Measuring ISP Performance in Broadband America: A Study of Latency Under Load

Dr. David P. Reed and Dr. Levi Perigo[1]
Computer Science Department
University of Colorado Boulder

Abstract

This study analyzes Latency Under Load (LUL) test data collected by the FCC's Measuring Broadband America program over the past 5 years. To the best of our knowledge, there have yet to be any public studies examining these data. Our analysis of LUL test results found variance in latencies among Internet Service Provider (ISP) technologies with cable and fiber-to-the-home (FTTH) networks performing substantially better than DSL networks. Overall, for the ISPs that we reviewed, cable, FTTH, and downstream DSL LUL is at levels that would not severely impair user experience (i.e., average latency is under 250 ms), though this is not the case for upstream DSL latency under load which would impact user experience with an average LUL of 656 ms. The LUL data also indicated that the worst case, *maximum* latencies measured under traffic load for all of the fixed terrestrial-based ISP technologies would not be adequate to maintain high-quality connectivity for common latency-sensitive Internet applications, such as multiplayer gaming or VoIP. Our longitudinal analysis, however, indicates that LUL data over the past 5 years has improved at an average rate of about 10% per year over cable networks due to advances in protocols and increased service speeds. The test results even provide an example of how one such protocol advance – the use of Advanced Queue Management (AQM) over a DOCSIS cable network – resulted in an immediate downstream latency reduction of 48% after its deployment.

Given the real-world nature of the LUL performance test with regard to network traffic loading as compared to idle latency measurements, further research attention and analysis of latency performance under traffic load is warranted given the increasing importance of broadband network performance and service quality on societal policy objectives.

Introduction

Broadband availability and performance has garnered increasing public attention over the last decade as governments recognize the substantial benefits of ubiquitous broadband infrastructure conveyed upon modern society. Universal, high-quality Internet access creates enhanced economic development and public safety for communities, improved health care and educational opportunities, and better quality of life, to name but a few of these benefits. Recognition of the critical role this infrastructure now plays in society has led most nations to adopt universal service programs to subsidize the deployment of broadband networks in underserved or unserved areas to bridge the Digital Divide and the problems that it causes. Along with this growing government support and in recognition of the important societal role it plays, some government regulators have created programs to gather information regarding the performance of broadband networks to monitor that their constituencies are capturing the full benefits of this critical infrastructure. The Federal Communications Commission (FCC) has been

[1] This research was supported from a gift of the Comcast Innovation Fund.

supporting one prominent effort since 2011 to gather data on broadband performance throughout the United States. This effort, known as the Measuring Broadband America (MBA) program, has published 10 reports over this time. [2]

An important focus of the MBA program has been to collect data associated with the speed of broadband services offered by Internet Service Providers (ISPs) to insure actual, offered speeds are consistent with advertised speeds pitched to consumers by the ISPs. This central focus is understandable, as the FCC's definition of broadband relies entirely upon the speed of the service and does not include other service attributes such as latency, packet loss, or transmission technology. Yet as the ISP market has matured, and the role of broadband to achieve societal objectives associated with information access have emerged, the potential importance and role of additional broadband service attributes beyond speed are now receiving additional attention. [3]

Network latency is one of the broadband service attributes that is receiving increasing attention as it has a significant impact on application performance. Latency has been an accepted performance metric of interest since the inception of modern data communications based on digital technology. Stated briefly, network latency represents the time that it takes for data packets to travel from one network host to another network host. Since data packets cannot instantaneously be sent from the source to its destination, the network latency metric provides a measure of the total delay experienced by the packet as it is transmitted through many different network nodes in order to arrive at its intended destination. The FCC MBA program has measured the network latency encountered on ISP networks in baseline network conditions (generating what is known as the "idle" latency of the services) and regularly reported these results in its reports.

The problem with the idle latency measurements is that they fail to reliably measure network performance when reasonable traffic loads are present on the network, which can lead to misleading characterizations of the network delays present on the network when it is operating under normal peak load conditions. Latency measurements taken when significant network traffic is present provides a more realistic measure of network delays. The FCC MBA program actually does collect data for network latency under load (LUL), though they have yet to include any discussion of these measurements in the current series of MBA reports.

Given the current setting of heightened interest in broadband performance to meet policy objectives, this paper examines the LUL data collected by the FCC's MBA program to review how ISPs are performing under traffic load conditions. Latency under load has been an overlooked performance metric collected by the FCC that could shed interesting insights into the performance of different broadband ISPs included in the MBA program.

A full analysis of the LUL data for all the ISPs that has been collected by the MBA program falls beyond the scope of this paper. Instead, we will analyze the LUL data for a subset of the ISPs with data collected within the MBA program to identify important variations by ISP in network LUL performance by local access network technology employed (e.g., DSL, cable, and Fiber-to-the-Home) and the service speed actually provided to the ISP customer (e.g., actual

---

[2] See Measuring Broadband America. (2011, July 28). Federal Communications Commission at https://www.fcc.gov/general/measuring-broadband-america.

[3] For a discussion of the broad implications of the Definition of Broadband for the purposes of regulation, see David P. Reed, "Is Speed Enough? Examining the Definition of Broadband and Its Implications for Public Policy", *forthcoming*, 49th Annual Telecommunications Policy Research Conference, September 22-24, 2021.

download speeds between 10 Mbps – 1 Gbps or actual upload speeds between 2 – 45 Mbps). We also show how LUL data has changed over time for one cable ISP and use this historical view to explain how technical innovations can improve network performance over time.

The outline of the paper is to begin with a brief review of the MBA program database used for this analysis. We then review our research methodology, followed by the results of our analysis based upon implementation of this research approach. A conclusion section summarizes the implications of our study.

FCC's Measuring Broadband America Data

We rely upon data collected and provided by the FCC as described and downloaded from the FCC MBA website.[4] The MBA program employs the SamKnows hardware platform to execute different tests designed to measure key performance indicators (KPIs) of a broadband connection. We will not go into depth in this paper describing the details of the methodology and platform employed by the FCC to collect its KPI data. A detailed explanation of the FCC's approach can be found in the Technical Appendix of the 10th MBA Report (*Technical-Appendix-Fixed-2020.Pdf*, n.d.).[5] With the exception of using some selected data from prior years needed to generate a historical view of trends in LUL values over time, all data reported in this paper was gathered by the FCC during 2020 assuming continuation of the methodology as described in the Technical Appendix.

To gather test data, the MBA system employs 54 measurement servers (operated by M-Lab, Level 3, and StackPath) that are geographically distributed across the core network to run the necessary tests including the LUL test which is of interest for our study. These servers measure broadband performance between consumers' locations and the measurement servers giving the lowest round-trip times (RTT) to the consumers' network addresses. The result provides "off-net" test data since these servers are located outside the network boundaries of the participating ISPs.

The MBA program also collects data from measurement servers operated by participating ISPs to provide additional insight into broadband service performance within an ISP's network. These servers generate "on-net" test data since they are located inside the network boundaries of the participating ISPs. While the MBA raw data includes both off-net and on-net testing results, we only present off-net results in this paper to ensure a consistent approach for gathering data measurements.

The LUL test reported in the MBA raw data reflects the average RTT obtained by sending a series of evenly spaced User Datagram Protocol (UDP) packets between the consumer location and the test server in both the downstream and upstream directions. The test occurs at the same time that upstream and downstream speed tests are running and measures throughput by performing multiple simultaneous HTTP Get and HTTP POST requests. The MBA upstream and downstream speed tests use three concurrent TCP connections (and therefore three concurrent HTTP requests) to ensure that the connection is saturated or under load. The simultaneous UDP

---

[4] See Measuring Broadband Raw Data Releases—Fixed. (2019, December 5). Federal Communications Commission at https://www.fcc.gov/oet/mba/raw-data-releases.

[5] See Technical-Appendix-fixed-2020.pdf. (2020). Retrieved July 14, 2021, at https://data.fcc.gov/download/measuring-broadband-america/2020/Technical-Appendix-fixed-2020.pdf.

LUL test runs over the same 10 seconds of the speed tests and measures RTT (recording the mean, minimum, and maximum) and number of packets lost. UDP packets are spaced 500 milliseconds (ms) apart during the test. As stated in the Technical Appendix, the standard test schedule calls for the LUL test to run from a single off-nest test node every hour of the first three days of the month, generating a volume of 5.8 Mbytes of daily test traffic.

The raw data for downstream and upstream LUL measurements is contained in the data files named *curr_dlping.csv* and *curr_ulping.csv*, respectively. Both of these data files include the following schema of row level results relevant for use in our analysis:

- **Unit_id** – Unique identifier for an individual unit or volunteer
- **dtime** – Time that test finished
- **target** – Target hostname or IP address
- **rtt_avg** – Average round-trip time
- **rtt_max** – Maximum round-trip time

In order to compare LUL data to service speeds, our study also will draw upon raw data from the download and upload speed tests described earlier that measure the download and upload throughput by performing multiple simultaneous HTTP GET and HTTP POST requests to a target test node.

The raw data for download and upload measurements, assuming use of IPv4, is contained in the data files named *curr_httpgetmt.csv* and *curr_httppostmt.csv*, respectively. Both of these data files include the following schema of row level results relevant for use in our analysis:

- **Unit_id** – Unique identifier for an individual unit or volunteer
- **dtime** – Time that test finished
- **target** – Target hostname or IP address
- **bytes_sec** – Running total of throughput, which is sum of speeds measured for each stream (in bytes/sec), from the start of the test to the current interval

The next section of the paper will discuss how these data files are used for our analysis.

Finally, in order to examine how LUL varies across local access network technologies we focused our analysis on the MBA LUL data collected from three ISPs, each representing a different network technology. The three ISPs that we selected for this study are CenturyLink to represent DSL, Verizon Fiber to represent fiber-to-the-home (FTTH), and Comcast to represent cable. Table 1 shows some relevant statistics associated with these three ISPs relative to their participation in the MBA testing program.

ISP	Tech-nology	Sample Size	% of All Volunteers	# of On-Net Test Nodes	Identified Speed Tiers During Tests
CenturyLink	DSL	571	19.5%	14	Download Speeds: 1.5, 3, 7, 10, 12, 20, 40 Mbps Upload Speeds: 0.768, 0.896, 2, 5 Mbps
Comcast	Cable	276	9.4%	37	Download Speeds: 60, 150, 250 Mbps Upload Speeds: 5, 10 Mbps
Verizon Fiber	FTTH	177	6.0%	2	Download Speeds: 75, 100 & 1000 Mbps Upload Speeds: 75, 100 Mbps

Table 1: Participation of Selected ISPs in MBA Testing Program

Research Methodology

This section describes the research methodology that we use to perform our analysis of the LUL data collected in the FCC's MBA program. Our approach is to apply a data cleaning process to the LUL data in order to provide graphical visualization of these measurements as a function of time, specific ISP, and technology type.

As described in the last section, our analysis utilizes the raw data from two of the performance tests included in the MBA program: the LUL test and the download/upload speed test. For our analysis of raw LUL data, Table 2 below describes the simplified approach we have implemented for first cleaning the data to remove artifacts such as zero entries in the data and then filtering the data by ISP.[6] Our filtering approach excludes on-net measurements by removing all data entries associated with target names associated with the ISP. For example, the target name of on-net data for Comcast include a URL name that includes the word Comcast. Once the data is filtered and cleaned, we then calculate monthly averages of the maximum and average RTT data for each ISP under study.

Step	Task	Description
1	Download raw LUL data files from FCC MBA web site	• Curr_dlping (per month) for downstream • Curr_ulping (per month) for upstream
2	Order **RTT_avg** and **RTT_max** data fields from lowest to highest values	• Remove all zero values and obvious outliers (we found zero entries were ~1% or 8K-14K of the entries)
3	Filter data by selected ISP and off net (excluding on-net data)	• Comcast for cable • CenturyLink for DSL • Verizon for FTTH
4	Calculate metric	Average all **RTT_Avg** and **RTT_Max** data in a month for each ISP

Table 2: Data Cleaning and Processing of LUL Data

Table 3 below describes the simplified approach we have implemented for cleaning and filtering the raw speed test data. Again, our filtering approach excludes zero value data entries and on-net measurements. Note here that our mapping of LUL data to speed test data is constrained to a single ISP. Once the data is filtered and cleaned, we then calculate average monthly throughput for each User ID (i.e., every volunteer subscriber with the ISP) present in the data set.

---

[6] More specifically, the filter of the raw file by "on-net" ISP began with a find/replace of all targets with "**ISP_name*", leaving the file with only entries with a target of *ISP_name*, which generated a list of all of the ISP's on-net user-IDs (UIDs) in raw data file. Next, the entire "unfiltered" data was sorted by the ISP UIDs (generating both on-net and associated off-net entries). This result was filtered again by removing all the "on-net" entries, leaving a file with only the ISP UIDs that were off-net. This process was repeated for each ISP included in the study for each month of both upload and download data available.

Step	Task	Description
1	Download raw speed test data files from FCC MBA web site	• Curr_httpgetmt (per month) for download speed tiers • Curr_httppostmt (per month) for upload speed tiers
2	Order **bytes_sec** data field from lowest to highest values	• Remove all zero values • Sort data by User ID
3	Filter data by selected ISP and off net (excluding on-net data)	Comcast for cable
4	Calculate metric	Average all **bytes_sec** data in a month for each User ID

Table 3: Data Cleaning and Processing of Speed Test Data

Before conducting the analysis, the final step in our methodology is to merge the LUL and speed test data by common User ID so that we can associate the LUL data with speed test data for each User ID (volunteer). At this point the data is now ready for presentation and analysis. Note one delimitation of our study is that we did not further clean the data to identify whether a particular user changed speed tiers during the course of the year. The full cleaning process employed by the FCC MBA program to generate "verified" data includes a step to remove data from those User IDs where a change in speed tier can be identified.

Latency Under Load Analysis

This section presents the results of our analysis of the LUL data collected in 2020 by the MBA program.

We begin by first looking at the downstream and upstream LUL data by technology as shown in Figures 1 and 2, which plot LUL versus average monthly RTT during 2020. Recall that the DSL, cable, and FTTH technology plots correspond to data collected from CenturyLink, Comcast, and Verizon Fiber, respectively. MBA LUL data in 2020 from January is not available, nor from September and October of 2020 until the FCC completes its next MBA Report based upon this data.

There are four interesting observations regarding these results. First, there is a substantial difference in the magnitude of downstream and upstream LUL for all technologies. The average LUL for all data available in 2020 (averaging across all months shown) is 149 ms/356 ms in the downstream/upstream. The reasons for this disparity are slower upstream service speeds and use of multiple access protocols in the upstream that can cause network delays.

Second, there is substantial variation in downstream and upstream LUL by technology type. The average downstream cable LUL (103 ms) and FTTH LUL (148 ms) is 91 ms (or 47%) and 46 ms (or 24%), respectively, lower than the average DSL downstream LUL of 194 ms. The difference is even larger in upstream LUL where the average upstream cable LUL (176 ms) and FTTH LUL (207 ms) is 509 ms (or 74%) and 478 ms (or 70%), respectively, lower than the average DSL upstream LUL of 686 ms. The lower service speeds of DSL are likely the major contributor to this discrepancy. The results also show total average cable downstream LUL to be 45 ms lower than FTTH. The explanation for this difference is unclear, and we do not read much into this difference. We note that the FCC did remove Verizon users with 1 Gbps FTTH service speeds in the last MBA report due to potential concerns regarding the integrity of the MBA performance tests at a service speed of 1 Gbps.

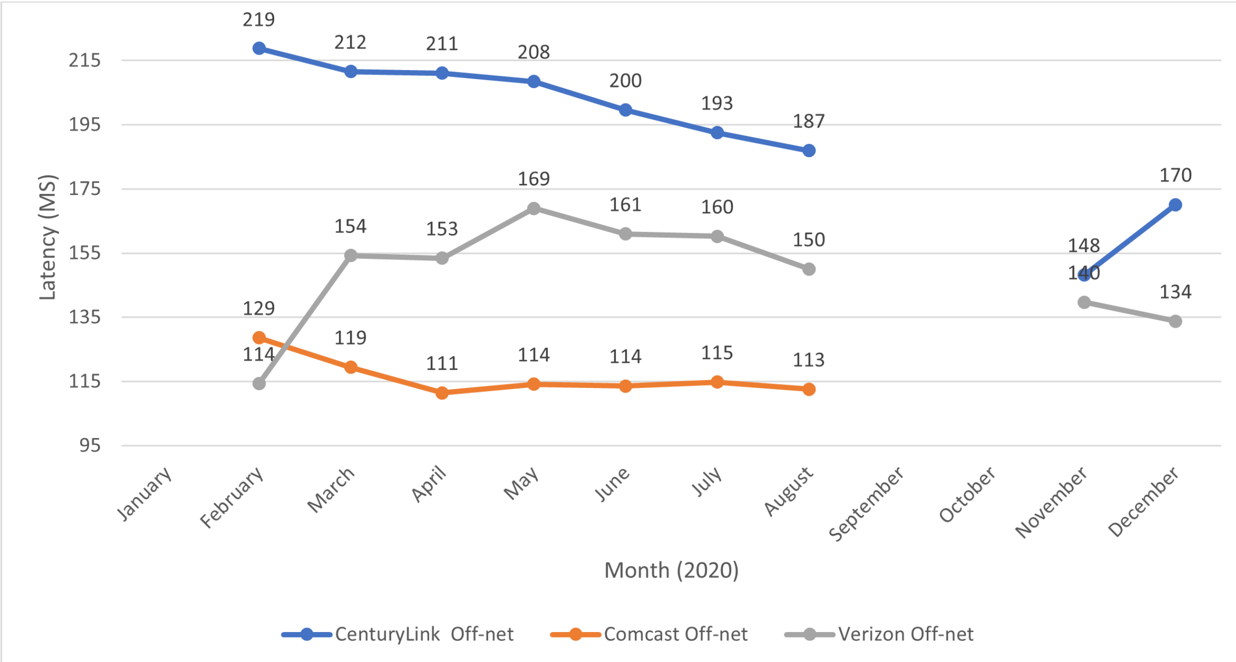

Figure 1: Downstream Average Latency Under Load By Technology (2020, September/October Not Available)

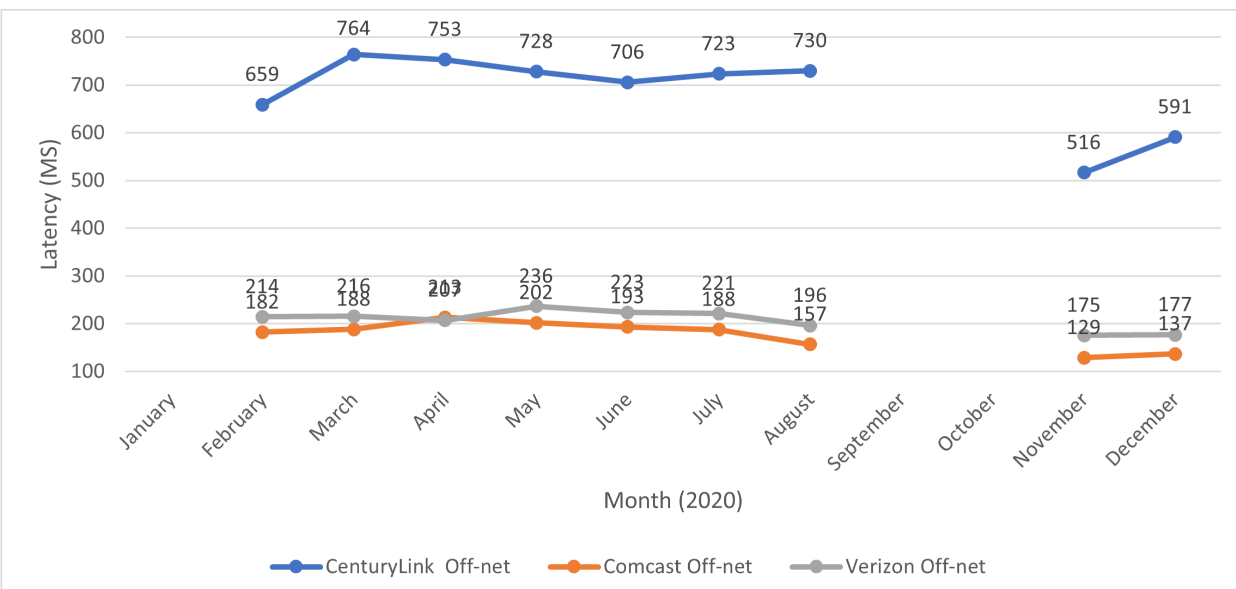

Figure 2: Upstream Average Latency Under Load By Technology (2020, September/October Not Available)

Third, the overall magnitude of the LUL confirms the efficacy of the loading associated with the test since LUL values are much higher than the idle latency measurements included in the MBA program. While reporting the values of idle latency in the MBA data for 2020 are beyond the scope of this study, the prior MBA Report based on 2019 data reported “Cable latencies ranged between 16ms to 28ms, fiber latencies between 5ms to 11ms, and DSL between 21ms to 61ms."[7] All of these ranges include latency measurements substantially below the LUL measurements. Moreover, using a latency of 250 ms as a threshold above which user experience concerns increase significantly due to notable deterioration in the quality of real-time applications such as multiplayer gaming or VoIP, the network performance of cable, FTTH, and downstream DSL under load is likely to be generally acceptable. The LUL latency of upstream DSL at 656 ms, however, is likely a reflection of service degradation that would be notable to DSL users of real-time services during periods of peak load.

Fourth, the downstream LUL values for the cable technology (Comcast) show a substantial decline between August and November of 2020. In fact, the monthly average of downstream LUL decreases by 48% if we compare the average monthly LUL between February – August against data between November – December. The explanation for this substantial improvement in network performance, based upon information provided to us directly from Comcast technical staff, is the launch of Active Queue Management (AQM) technology in the Cable Modem Termination System (CMTS) equipment of the ISP in the August/September time frame which was fully functional by November.[8] This snapshot in time provides a unique view of the positive impact that new innovations in IP protocols and equipment can have on network performance. In this case, improvements in how the CMTS routing manages packet buffers under load resulted in a decrease of almost 50% in LUL performance.

Figures 3 and 4 plot the downstream and upstream LUL performance by technology versus average monthly maximum RTT (rather than average monthly RTT shown in Figures 1 and 2). Comparing the average versus maximum RTT provides insight into the range of variability that is possible for network delay under higher traffic load conditions. The overall maximum LUL for all data available in 2020 (averaging across all months shown) increased to 244 ms (up 65% from 149 ms) in the downstream and 561 ms (up 57% from 356 ms) in the upstream compared to the average RTT measurements. As user experience concerns increase with higher latency, particular above 250 ms, this result provides a measure of the upper bound on latency that is likely to be encountered in the normal operations of these ISPs and its likely impact on user experience associated with real-time applications.

These measurements reflect similar variations in downstream and upstream LUL by technology type shown earlier, though the spread decreased a bit as the average maximum downstream/upstream LUL for cable (up 85%/97%) increased more than DSL (up 53%/47%) or FTTH (up 66%/59%). With the overall average maximum RTT for downstream/upstream LUL of 191/347 ms for cable, 296/1006 ms for DSL, and 246/329 ms for FTTH, all users are likely to experience performance issues of time-sensitive applications during a maximum LUL event.

---

[7] See Measuring Fixed Broadband—Tenth Report. (2021, January 4). Federal Communications Commission at https://www.fcc.gov/reports-research/reports/measuring-broadband-america/measuring-fixed-broadband-tenth-report

[8] For a detailed description of the new Proportional Integral Controller Enhanced (DOCSIS-PIE) AQM technique implemented by Comcast, see Allen Flickinger, Carl Klatsky, Atahualpa Ledesma, Jason Livingood, Sebnem Ozer. *Improving Latency with Active Queue Management (AQM) During COVID-19.* (2021). Retrieved July 30, 2021, from https://arxiv.org/pdf/2107.13968.pdf .

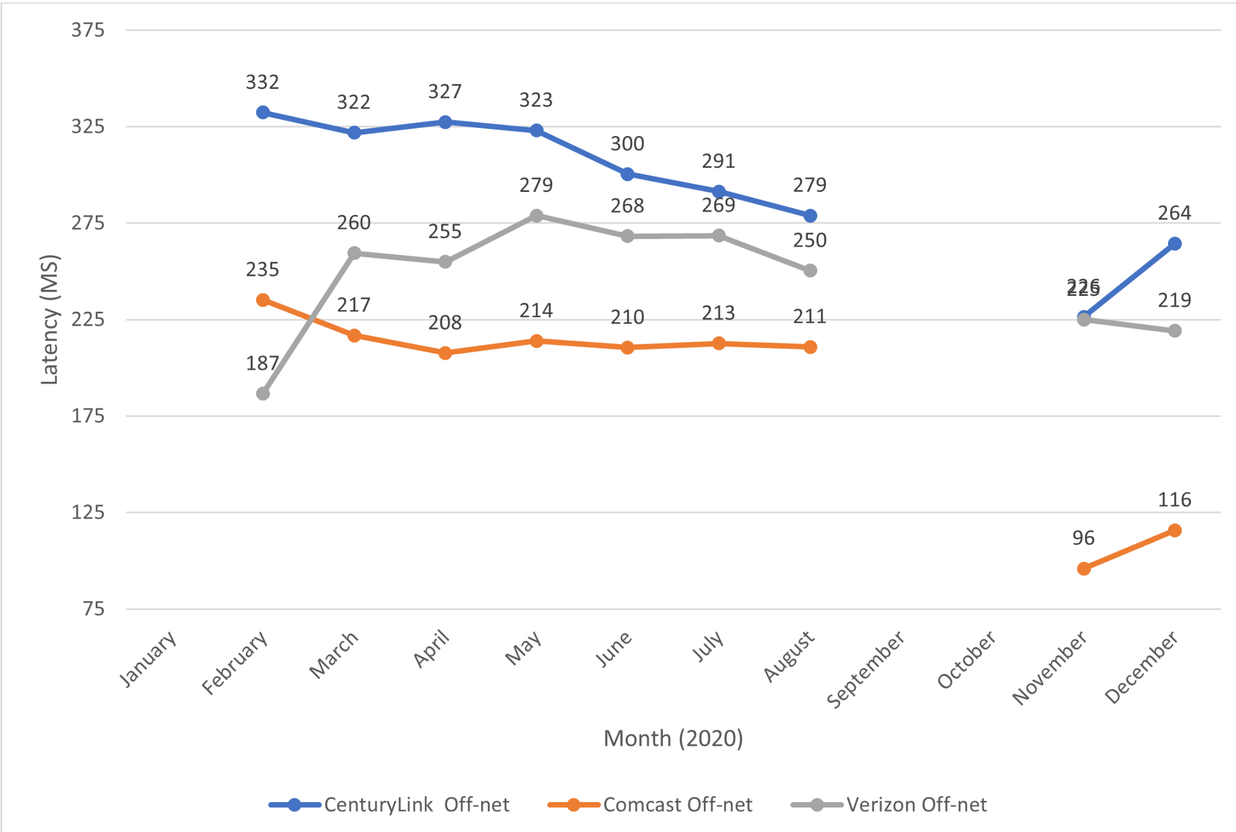

Figure 3: Downstream Average Maximum Latency Under Load By Technology (2020, September/October Not Available)

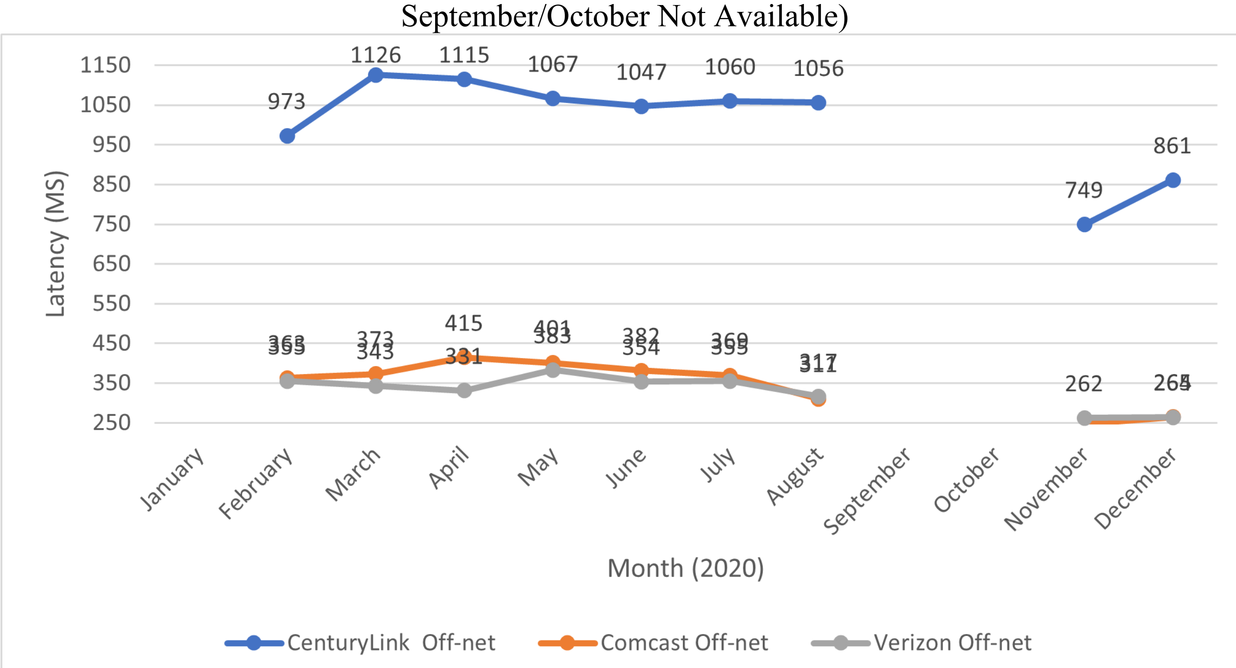

Figure 4: Upstream Average Maximum Latency Under Load By Technology (2020, September/October Not Available)

Given this broad categorization of LUL performance across ISP using different local access network technologies, the next step in our analysis is to take a closer look at the LUL data reported for one of our selected ISPs, Comcast, to draw additional insights on the implications of this data on network performance.

Figure 5 provides a historical view of LUL data for Comcast in the month of November over 2015 – 2020, which are the years during which the MBA program has collected data on the LUL performance test. The figure shows how the average and maximum RTT during the LUL test have changed for the past six years. One clear observation from these results is the improvement achieved in LUL performance over this time. These results suggest that LUL performance for both average and maximum RTT has been improving on average about 10% every year. Reasons for this improvement over this time period are likely due to increasing service speeds along with improvements in the IP protocols and equipment (such as the impact of using AQM in the CMTS).

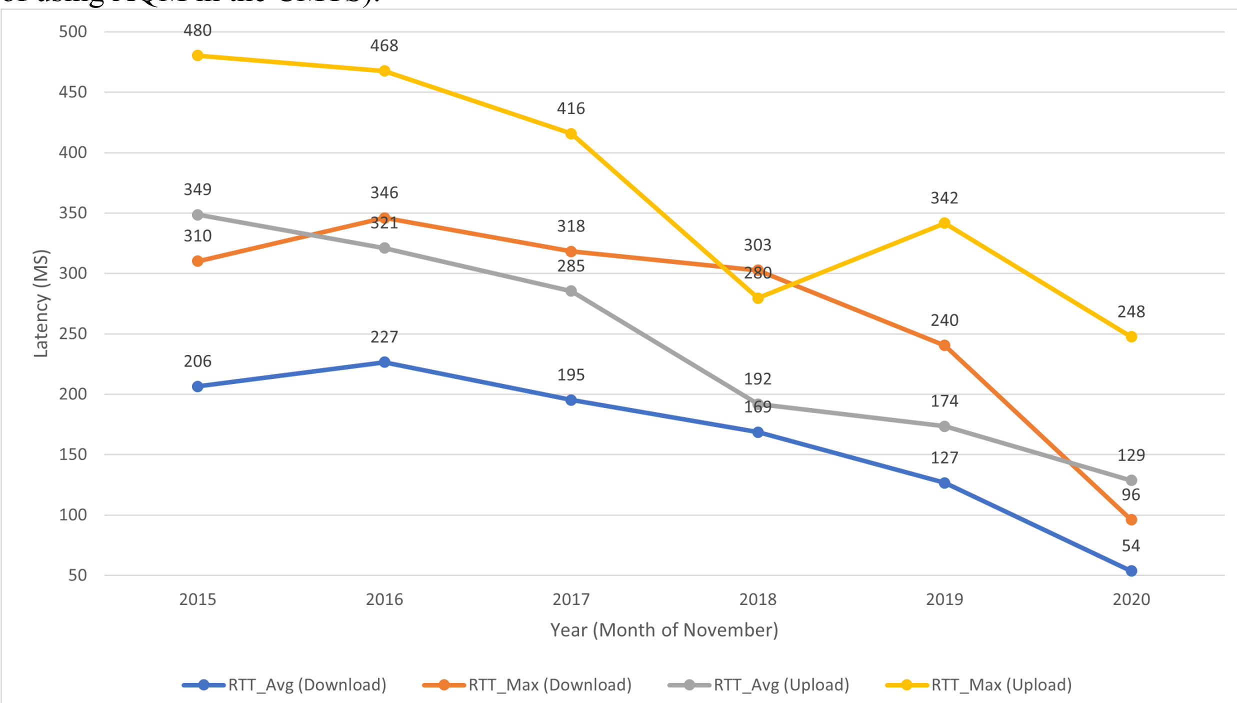

Figure 5: Historical Downstream and Upstream Latency Under Load (Nov, 2015 - 2020)

Throughout our analysis we have been implying that network latency decreases with ISP service speed without providing any evidence of this relationship. To show this is indeed the case, Figures 6 and 7 show scatter plots of the downstream and upstream average RTT LUL data, respectively, for Comcast in July and November of 2020. These graphs clearly show how LUL decreases with increasing speeds along the x-axis. Looking more closely at the scatter plot of downstream LUL data in Figure 6, one also can see clusters of data points around the downstream speed tiers offered by Comcast. According to the FCC's Technical Appendix, Comcast download tiers in 2019 were 60, 150, and 250 Mbps and the upload tiers were 5 and 10 Mbps. The clusters in Figure 6 seem to indicate downstream speed tiers of roughly 25, 90, 225, and 350 Mbps, which does not align with the 2019 tiers in any recognizable pattern. Data clusters in Figure 7 appear to indicate upstream speed tiers of 3, 6, 12, 18, and 24 Mbps, with the data clusters around 6 and 16 Mbps aligning to the 2019 tiers of 5 and 10 Mbps.

Choosing these two months to graph LUL data not only demonstrates the relationship between latency and speed over a cable network, but also allows us to see how the implementation of AQM impacts the downstream latency. Looking at Figure 6, the improved LUL performance provided by AQM in November as compared to July is very clear. In Figure 7, there does not appear to be any difference in LUL performance between July and November but this is to be expected as the AQM technology implemented by Comcast was in the downstream direction.

Finally, Figure 8 shows a different look at the comparison of the before and after LUL performance of the new AQM technique on the Comcast network. This chart shows the distribution or percentage of LUL values (both average and maximum RTT) falling into one of 3 categories: 0 -149 ms (acceptable quality for most real-time services), 150 – 249 ms (some deterioration in quality), and above 250 ms (unacceptable quality). The data from July shows the downstream performance of the Comcast network before the new AQM technique was implemented. So, for example, 81% of the average RTT LUL samples in July were below 149 ms (and 10% above 250 ms), and 69% of the maximum RTT LUL samples were below 149 ms (and 20% above 250 ms).

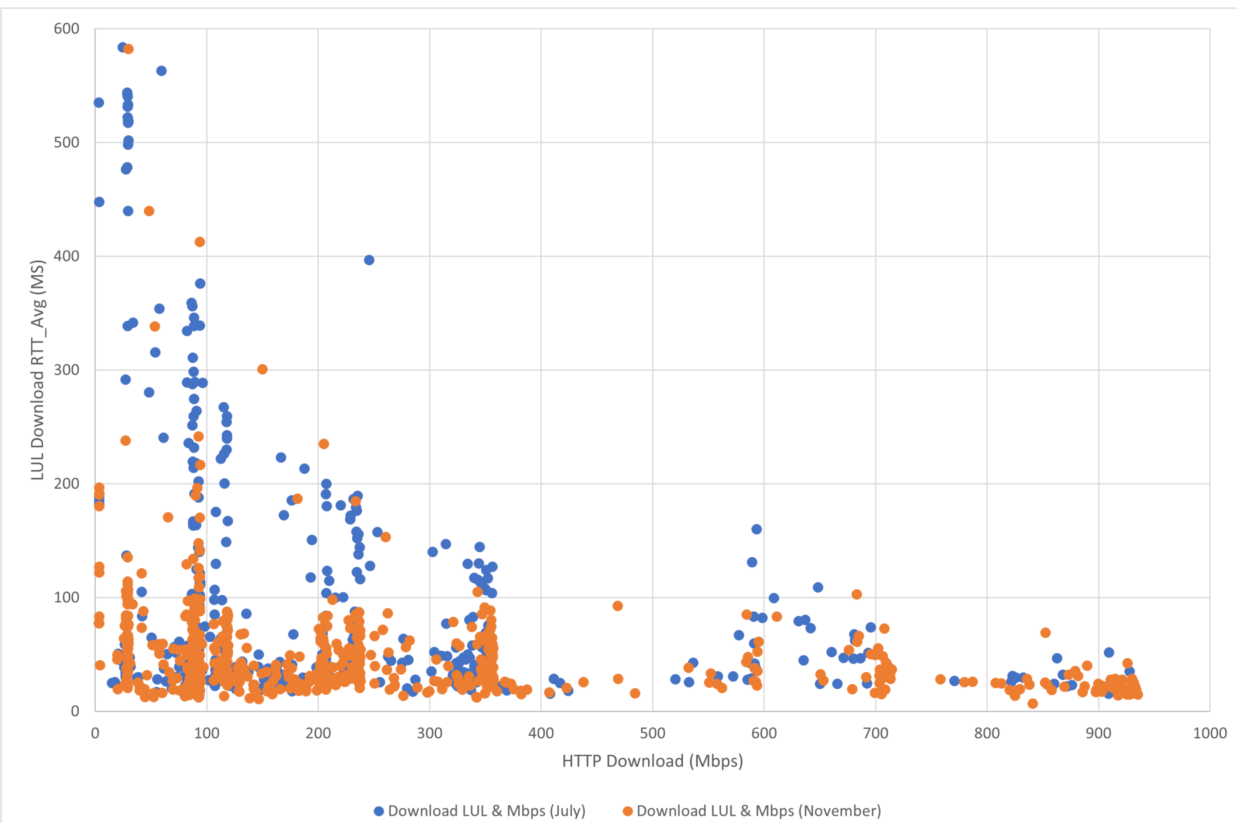

Figure 6: Scatter Plot of Downstream LUL Average for RTT and HTTP for Comcast (July and November 2020)

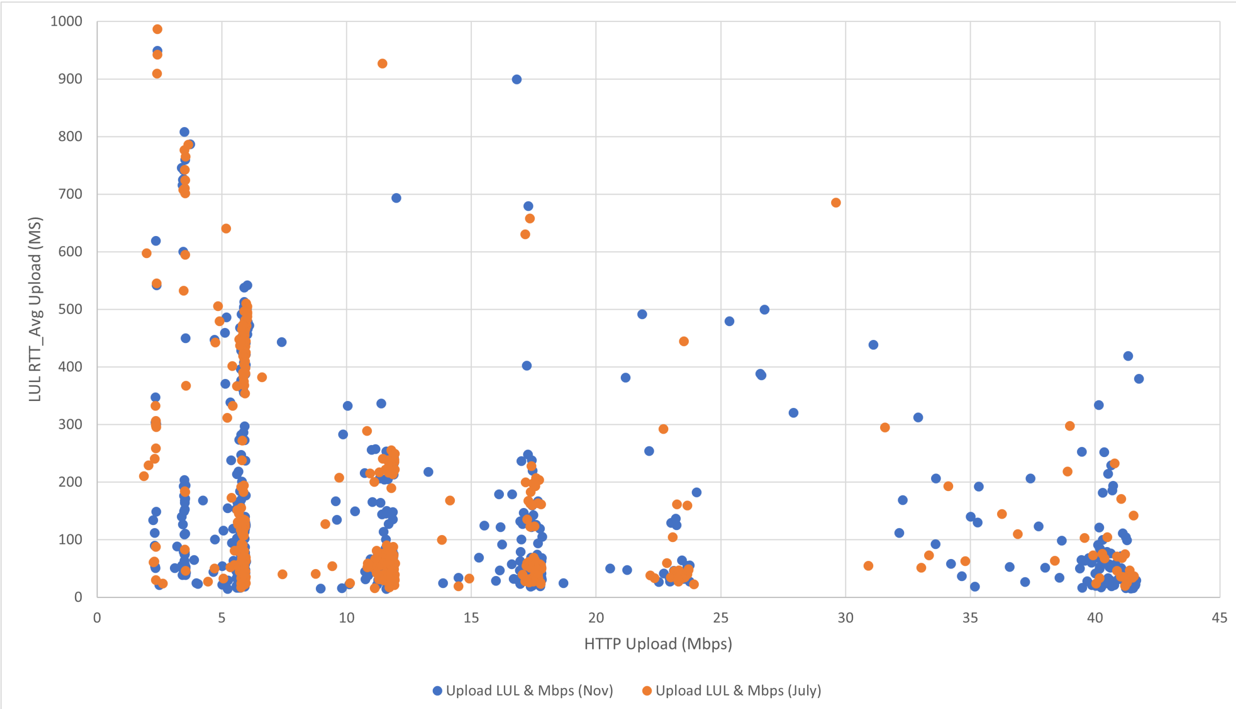

Figure 7: Scatter Plot of Upstream LUL Average for RTT and HTTP for Comcast (July and November 2020)

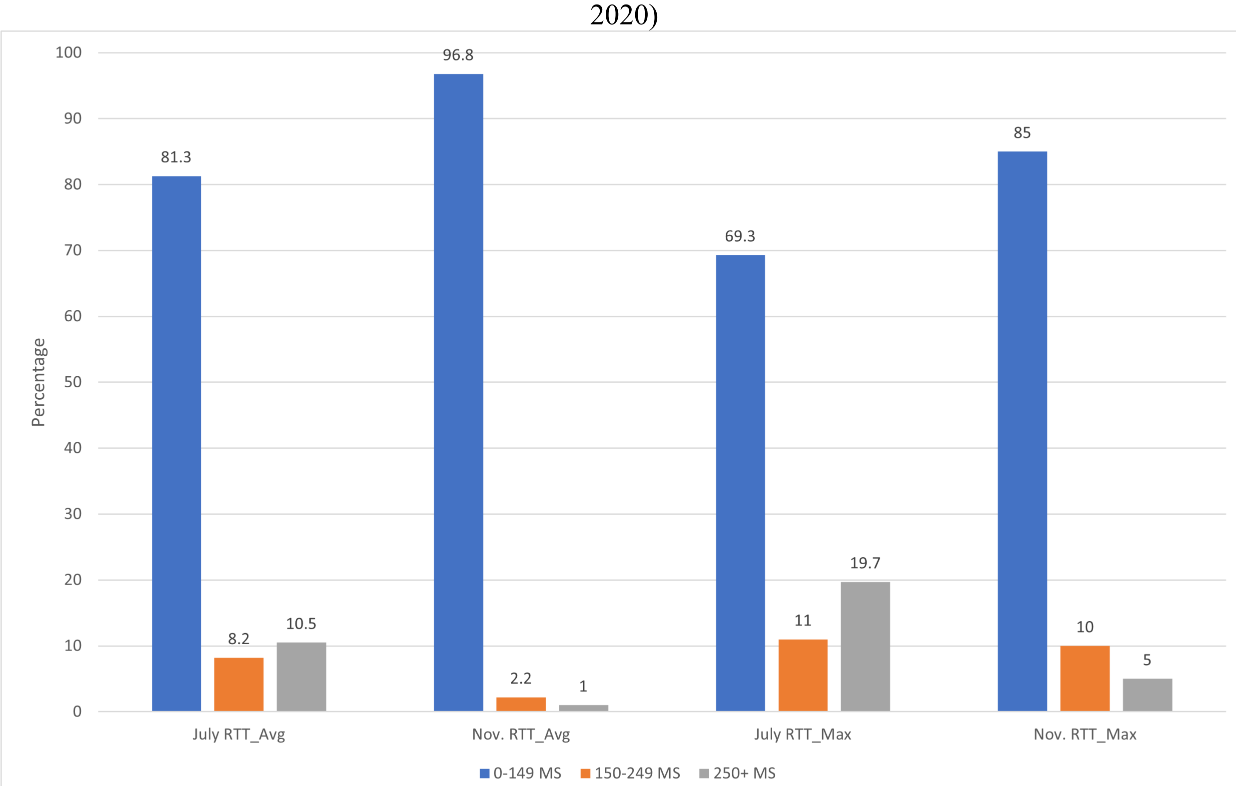

Figure 8: Distribution of Downstream Latency Under Load for Average and Maximum RTT, Before (July) and After (November) the Use of New AQM for Comcast (2020)

After the AQM was activated in November, however, the 97% of the average RTT LUL samples were below 149 ms (and only 1% above 250 ms), and 85% of the maximum RTT LUL samples were below 149 ms (and 5% above 250 ms). The MBA data indicates there has been a dramatic improvement in network LUL performance, particularly in reducing the number of LUL samples above 250 ms.

Conclusions

This study analyzes the latency under load data generated by a broadband performance test included in the FCC's Measuring Broadband America program. Our analysis of this often-overlooked performance metric from the MBA dataset has provided several interesting insights into the performance of a selected set of ISPs that participating in the MBA program. These insights include the following:

1. LUL data provides additional information than delivered in prior MBA Reports. The larger latency reported by the LUL metric indicates how network performance changes under increased traffic loads, particularly when compared to idle latency measurements that the FCC has reported on in the past. The consistent measurement approach implemented by the MBA program allows for a useful opportunity for comparative evaluation of how well the networks of participating ISPs perform under the additional stress presented by this performance test. As broadband performance in terms of latency and reliability continues to be an escalating priority for policy makers given their increasing role in the deployment and regulation of broadband infrastructure, more refined performance tests will be sought after to understand the performance of this critical infrastructure more fully in times of crisis or during peak usage periods.
2. Even for symmetric speed tiers, there appears to be a substantial difference in the magnitude of downstream and upstream LUL for all technologies. This may be due to the use of multiple access protocols to manage access to shared bandwidth in the upstream, though further investigation is needed to investigate more the reason for this disparity.
3. There is substantial variation in downstream and upstream LUL by technology type. Varying service speeds across technologies is the main explanation for the differences. Given our analysis, the latency performance of high-bandwidth cable and FTTH network systems was demonstrably better than the DSL platform that is close to end-of-life status in support of the maturing broadband markets.
4. The maximum round-trip times captured in the LUL tests show the upper bound on potential latency experienced by users, and this data indicates that users caught in a maximum latency event would experience noticeable degradation of real-time applications. The LUL data provides a view of how often latency might exceed important performance benchmarks when real-time services, such as VoIP, are highly impaired.
5. LUL data collected over the last 6 years show that LUL performance by a cable ISP (Comcast) improved roughly 10% per year on average. The reasons for this are due to increasing services speeds and improvements in IP protocols and devices over this time period. Comcast LUL data also indicated the benefits of implementing new AQM technology in the downstream which immediately improved average latency under load performance by 48% after its implementation. Further break down of the results indicated that the most significant improvement occurred in reducing the number of very high LUL values above 250 ms. The AQM implementation provides a concrete example of how

latency can be improved over time through innovative improvements to technical protocols.

6. Finally, LUL data shows the relationship between latency and service speed. As service speeds increase, latency decreases. Given the long-term trend of broadband service speeds that have been consistently increasing over time, the prospects for improved latency performance would appear to be likely as the average service speeds increase above the current definition of 25 Mbps, to 100 Mbps or above in the future.