.

Variance during a session of between 25% and 50% can be considered typical.

2.5 Time of Day Variation

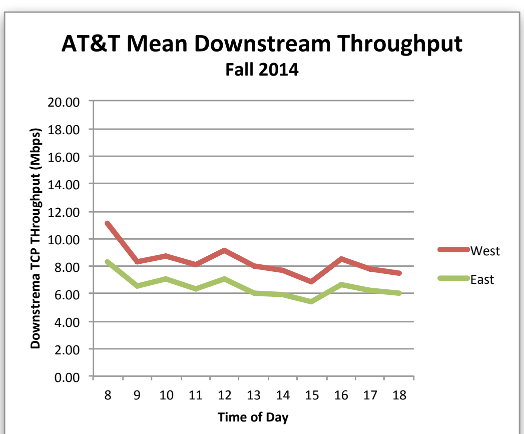

The least important variation is by time of day. Each CalSPEED measurement records the time of day of the measurement. As the chart below demonstrates for AT&T for Fall 2014, the mean downstream throughput shows some variation with time of day, but the variation is on the order of 10% - much smaller than the other sources of variation.

All the carriers show a similar pattern of largely constant average throughput during the day, with a modest decrease from morning towards evening.

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

3. Mobile Broadband Has Stopped Improving

CalSPEED was designed to support comparison over time of network performance. We have tracked four major trends over time: changes in throughput, latency, jitter and service quality. The Fall 2014 data suggests that the capacity and quality of mobile broadband has (at least) stopped improving.

A speculation on this pause in mobile performance improvement might be mobile offered load catching up to network capacity. When offered load approaches or exceeds network capacity, measured performance will stop improving and might begin to degrade if additional capacity is not brought online.

3.1 Throughput

One straightforward summary measurement is the mean across all measurement locations, for both user devices[3] and geographic measurement servers of the downstream and upstream TCP throughput. The following chart documents the change in upstream and downstream TCP throughput by carrier.

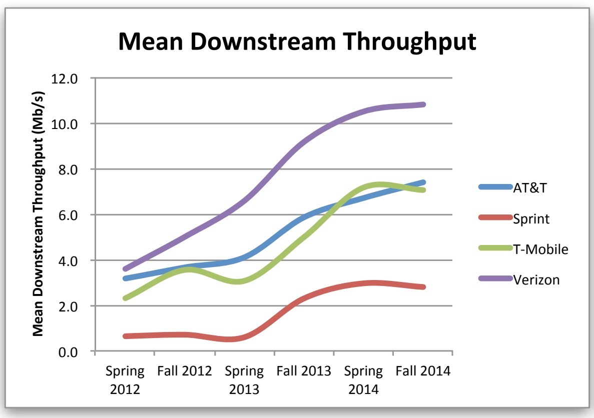

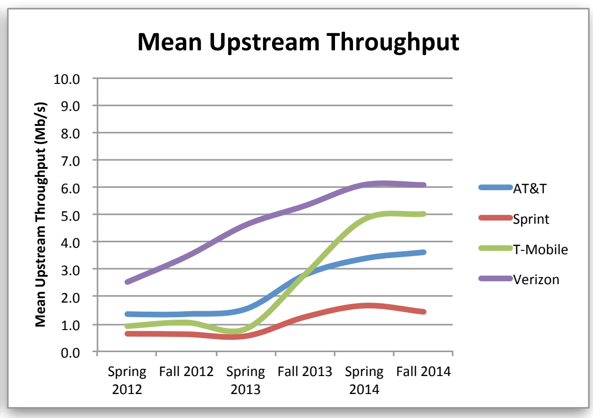

Note that since the last measurement round in the Spring of 2014, performance increase has stalled in the case of Verizon downstream and T-Mobile upstream, decreased in the cases of Verizon upstream, T-Mobile downstream, Sprint downstream and upstream and continued to improve for AT&T downstream and upstream.

3.2 Latency

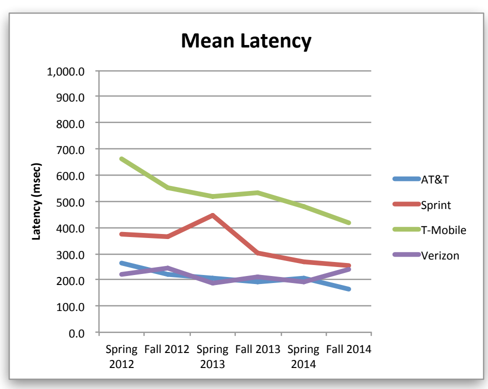

The analysis of overall average latency for each of the carriers shows a similar mixed story as noted for throughput.

The historic trend has been for decreasing latency over time. Since the Spring of 2014 however, Verizon's latency has increased while AT&T, T-Mobile

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

and Sprint continued trends of decreasing latency.

3.3 Jitter

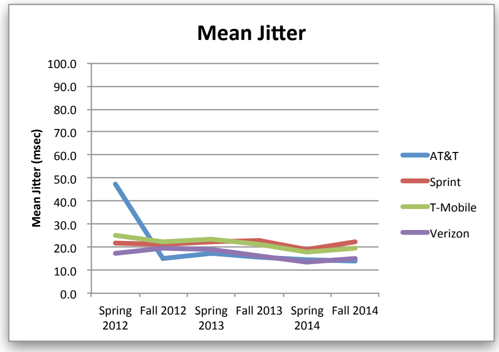

Jitter, with the exception of AT&T, has degraded since Spring 2014.

3.4 Service Quality

CalSPEED has several metrics of quality:

- overall rates of TCP connection failures (percentage of failed TCP connections)
- packet loss; and
- MOS (VoIP Mean Opinion Score)

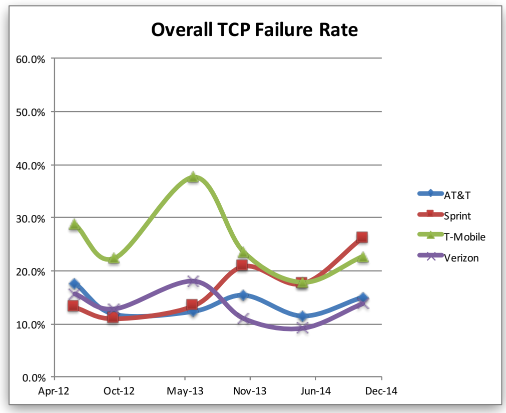

With the Fall 2014 survey, an overall increase in TCP connection failures can be seen for all carriers in California.

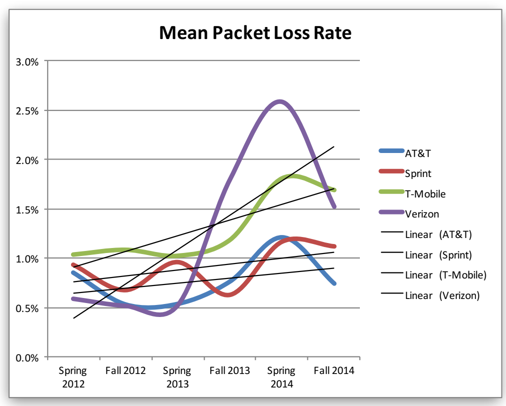

While packet loss rates for Fall 2014 improved for all carriers, a longer term trend towards increasing packet loss continues.

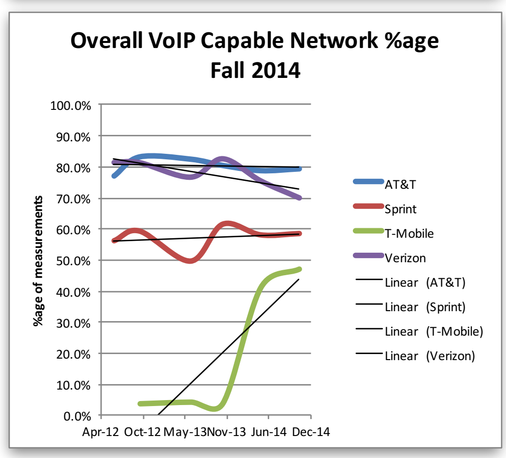

VoIP MOS is a leading indicator of network quality as it integrates packet latency, packet loss and jitter into one metric. While T-Mobile continues a dramatic trend towards increasing network quality, Sprint is stagnant and the two leading VoIP quality networks in California (AT&T and Verizon) show a continuing trend towards overall decreasing MOS quality.

4. Rural Quality Continues to be Materially Lower than Urban

The digital divide between urban and rural continues as measured in Fall 2014.

In the chart to the right we can see that the downstream throughput performance gap continues for all carriers between urban and rural demographics.

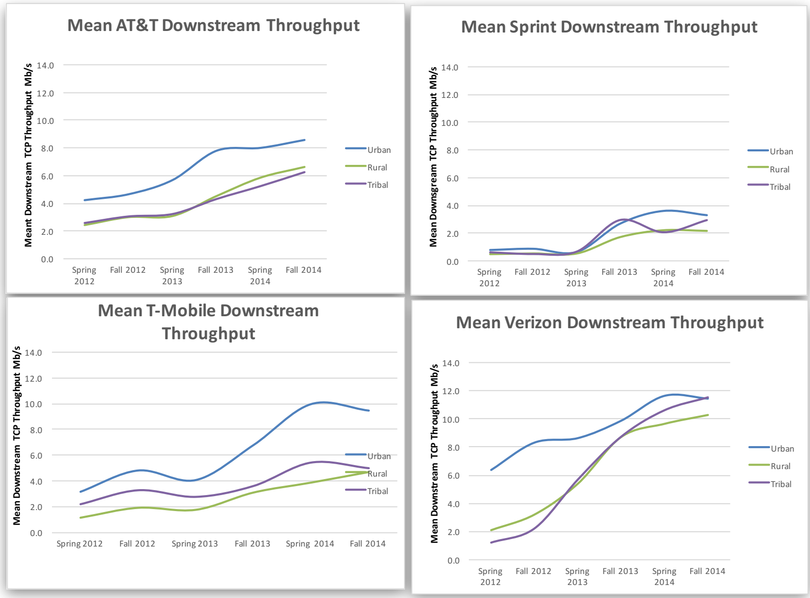

Similarly, latency for rural demographics continues to exceed urban latency for all carriers. For Verizon, the most extensive rural carrier, latency for both urban and rural users increased in Fall 2014.

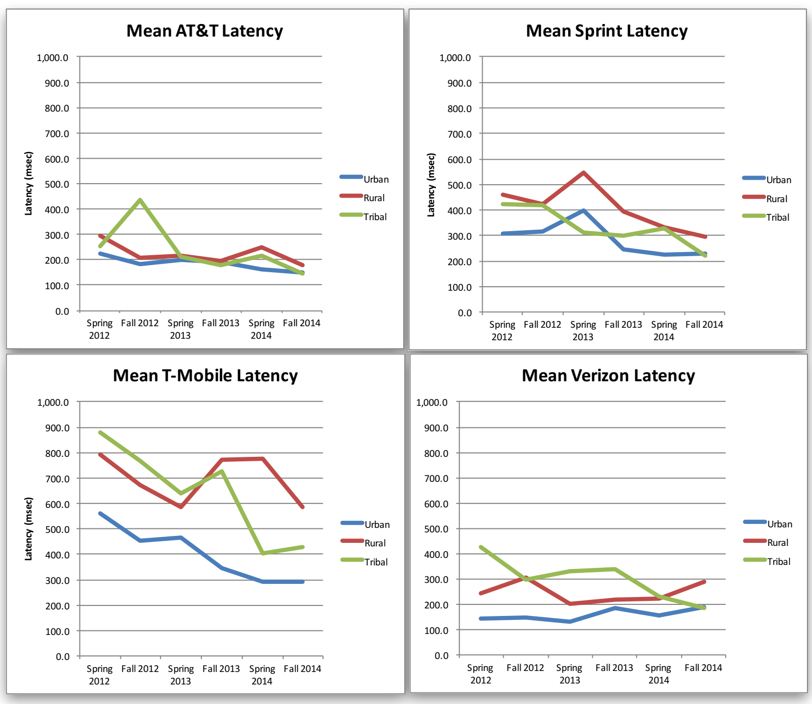

For three of the four networks jitter degraded for rural users more than for urban users. For Verizon, jitter improved for urban users by about 2% but degraded for rural users by over 20%.

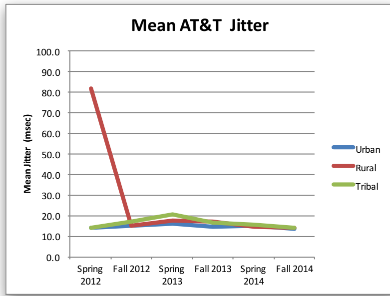

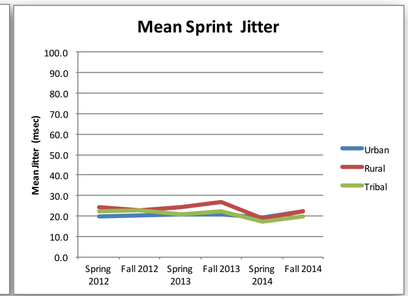

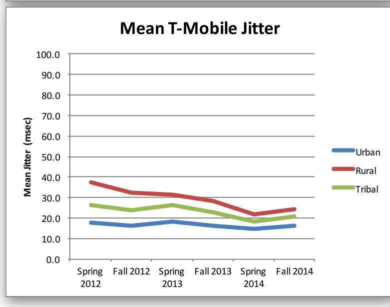

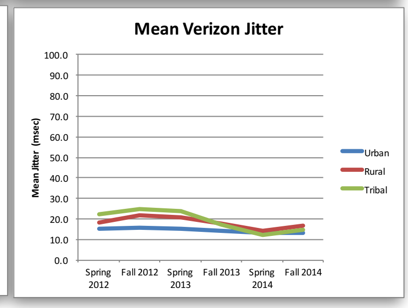

The overall rate of TCP connection failure has increased for all carriers in Fall 2014. However, the connection failure rate has increased more for rural users.

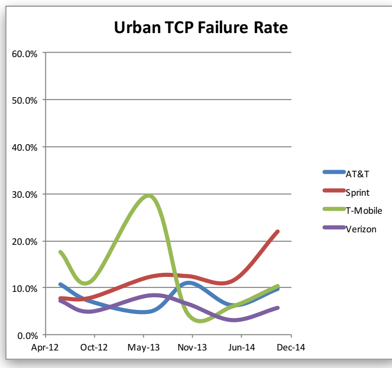

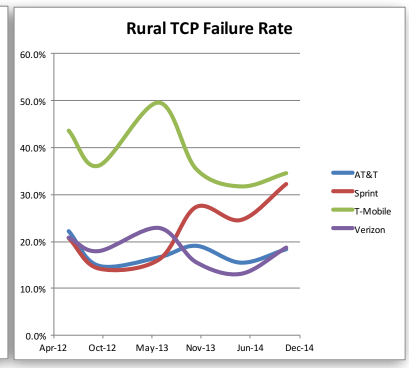

5. Rural Mobile Technology Deployment Slows

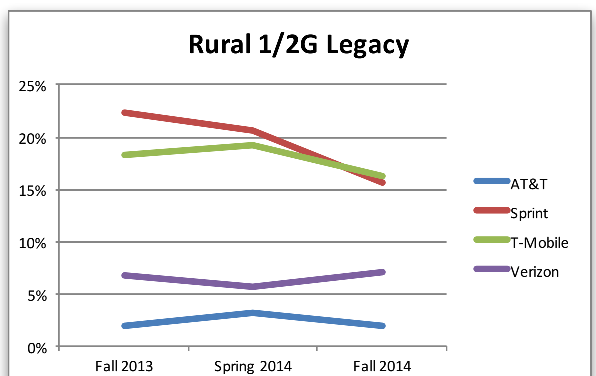

The pace of rural deployment of modern mobile access technology appears to be peaking in California. In the chart to the right, we can see that the percentage of sample locations with 1G or 2G service has stabilized for Verizon (~7%) and for AT&T (~2%) while still high and decreasing for the less deployed networks of Sprint (~16%) and T-Mobile (~17%).

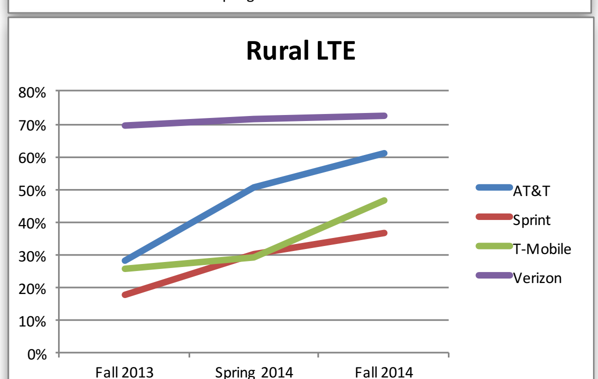

Similarly, the percentage of rural sample locations indicating LTE service has stabilized for Verizon at ~72% while still improving for AT&T (~60%), Sprint (~36%) and T-Mobile (~47%).

These trends suggest that legacy mobile access technology will remain in a minority of locations while LTE deployment will level out without being deployed completely throughout the state.

6. User Device Variation can be VERY Wide

User network experience varies with the specific device used for connection to the network. CalSPEED has used both an Android smartphone and a Windows laptop with a USB modem for each carrier to give some sense of the diversity of user experience.

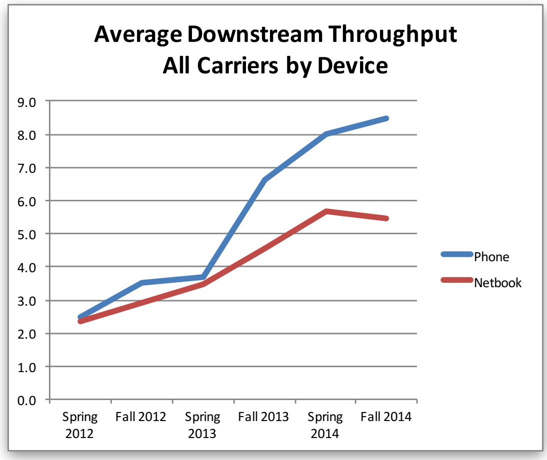

Looking at all carriers, we can see that for the first 18 months of the CalSPEED survey (thru Spring 2013) smartphones and USB modems were increasing in performance at similar rates. However, beginning in the Fall of 2013 and dramatically continuing through the Fall of 2014 smartphones have increased performance at a MUCH higher rate than USB modems - with smartphones now almost 2x the performance of USB modems - on average.

It is unclear as to the cause, other than to speculate that USB modems have fallen behind in technology upgrades - and now do not match the performance in the much more widely used smartphones.

Reflecting this divergence and decreasing market share of USB modems, CalSPEED will be changing devices beginning in Spring 2015 - replacing USB modems for laptops with tablets. We will retain phones as a constant across all survey periods.

7. Broadband Coverage Degrades

CalSPEED measures comparative coverage for each carrier's performance within the announced coverage area that meet current standards for broadband service. The chart to the left below documents the percentage of sample locations, within the announced footprint of each carrier, that meet or exceed the current California standard for sufficient broadband service - 6 Mbps downstream AND 1.5 Mbps upstream. Areas that do not meet this standard are eligible for

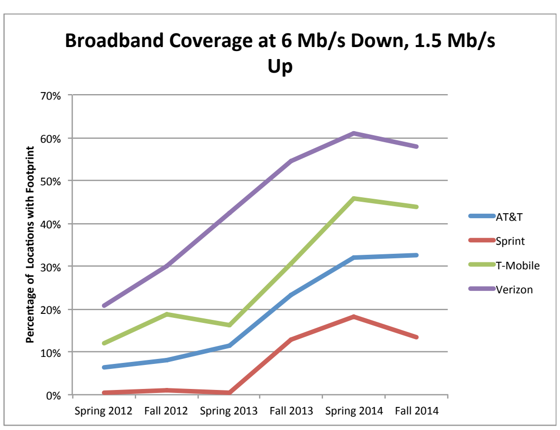

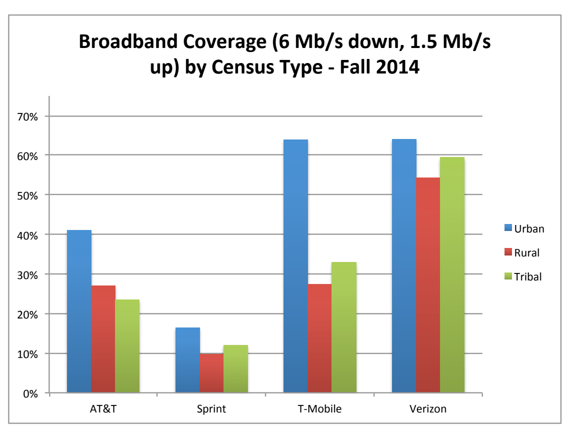

broadband infrastructure subsidies. Note the decrease in broadband coverage in Fall 2014 for

Verizon, T-Mobile and Sprint.   The chart to the right breaks down the Fall 2014 survey by demographic area.   We can see that for all carriers, a higher percentage of urban locations meet the 6/1.5 Mbps standard than do rural locations.

The FCC has determined the federal standard for broadband as 25 Mbps down and 3 Mbps up. The following chart documents the percentage of sample locations in Fall 2014 that meet that federal standard - under 10% for the best carrier Verizon and well under 1% for the lowest quality carrier - Sprint.

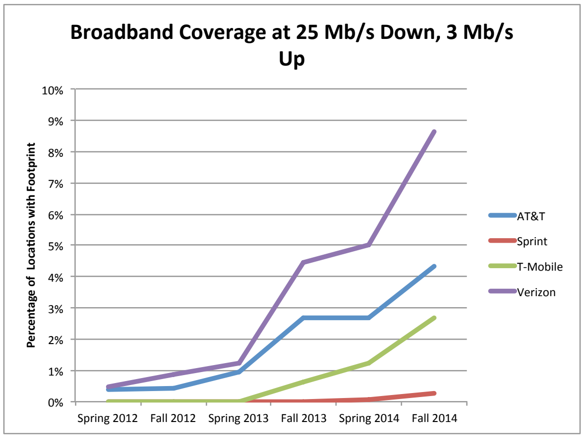

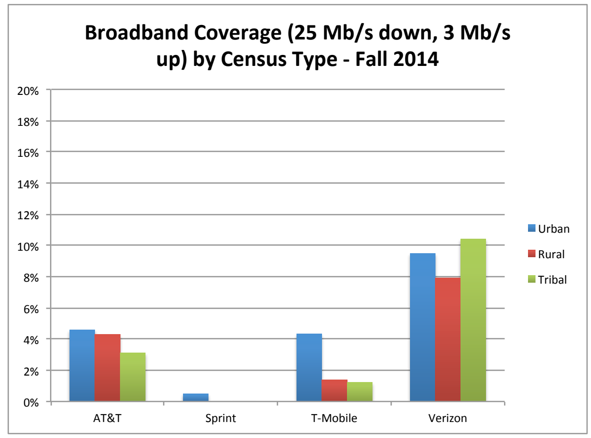

10. Conclusions

This paper has examined the key findings of the sixth measurement rounds of CalSPEED covering 36 months of measurement between Spring 2012 and Fall 2014.   There have been rapid changes during that time and the data provide a solid foundation for five key incremental findings since the Spring of 2014 analysis.

- Mobile broadband's overall performance and quality has stopped improving and shows signs of degradation.
- Mobile broadband continues trends of wide variation across California among carriers, locations of services, the growing digital divide between urban and rural,
- Quality degradation is particularly noticeable in rural areas - in which quality metrics can be 2x worse than in urban demographics.
- Penetration of rural LTE shows signs of stalling.
- There is substantial variation between user devices on the performance and quality of service.

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

How CalSPEED Measures

CalSPEED performs the following sequence of measurements to gather its information:

1. ICMP ping to the West server for four seconds for connectivity checking. If the ICMP ping fails, CalSPEED presumes that there is no effective connectivity to the Internet and records that result.
2. iPerf TCP test (4 parallel flows) to the West server - both downstream and upstream. CalSPEED uses four parallel flows to ensure that the maximum capacity is measured during the test.
3. ICMP ping to the West server for 10 seconds to measure latency to the West server.
4. UDP test to the West server. CalSPEED constructs a UDP stream of 220 byte packets to emulate a VoIP connection with 88kb/s throughput. This UDP stream is used to measure packet loss, latency and jitter.
5. iPerf TCP test (4 parallel flows) to the East server to measure downstream and upstream TCP throughput.
6. ICMP ping to the east server for 10 seconds to measure latency to the East server.
7. UDP test to the East server to measure packet loss, latency and jitter with a simulated VoIP data stream.

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

CalSPEED uses two device measurements - a current smartphone and current USB datastick for laptops. Both are upgraded for each measurement round to match the latest wireless technology deployed by each carrier.

Open Source. CalSPEED is an open source network performance measurement tool that is in turn based on an industry standard open source performance measurement tool - iPerf[4]. iPerf provides the foundation network measurement engine for both the TCP and UDP protocols. CalSPEED packages this engine in both Windows and Android client tools for measuring and recording mobile network performance.

End-to-End User Experience. A foundation assumption of CalSPEED, uniquely among network measurement tools, is an attempt to replicate the end to end user experience. In particular, CalSPEED recognizes that the Internet resources that a typical user accesses are scattered across the entire Internet … and despite the use of content delivery networks to speed Internet performance by caching frequently accessed content, are not always "local" to the user. Many measurement tools focus on evaluating just the local radio access network - the last few miles - and not the backhaul network to the ultimate server resource used. CalSPEED instead tests the complete network path, from the client device, through the local access network, through the Internet backbone, to several ultimate server destinations.

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

CalSPEED emulates this user experience with two fixed servers - one physically located in Northern California and the other in Northern Virginia - both in the Amazon AWS cloud. CalSPEED reports performance both to each individual server and the average between them. Not only does this method measure the different local access methods, but provides a network interferometry that gives insight into the different backhaul strategies chosen by carriers. We find carrier unique 2:1 differences in end to end latency and jitter and material difference in upstream and downstream throughput between the two servers.

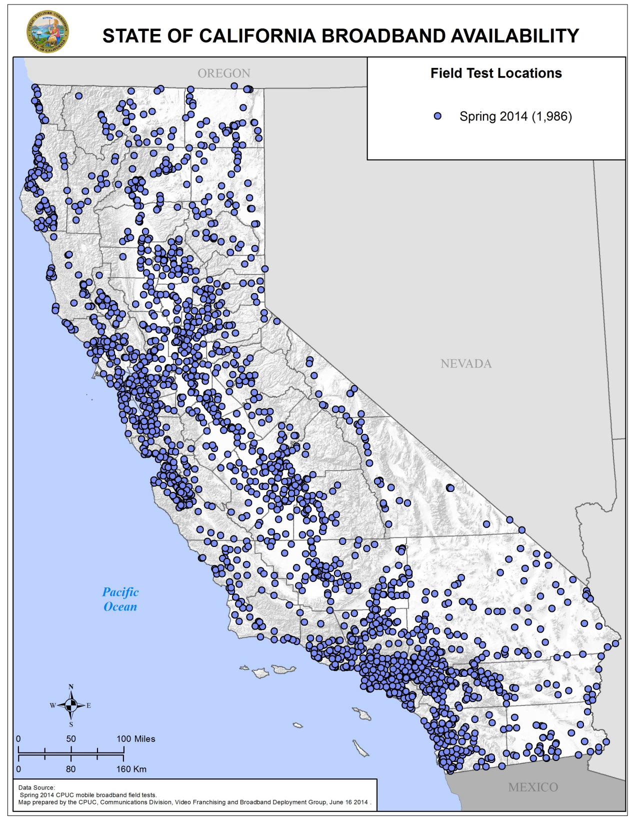

These differences in fundamental network performance illustrate that location matters - Internet performance delivered to the user - the Internet user experience - will vary based on where on the Internet the desired server is located. And desired servers are scattered across the Internet, not just close to every user. Measurement to a local server only results in an overly optimistic expectation of service quality than a typical user will actually experience.

CalSPEED measures a complete portfolio of network metrics including end-to-end packet latency, bidirectional TCP throughput, UDP packet loss and jitter.

<u>Just the Facts.</u> 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.

<u>Not Just for Crowds.</u> 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[5] 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

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

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.

Maps for decision-makers not just for information. We then take the measurement data and create geospatial kriging[6] maps interpolating CalSPEED measurements of (but not limited to) latency, downstream and upstream throughput, jitter and packet loss over the entire state.

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

CalSPEED has now had six rounds of sampling California (Spring 2012, Fall 2012, Spring 2013, Fall 2013, Spring 2014, and Fall of 2014) and is shortly to finish a seventh round (Spring 2015). In each sampling round, we have surveyed the entire state and all four of the major wireless carriers - AT&T Mobility, Sprint, T-Mobile and Verizon Wireless.

[6] http://en.wikipedia.org/wiki/Kriging

Appendix B: Terms

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

Term	Definition
Downstream	The Internet direction from a server to a client.
East Server	Test server located on the East Coast in Northern Virginia
Jitter	The variation in end to end packet latency between user and server.
Kriging	A geostatistical technique for interpolating data from a sample set.
Latency	The end to end round trip delay for a single packet to traverse the Internet from user to server and back.
MOS	Mean Opinion Score. A measurement of VoIP quality
Packet Loss	The rate of loss of packet delivery end to end.
TCP	Transmission Control Protocol. The essential end to end protocol for the Internet that creates a reliable, sequentially delivered byte stream from 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 measures 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