

Mapping the Digital Divide: Before, During, and After COVID-19

Francesco Bronzino
Université Savoie Mont Blanc

Nick Feamster, Shinan Liu, James Saxon
University of Chicago

Paul Schmitt
Princeton University

Abstract

The digital divide—and, in particular, the homework gap—have been exacerbated by the COVID-19 pandemic, laying bare not only the inequities in broadband Internet access but also how these inequities ultimately affect citizens' ability to learn, work, and play. Addressing these inequities ultimately requires having holistic, “full stack” data on the nature of the gaps in infrastructure and uptake—from the physical infrastructure (e.g., fiber, cable) to speed and application performance to affordability and neighborhood effects that ultimately affect whether a technology is adopted. This paper surveys how various existing datasets can (and cannot) shed light on these gaps, the limitations of these datasets, what we know from existing data about how the Internet responded to shifts in traffic during COVID-19, and—importantly for the future—what data we need to better understand these problems moving forward and how the research community, policymakers, and the public might gain access to various data.

1 Introduction

The Federal Communications Commission (FCC) has regularly suggested that educational and economic opportunity increasingly depends on the availability of affordable, high-speed Internet access. One aspect of this is the “*homework gap*”, referring to the barriers that students face when working on homework assignments without reliable Internet access at home. As schools increasingly incorporate Internet-based learning into the curriculum, students become increasingly dependent on reliable Internet access to succeed.

In 2009, the FCC's Broadband Task Force reported that about 70% of teachers assign homework that require access to broadband Internet access, and that about 65% of students use the Internet at home to complete their homework, including submitting assignments, working on shared projects, and communicating with teachers. As of 2018, 24 million Americans have no access at all to high-speed Internet access, and 12 million Americans face the so-called homework gap; 25% of US households with children under 8 years old do not have

access to broadband Internet access. Many of those who do have access do not have access to unlimited Internet access plans.

With the rise of the COVID-19 pandemic, the problem has never been more immediate or important: Questions of inequity have become more acute in recent months, as a global pandemic has caused shifts in patterns of education, work, and life that involve increased dependency on high-speed broadband Internet access from home—for example, with increased dependence on remote learning and remote caregiving—at the very time that the network itself is experiencing increased load.

An important step towards addressing this inequity involves developing a deep, holistic understanding of available broadband Internet access in the United States. Beyond simply “speed tests”, reliable, usable Internet access depends on a holistic set of factors, including (1) physical connectivity (location of fiber and cable); (2) network throughput, latency, and packet loss; (3) application performance metrics; (4) other various non-technical data, including pricing data, information about subscription rates, penetration, and adoption, and so forth.

First, we explore the current nature of these metrics, surveying the available data sources on these topics as well as the various shortcomings of various datasets. We focus on the current state of various data sources, including the state of physical connectivity, network throughput and performance metrics, interconnection measurements, and application performance metrics, noting in particular what these datasets can—and *cannot*—tell us about the current state of the digital divide.

Second, with the above metrics and analysis on historical data as a baseline, the second part of the paper explores how the Internet responded under unprecedented shifts in traffic load as a result of the COVID-19 pandemic. We study the effects of the shifts in traffic load resulting from the COVID-19 pandemic on network utilization and network performance, noting how various aspects of the network infrastructure adapted to shifts in traffic load during COVID-19.

Third, we highlight gaps in existing datasets, in particular how existing public datasets do not shed light on application quality of experience and, ultimately, user experience. Towards filling this gap, we propose mechanisms for gathering data along these dimensions, first at a hyper-local, neighborhood level, but ultimately with an eye towards scaling up a national infrastructure that can provide a more holistic view into these dimensions.

2 How We Map Internet Coverage

The coronavirus pandemic forced daily routines into virtual spaces. As both work and learning went remote, policymakers around the country have sought to address ongoing inequities in Internet connectivity, suddenly more acute than ever. Data sources on Internet connectivity have substantial limitations in measuring inequity. In this Section, we highlight these limitations while demonstrating what can nevertheless be done with the data.

2.1 Existing Metrics

Various metrics exist on the state of broadband Internet connectivity. Each of these metrics provides a partial view of the state of connectivity—or, in some cases, lack thereof. We briefly survey these metrics below, organizing the discussion into (1) performance measurements (of which speed tests are perhaps the most common instance); and (2) survey-based methods.

2.1.1 Performance Measurements

A common way to measure Internet connectivity is via performance measurements. There are various types of performance measurements: the most common of these are “speed tests”, which actively measure the performance of the network, in many cases by saturating a network link. Speed tests can be router-based or client-based; client-based speed tests can be browser-based or native. Most commonly, speed tests measure network characteristics such as throughput, latency, and packet loss. Some speed tests also measure the performance of specific applications, such as web page load time or video streaming, although those types of tests are typically less common in active speed testing applications.

Router-based speedtests, which measure the performance of the network directly from the router—particularly those that use multiple connections to test the throughput of the connection—typically produce more accurate results and are now a commonly accepted standard way of measuring the performance of the access ISP. These tests were pioneered in 2010 and remain the standard measurement ap-

proach of the Federal Communication Commission’s Measuring Broadband America (MBA) program. Measuring from the router mitigates some of the problems of client-based testing

2.1.2 Passive Quality Measurements

In addition to speed tests, an emerging class of performance measurements are passive application performance inference tools, which passively survey network traffic and infer the quality of experience of an application based on various properties of the traffic (e.g., segment arrival time, downstream throughput utilization). Quality of experience metrics generally entail the following four quality of experience metrics: (1) startup delay, or the time to start playing the video; (2) rebuffering, or whether the video has experienced an interruption in playback; (3) video resolution; (4) changes in resolution. The last two metrics concerning resolution are related, although because the problem is based on statistical inference, sometimes models can more accurately detect a change in resolution than the actual resolution (or vice versa).

Passive measurement techniques introduce various tradeoffs. On the one hand, because they involve inferring performance directly from existing network traffic, they can provide an accurate reflection of the performance of applications as they are actually used, as well as an estimate of quality parameters that correspond to the performance of that application. The passive nature of the measurements also eschews the need for bespoke active measurement tests that are specific to applications (e.g., web, streaming video, gaming), as well as dedicated test infrastructure for each of those tests that is representative of the user experience when using those applications. On the other hand, these methods face various deployment barriers, as we discuss in the subsequent section.

2.1.3 Interconnection Data

Previous efforts have made it possible to explore utilization of Internet interconnection points between access ISPs and content providers. In particular, previous work has analyzed the utilization and provisioning of interconnection measurement points using IPFIX-based traffic measurements and SNMP polling data, respectively. Data from the Interconnection Measurement Project in particular makes the following statistics available: (1) timestamps at five-minute intervals; (2) geographic region; (3) anonymized partner network (e.g., content delivery network, application service provider);

(4) access ISP; (5) total ingress bytes; (6) total egress bytes; (7) capacity.

Such data makes it possible to observe both utilization rates and how ISPs have provisioned and augmented capacity over various periods of time. The data provides helpful perspective in understanding the extent to which provisioning between access ISPs and peer networks is sufficient to sustain traffic demands, as well as how providers have responded over time to increases in traffic demands. As we will see in Section 3, this data provides helpful perspective in understanding how certain service providers responded in the face of considerable shifts in traffic demands as a result of the COVID-19 pandemic.

2.1.4 Surveys

In the American Community Survey (ACS) the US Census Bureau asks respondents whether they have broadband Internet; for the National Telecommunications and Information Administration (NTIA) supplement to the Current Population Survey (CPS), they gather both data on connectivity and online behaviors. While the US Census Bureau is the “gold standard” of survey methodology, the question of broadband access is coarse: the presence of a nominal, 25/3 Mbps contract. Wealthy Americans have enjoyed higher bandwidths over time; it is not possible to understand relative changes on the *intensive* margin (consumption or bandwidth) using the ACS or the CPS. Unfortunately, subtlety is not required: neighborhood economic indicators are strongly predictive of broadband contracts. In Chicago for instance, the tract-level correlation between log median household income and the share of households with a broadband contract is 0.82, as can be seen in Figure 1. The CPS/NTIA data are not designed for use below the state-level, but microdata allow breakdowns by race, ethnicity, and poverty status.

2.2 Limitations of Existing Methods

The existing methods for characterizing access network performance face various limitations. We now survey various limitations concerning test methods, sampling, and deployment barriers.

2.2.1 Test Environment and Methods

Network environment. Client-based speed tests have a number of limitations. One of those limitations is that, as access network speeds have increased, the speed test itself is less likely to be measuring the speed of the access link, and far more likely to be measuring the speed of some other aspect of the end-to-end path, as other elements along the net-




Figure 1: Census tract median household income versus share of households with broadband Internet, in Chicago. A lowess curve is drawn. The linear correlation is 0.82.

work path introduce performance bottlenecks. Some notable issues include the following:

- At speeds above 100–200 Mbps, browser-based speedtests can produce inaccurate results, due to the fact that the browser itself becomes a performance bottleneck.
- On older phones, the cellular radio itself can introduce a performance bottleneck, if the WiFi radio is not capable of transmitting at rates that exceed that of the access network.
- In the case of a home network, the home WiFi network itself can introduce performance bottlenecks due to old equipment, a poorly placed WiFi access point, or other factors.

Test design. The design of the test itself can also pose limitations. The Network Diagnostic Test (NDT), used as the underlying tool in the Measurement Lab tests, for example, has long used a single TCP connection to test network speeds. Such a test has failed to successfully saturate links as low as 10 Mbps, particularly if the link has any amount of packet loss. Despite the known shortcomings of the tool, it has continued to be used in an attempt to demonstrate the performance of ISP networks, when today more often than not the tool itself is the bottleneck.

Test infrastructure. The infrastructure that supports the test may also introduce limitations, if it is incapable of supporting the speeds required to test the speed of the access


link for a large number of tests. In the past, the Measurement Lab infrastructure has suffered from a lack of available server endpoints (especially as compared to approaches that use similar methods such as Ookla, which have an order of magnitude more servers). The same infrastructure has also suffered performance limitations due to the underlying network infrastructure which resulted in packet losses under switch microbursts, resulting in data from the FCC’s MBA program that had to be invalidated.

2.2.2 Sample Limitations


Sampling bias. Objective measures of real-world performance could shed light on differences in online behaviors or infrastructures by community. For this purpose, we can examine data from SamKnows whiteboxes or distributed measurements from Ookla speedtests, newly-released as Open Data [12]. The SamKnows data includes many measures of performance including bandwidth and consumption, and the Ookla speedtest measures bandwidth and latency. The challenge with these data is that they are highly non-representative. The sampling frame for the SamKnows data was stratified by ISP and speed tier (SamKnows worked with ISPs to under-represented speed cells), rather than according to population demographics [6]. While no demographic indicators are provided, block group or tract of residence is available for 94% of the households included in the 2018 report. If we aggregate block groups to the tracts that contain them, we can compare tract median household income between SamKnows residences to the American population (Figure 2). This reveals a significant offset: a log difference of 0.13, implying incomes 14% higher in SamKnows “neighborhoods.”

A similar effect is apparent in the Ookla data. The Ookla data are provided publicly as quad tiles (zoom level 16), which are roughly half a kilometer at the equator. These do not map neatly to Census geographies, but given the small tiles we can allocate them approximately. We assign each device from each quad tile in Chicago to a random point within its tile; we then merge these generated points to Census tracts. This rough measure of the population of devices can then be contrasted with the Census, by median household income, as above.

Just as for the SamKnows data, the devices are recorded in neighborhoods that are markedly wealthier than the general population of Chicago. The obvious reason for this is that the general population does not all have Internet. However, if we weight Census populations by the share of households with Internet, the gap remains significant. This suggests that people in affluent areas are more likely to run speed tests, even



(a) Ookla / Chicago




(b) SamKnows / US

Figure 2: Normalized distribution of Census tract median household income, weighted population counts, or for Ookla devices or SamKnows WhiteBox households. The dotted lines weight Census populations by the share of neighborhood households with a broadband Internet contract.


conditional on having Internet. One can further see these neighborhood effects when exploring this data geographically, as shown in Figure 3.

In short, both state-of-the-art distributed measurements and official FCC samples are systematically blind to vulnerable populations, who do not run speedtests or volunteer for FCC measurements.


Though we emphasize the significant limitations of the data, it is possible to contrast performance over the course of the coronavirus pandemic, by neighborhood, for SamKnows households. We first divide the SamKnows sample into three categories by neighborhood income – bottom quartile, middle quartiles, and highest quartile of the *SamKnows sample*. We then plot consumption and speed over time. Within each



(a) Chicago



(b) Philadelphia



(c) Washington

Figure 3: The number of Ookla Speedtests performed during Q3 2020 for three major U.S. cities shows how the test samples vary by geography (and demographic).




Figure 4: Bandwidth measured by SamKnows data for low (first-quartile), medium (middle quartiles), and high-income tracts. Quartiles are of the SamKnows population. For each segment of the population the three lines show the first, second and third quartiles of bandwidth.




Figure 5: Received and transmitted data, measured by SamKnows for low (first-quartile), medium (middle quartiles), and high-income tracts. Quartiles are of the SamKnows population. For each segment of the population the three lines show the first, second and third quartiles of consumption by day.

income group, we plot the 1st quartile, median, and 3rd quartile of the distribution, from November 2019 until the most-recent release, August 2020. These effects are shown in Figure 4.

Conditional on *having Internet*, households from wealthy neighborhoods do in fact purchase higher-bandwidth plans, as shown in Figure 5. On the other hand, it is somewhat surprising to find that from the perspective of consumption households in lower income communities tend to use moderately *more* Internet, across the distribution (i.e., at all three quartiles). These results remain somewhat inconclusive given the limitations of the data, but they underscore the

need for demographically-attuned measurement of Internet performance and use.

2.2.3 Deployment Barriers

On the other hand, passive collection of application performance information introduces a unique set of challenges since it requires collecting traffic from the network. This approach introduces both challenges of technical instrumentation and privacy. First, gathering this type of data from a home network requires collecting traffic in the home, which can require extensive instrumentation of the network itself. A typical home might require a modem, router, and WiFi access point, in many cases this functionality is combined in a single device—sometimes one provided by the ISP, but in




Figure 6: *Passive application performance inference requires extensive network instrumentation, posing a significant deployment barrier.*


any case one that is readily available. Unfortunately, these existing devices do not typically make it easy or possible to capture network traffic, especially at line rate.

As a result, capturing such traffic requires the level of instrumentation shown in Figure 6, whereby an additional switch and capture device are interposed between the router and the WiFi access point. Additionally, the user’s WiFi access point must be configured in bridge mode (the default configuration of many WiFi access points is often to act as the router). This configuration is necessary because traffic capture requires the capability to mirror all traffic seen on a switch port to a capture device, and consumer-grade routers and WiFi access points typically do not provide this capability. Such instrumentation can prove prohibitive, particularly for non-technical users.


Furthermore, this measurement technique and instrumentation introduce potential privacy concerns, given that the method requires capture of passive traffic. In our own previous work, we have designed such techniques to collect the minimum amount of data that is necessary to infer application performance characteristics, but the technique nonetheless may be perceived to be invasive by some. This concern is particularly acute in light of the sampling questions in earlier sections, given that portions of the demographic that are not as well-represented in other datasets may also

3 The Internet During COVID-19

Internet traffic patterns shifted dramatically at the start of the COVID-19 pandemic, as people increasingly began to work from home. Various reports have documented how different aspects of the Internet infrastructure responded during this time period [1, 3, 8, 10]. In this section, we briefly survey what we know about how the Internet responded from public data, additional information we know from semi-public datasets (e.g., interconnect data), and various things we do *not* know either due to a lack of available public



(a) *Absolute utilization.*




(b) *Normalized utilization.*

Figure 7: *99th percentile interconnect link utilization for two ISPs.*


datasets or, more generally, a lack of accepted method for measuring metrics of interest.

3.1 What We Know

Available datasets can tell us about how various network metrics, such as throughput and latency fared over the course of the pandemic, how these how these metrics varied across ISPs, and how the effects of the pandemic compare to other exogenous “shocks” (e.g., the holiday season, when home Internet use also tends to experience a surge). In this section, we briefly survey what we know from these datasets. These results are also covered in more detail in a corresponding technical publication [9].



(a) Peak download.



(b) Peak upload.

Figure 8: Peer link utilization for ISP A between January 15 to April 15, 2020.


Traffic volumes increased in general. Traffic surged about 20% in Europe for broadband networks [7]; the same study found that the fixed-line Internet infrastructure was able to sustain the 15–20% increase in traffic that happened

rapidly during a short window of one week. In the United States, downstream peak traffic has recently stabilized [11], but in the early weeks of the pandemic, it showed a growth of 20.1%. Wireless networks saw volume increases of up to 12.2% for voice and 28.4% for data by the top four providers were shown in an industry report [4]. Industry operators have self-reported on their network responses largely through blog posts [1, 3, 8, 10]. Facebook found that developing regions experienced more significant performance degradations [2]. Network latencies were approximately 30% higher during the lockdown in Italy [7]. According to an NCTA report, networks in the United States saw less congestion [11].


Traffic patterns increased at interconnection points. We consider peak-hour link utilization from the Interconnect Measurement Project as a measure of traffic demand [5]. We pre-processed the interconnect dataset and remove anomalous data points that are caused by failures in the measurement system. Figure 7 shows both the absolute utilization and the utilization normalized against the link capacity for two anonymized ISPs. For each ISP, we show the value corresponding to the 99th percentile link utilization for a given day. Figure 7a shows that ISP A experienced increase concurrent with the initial COVID-19 lockdowns (early March 2020); ISP B saw a smaller raw increase in utilization. Figure 7b also shows that both ISPs experienced significant increases in normalized utilization in March–April 2020.

Traffic ratios changed for some peers. We also compared the peak hour download and upload rates on all of ISP A’s interconnects on (1) January 15, 2020, and (2) April 15, 2020 (Figure 8). Traffic volumes to (and from) some peers change significantly—some by several orders of magnitude. Such drastic changes may be attributable to users working from home and connecting to services that would cause more traffic to traverse the peer link in the upstream direction. We confirmed these results with the operators at ISP A and report that they observed that streaming video traffic decreased from 67% to 63% of the total traffic, but video conferencing increased from 1% to 4% as a percentage of overall traffic.

Some ISPs experienced higher latency, particularly in the tail. Figures 9 shows the 95th percentile of round-trip latency across the SamKnows Whiteboxes in the FCC Measuring Broadband America before and during the COVID-19 lockdown. These results show that, overall 95th percentile latency across most ISPs remained stable. We also studied 99th percentile latency which did show some deviations from



(a) 95th percentile of ISP latency (Group 1).




(b) 95th percentile of ISP latency (Group 2).

Figure 9: Latency (95th percentile) for different ISPs.


normal levels during lockdown for certain ISPs. Notably, in many cases the same ISPs experienced deviations in latency during other periods of time, as well (e.g., during the December holidays), suggesting that even in the tail, latency effects of COVID were not overly pronounced. These results are likely explained, at least in part, by ISP capacity augmentation.

ISPs responded to changes in traffic demands by augmenting capacity. Using the Interconnection Measurement Project dataset, we calculate the total interconnect capacity for each ISP by summing the capacities for all of the

links associated with the ISP. To enable comparison between ISPs that may have more or less infrastructure overall, we normalize the capacity values for each using min-max normalization. We again filter out date values that are beyond two standard deviations from a rolling 60-day window mean. A linear regression yields a slope for each quarter that illustrates the best-fit rate of capacity increases over that quarter. We scale the slope value to show what the increase would be if the rate of increase was maintained for 365 days (i.e., a slope of 1 would result in a doubling of capacity over the course of a year). Figure 10 shows the resulting capacity plots. The overall trend shows how these two ISPs in the United States



(a) ISP A



(b) ISP B

Figure 10: Normalized interconnect capacity increases for two ISPs.

added capacity at more than twice the rate at which they were adding capacity over a comparable time period in the previous year. Second, both ISPs significantly added capacity in the first quarter of 2020—at a far greater rate than they were adding capacity in the first quarter of 2019.

3.2 What We Don't Know

Application performance. Available public datasets have little to say about how individual *applications* performed, as well as user quality of experience when using those applications. Our inability to shed light on application performance and user quality of experience is somewhat problematic, given that such metrics ultimately reflect how

well the Internet is performing for users within the context of the applications they are trying to use. Unfortunately, for reasons previously mentioned, obtaining application performance and quality of experience metrics at scale can be challenging because doing so requires either instrumentation of specific applications or (as previously discussed) invasive network instrumentation coupled with statistical inference techniques to infer application quality. An open challenge for the policy community is thus reconciling the importance of these types of metrics with the lack of these metrics in publicly available datasets and the general inability to gather such publicly available datasets at scale. An alternative approach may be to first design such measurements to be “hyper local”, as we discuss in Section 4, and then determine how to scale them more generally.

Measurements at fine spatial and temporal granularity.

Critics are often quick to point out that aggregates “hide” outliers. For this reason, much of our analysis in this paper and other related work has focused on the 95th and 99th percentiles, as well as exploration of outlier statistics. Aggregates can still be necessary and useful: not only can they serve as helpful summary statistics but also in some cases (e.g., with interconnection data), aggregation is a condition to publicly sharing data that might otherwise be private. The issue, then, is not to necessarily avoid aggregates altogether but rather to be very clear on what these statistics can and cannot tell us. One particular limitation of the interconnection data we have presented, for example, is that it does not tell us whether particular isolated or short-lived high-utilization events occurred on a particular interconnect; the private dataset allows us to affirm this, but sharing data at finer granularity with the general public would permit more thorough examination from the public. Second, and particularly important to the digital divide, many available datasets do not permit an exploration of the data at block-level geographic granularity, yet many of the issues with infrastructure, performance, and broadband Internet adoption exist at that granularity.

How different communities fared. Related to the lack of fine spatial granularity in existing datasets is the lack of information more generally about how various underserved communities have fared with respect to Internet availability, adoption, and performance. As discussed, existing survey-based datasets ask questions about availability of broadband Internet at extremely coarse granularity; that information is not collected in a way that is specific to how Internet uptake or performance fared during the pandemic, nor does it shed

light on the quality of experience these users have faced when attempting to use video conference applications (or other applications) as they have been forced to work, learn, and play at home, especially within communities that have long faced barriers to Internet infrastructure.

4 Looking Ahead

While existing datasets have allowed us to understand that the Internet infrastructure has held up very well during the pandemic, the current situation has amplified the disparities that have existed for a long time.

Ultimately, bridging this divide will require gathering more holistic data about not only network performance, but also better information about application performance, user quality of experience, and better data about uptake and adoption. Better data can ultimately drive investment and direct policy, but doing so effectively requires knowing where those problems lie, and what specifically is contributing to these gaps in connectivity. Previous work has explored these problem from one dimension at a time, from the perspective of Internet speed or surveys. Better data can also provide decision makers with insights so that they can make the best possible decisions about how to direct the resources and investments. An important aspect of this exercise is matching the data inference and insights to the questions that policy makers are asking and need to know. Too often available data is not responsive to the precise questions being asked by policymakers, yet is used anyway to advocate. In this case, we must seek data that go beyond merely providing evidence that a problem exists towards insights that give us a deeper understanding of the true causes of the problems.

Initially, our focus on improved datasets involves going deeper into a few communities in Chicago to deeply understand the many dimensions of this problem, including physical infrastructure, network availability and performance, application quality of experience, and other factors including adoption and use within communities. We aim to develop these new methods designing them in such a way that it serves as a rubric that can be replicated and scaled from a few communities to cities, and ultimately to nations.

Acknowledgments

This work was funded in part by National Science Foundation Award CNS-2028145. We thank CableLabs and members of the Interconnection Measurement Project consortium for access to interconnection utilization and capacity data.

References

- [1] AT&T. *COVID-19: Our Response*. <https://about.att.com/pages/COVID-19.html>. accessed October, 2020.
- [2] Timm Böttger, Ghida Ibrahim, and Ben Vallis. “How the Internet reacted to Covid-19 – A perspective from Facebook’s Edge Network.” In: *ACM IMC*. 2020.
- [3] Comcast. *COVID-19 Network Update*. <https://corporate.comcast.com/covid-19/network/may-20-2020>. accessed October, 2020.
- [4] CTIA. *The Wireless Industry Responds to COVID-19: Network Performance*. <https://www.ctia.org/homepage/covid-19#network-performance>. accessed October, 2020.
- [5] Nick Feamster. “Revealing Utilization at Internet Interconnection Points.” In: *Conference on Communications, Information, and Internet Policy (TPRC)*. Washington, DC, Sept. 2016, pp. 1–10. URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2756888.
- [6] Federal Communications Commission. *2016 Technical Appendix: Measuring Broadband America/Fixed Broadband*. Tech. rep. Washington, DC: FCC Office of Engineering et al., 2016. URL: <https://data.fcc.gov/download/measuring-broadband-america/2016/Technical-Appendix-fixed-2016.pdf>.
- [7] Anja Feldmann, Oliver Gasser, Franziska Lichtblau, Enric Pujol, Ingmar Poese, Christoph Dietzel, Daniel Wagner, Matthias Wichtlhuber, Juan Tapidor, Narseo Vallina-Rodriguez, et al. “The Lockdown Effect: Implications of the COVID-19 Pandemic on Internet Traffic.” In: *Internet Measurement Conference (IMC ’20)*. 2020.
- [8] Google. *Keeping our network infrastructure strong amid COVID-19*. <https://blog.google/inside-google/infrastructure/keeping-our-network-infrastructure-strong-amid-covid-19/>. accessed October, 2020.
- [9] Shinan Liu, Paul Schmitt, Francesco Bronzino, and Nick Feamster. “Characterizing Service Provider Response to the COVID-19 Pandemic in the United States.” In: *Passive and Active Measurement Conference (PAM)*. Brandenburg, Germany, Mar. 2021, pp. 1–12. URL: <https://www.pam2021.b-tu.de/>.
- [10] Martin McKeay. *Parts of a Whole: Effect of COVID-19 on US Internet Traffic*. <https://blogs.akamai.com/sitr/2020/04/parts-of-a-whole-effect-of-covid-19-on-us-internet-traffic.html>. accessed October, 2020.

- [11] NCTA. *COVID-19: How Cable's Internet Networks Are Performing: METRICS, TRENDS & OBSERVATIONS*. <https://www.ncta.com/COVIDdashboard>. accessed October, 2020.
- [12] Speedtest® by Ookla®. *Global Fixed and Mobile Network Performance Maps*. Based on analysis by Ookla of Speedtest Intelligence® data for Q3 2020. Provided by Ookla. Ookla trademarks used under license and reprinted with permission. URL: <https://registry.opendata.aws/speedtest-global-performance/>.