Measuring Broadband America: A Retrospective on Origins, Achievements, and Challenges

Eric W. Burger*
Commonwealth Cyber Initiative
Virginia Tech
Arlington, VA, USA
eric.burger@vt.edu

Padma Krishnaswamy
Office of Engineering and Technology
Federal Communications Commission
Washington, DC, USA
padma.krishnaswamy@fcc.gov

Henning Schulzrinne
Department of Computer Science
Columbia University
New York, NY, USA
hgs@cs.columbia.edu

This article is an editorial note submitted to CCR. It has NOT been peer reviewed.
The opinions expressed by the authors are their own and do not necessarily reflect the opinions of their respective organizations. Comments can be posted through CCR Online.

ABSTRACT

The "Measuring Broadband America" program, run by the United States Federal Communications Commission (FCC), continually measures and releases data on the performance of consumer broadband access networks in the US. This paper presents a retrospective on the program, from its beginnings in 2010 to the present. It also reviews the underlying measurement approaches, philosophies, distinguishing features, and lessons learned over the program's duration thus far. We focus on fixed broadband access since it is the program component with the longest history. We also discuss future directions and challenges.

CCS CONCEPTS

• **Networks** → **Routing protocols**; • **Security and privacy** → **Security protocols**;

KEYWORDS

Measuring broadband, public policy, home internet, network measurements

1 INTRODUCTION

The United States Federal Communications Commission (FCC), the U.S. national telecommunications regulatory agency, launched the Measuring Broadband America (MBA) program in 2011 in response to multiple United States national directives, including the US Congress's Broadband Data Improvement Act of 2008 [61] and the recommendations from the National Broadband Plan [62] activity in 2010. MBA measures and reports on the network performance of most large U.S. ISPs with consumer-facing broadband access offerings. MBA is among the earliest and longest-running efforts to characterize the performance of consumer broadband on a national scale.

Since 2013, the FCC has also offered a downloadable consumer mobile app for Android and Apple iOS to measure mobile network performance[1]. For reasons of space and since the measurement of fixed broadband has the longer history, we will only discuss the fixed component of the program.

The FCC designed MBA to provide a reliable, consistent and well-documented source of longitudinal[2] performance data on US broadband, emphasizing transparency as to the metrics, measurement infrastructure, data collection, and the measurement and analyses processes used. Its particular focus is the *access* network performance. It therefore goes to considerable lengths to exclude potential confounding factors due to performance impairments that typically lie outside the access service provider's control, such as those that caused by in-home Wi-Fi, the measurement agent itself, or transit network traffic loads. The FCC makes the raw data from MBA measurements publicly available quarterly at https://www.fcc.gov/oet/mba/raw-data-releases.

Through the raw data repository on collected measurements provided, researchers can perform longitudinal comparisons on the characteristics of United States broadband services over the life of the program. The fixed MBA test suite currently consists of 18 metrics, including download and upload throughput ("speed"), latency, and packet loss [39, 41], along with the performance of essential Internet service adjuncts such as DNS resolution. Application performance tests include a composite Web access metric, reflecting page load time, as well as a VoIP jitter metric and a CDN latency metric. The test agents also collect daily total household Internet usage byte counts, as well as byte counts for the measurement data.

The FCC also issues a (roughly) annual report that presents the results from a subset of the eighteen measurements used by the program [44] (see Section 3.6) and in the case of one of these (download "speeds"), it also compares measured with advertised levels of performance for the ISPs that participate in MBA tests. The MBA reports currently portray carriers' Internet access performance through four metrics selected from the full eighteen: throughput, for both upload and download directions, latency, packet loss, and a composite measure for web performance. This set of metrics provides a basic consumer-focused summary of the performance of broadband services offered by the service providers that are part of the MBA sample. These reports are occasionally augmented; for instance, the 11th MBA report discussed the impact of COVID-19 on network performance.

The MBA program has met the original program objective to provide data and analyses derived from ongoing agency-managed

*Corresponding author

[1] https://www.fcc.gov/general/measuring-mobile-broadband-performance

[2] Wikipedia defines longitudinal data as "multi-dimensional data involving measurements over time"

evaluations of broadband access network performance across the United States, using a consistent, well-defined, open methodology. While maintaining this focus, MBA metrics have evolved to keep pace with Internet access services themselves, including the progression in service technologies from DSL to cable to fiber and accompanying changes in offered service speed and other characteristics. In addition, through a program component entitled “MBA Assisted Research Studies” (MARS), the FCC facilitates the re-use of the MBA measurement infrastructure by U.S. academics, in the public interest of progressing research on evolutionary trends in Internet services and architectures [45].

We next describe the origins and objectives of the program in Section 2, offer context in Section 3, provide a brief summary of how the program evolved in Section 4, draw lessons from the past decade in Section 5, look into possible future directions in Section 6, and summarize in Section 8.

Disclaimer: Views expressed are those of the authors and do not necessarily represent positions or policy of the FCC.

2 PROGRAM ORIGINS AND OBJECTIVES

As expressed in Section 706 of the Telecommunications Act of 1996 [60], the US Congress recognized that broadband communications, providing fast, always-on Internet access to the public, would be an essential enabler of national technological and economic progress. The United States was not alone in arriving at this conclusion. Over the next five years, numerous countries worked on or issued national broadband plans [31].

The Broadband Data Improvement Act of 2008 (BDIA 2008) [61] and the American Recovery and Reinvestment Act of 2009, which mandated the publication of a National Broadband Plan (NBP), were directed towards enabling and facilitating broadband and Internet adoption across the country. Also in that time frame, several commercial entities had published reports on ’actual’ (i.e. measured) broadband speeds, either incidental to their principal business, such as Akamai’s *State of the Internet* reports [8], or as a primary focus, such as Ookla’s Speedtests [46]. In such cases, important details, such as measurement methodology and sampling approach, were considered proprietary or business-sensitive, and thus unavailable for review by researchers, policy makers, and other interested parties. Multiple sources [13, 32] noted there were significant mismatches between advertised Internet access speeds and the corresponding measured performance experienced by consumers of broadband internet services. Yet these accounts were either anecdotal or, if empirically derived, opaque as to the methodology, statistical specifics, and other details of the associated measurements. Thus, there was a need for publicly available broadband performance data based on clearly defined, technically sound, and transparently conducted measurements and analyses.

In BDIA 2008, the US Congress required the FCC to better inform the US broadband community, including consumers, public interest groups, industry, academia, and the FCC itself, about the measured performance of broadband Internet access services. Collectively, BDIA 2008, together with the America COMPETES Reauthorization Act of 2010, and the National Broadband Plan, led to the introduction of the MBA program in late 2010. The FCC issued a Request for Quotation, and a contract was awarded to SamKnows Inc., to launch the *Measuring Broadband America* (MBA) program. During its initiation phase, members of the public and academic researchers provided comments on the proposed measurement solution [24]. The first MBA report was issued in early 2011 [33].

3 CONTEXT, GUIDING PRINCIPLES AND RELATED DESIGN REQUIREMENTS

3.1 Measurement Initiatives on Large Scale Networks

Numerous initiatives supporting work on Internet measurements have emerged over the last two decades. All have had distinct areas of emphases, with differing frameworks underlying their infrastructure, methodologies and funding. To put MBA measurements in context, we review a few other large-scale projects involving measurements, on which information is publicly available. As most such projects are or were exclusively focused on research, with overarching objectives centered on evolving new network architectures or services, these differ from the MBA case. In each instance, the resulting measurements are closely tied to the corresponding framework’s capabilities and principal aims. The measurements undertaken vary. Yet, to paraphrase one of the observations in [6], differences in measurement tools or systems are valuable, as each yields different insights, with no one test approach being necessarily better or worse than the rest. All infrastructures do not support every sort of measurement.An excellent overview of the Internet measurements arena can be found in [4]. The MBA program is specifically designed to obtain a measurements-based perspective on the quality of consumer broadband Internet access in the United States. This is in contrast to measurement initiatives featured in PlanetLab [17], GENI[28],[3] FABRIC [5], and RIPE Atlas.[4] With the exception of RIPE Atlas, the others mentioned either had, or still have, a research focus. PlanetLab was a global research network that supported the creation of new network services. Begun in 2002, it was aimed at countering “Internet ossification” and facilitated numerous highly significant research papers[5] on network architectures and measurement [1]. PlanetLab ceased operation in the US in 2020 but continues as PlaNetLab Europe. PlanetLab significantly informed and influenced its successor, the Global Environment for Network Innovation (GENI). GENI, also funded by the US National Science Foundation (NSF) starting in 2007, provides shared infrastructure with the goal of promoting advances in network science. GENI is in the process of evolving to FABRIC [5].Some of the measurements work supported on GENI include [2, 12, 59].

Of these other measurement-centric programs, RIPE Atlas resembles MBA most closely. Both rely on dedicated hardware-based measurement agents and perform active measurements. However, they differ significantly in scope and measurements. RIPE Atlas focuses on multi-country ISP operations and research projects. Réseaux IP Européens (RIPE) was formed in 1989 as an informal body of European ISPs looking to coordinate and exchange technical and operational information. The RIPE Network Coordination center (RIPE NCC) was formed in 1992. Today, RIPE functions as the Regional Internet Registry for Europe, parts of the Middle East and

[3]https://www.geni.net/
[4]https://atlas.ripe.net/about/
[5]https://planetlab.cs.princeton.edu/biblio.html

Central Asia. RIPE NCC started RIPE Atlas to help its participant countries' Local Internet Registries and ISPs monitor the state of the Internet in real time as well allow the volunteers hosting its roughly 11,000 globally deployed probes to run experiments on wide-scale Internet phenomena of interest to them, subject to a resource-management approach.[6] Although RIPE Atlas, like MBA, uses active measurements, and a standalone hardware-based test agent, its focus is on tracking global end-to-end Internet reachability (ping-based) and connectivity via traceroute, as well as other Internet infrastructural services such as DNS, NTP, SSL/TLS and HTTP. RIPE Atlas also offers utilities to track live BGP routing characteristics. In July 2022, RIPE Atlas covered 170[7] countries, while MBA is focused on consumer broadband performance in the United States. The RIPE Atlas system gathers a more limited set of measurements, e.g., round-trip times to root servers, reachability to such servers, and other measurements for operational tracking of Internet infrastructure services, compared to the more detailed *access* network performance data collected in the MBA test suite.

Ookla [46], and the Measurement Lab (MLab),[8] among others, are also engaged in the collection of Internet measurements. The data from these sources are convenience samples. They rely on visitors to websites or users of downloadable tools to conduct measurements, typically one-off. By their nature, these efforts cannot take the performance characteristics of the measurement agent into account, may be affected by in-home Wi-Fi performance limits and do not know whether other users are using the internet at the same time. Measurements cannot ensure uniform time-of day coverage and may be subject to sample bias, e.g., because people may be more likely to run speed tests when they experience performance issues. The MBA effort sought to avoid these and other confounding factors, as they can significantly impact the measurement of the access network, typically making the access network appear to perform significantly worse than it actually is.

3.2 Guiding Principles

The principal purpose of the MBA is to implement consistent, objective, clearly-specified measurements on United States broadband access networks, generating a dataset that can support longitudinal perspectives on performance. This informs all aspects of the MBA design, including its test approach, the selection of the measurement panel (ISP tiers and subscribers sampled), and the statistical and data processing pipelines used in the measurement analysis.

We now describe several characteristics of the tests and system design used in the MBA program. The particulars of the metrics are not explored here; they are covered in the Technical Appendix accompanying each MBA report, the most recent one of which is based on data collected in the fall of 2020 [44].

3.3 Focus on Access Network Performance

MBA aims to measure the performance of broadband Internet access services within the service provider's network. To do this it tries to operate its measurements between the relevant network 'edges': the service delivery point to the home and a provider network edge at the closest (in terms of latency) tier-1 Internet exchange point. This focus is because unlike measurement platforms that only offer browser-based tests initiated from a user-operated device such as a tablet or phone, the MBA platform aims to isolate measurements to the portion of Internet access network architecture that the ISP being tested is solely responsible for. The MBA measurement infrastructure's topology is therefore designed to eliminate or minimize confounding factors, including those from other sections of the end-to-end Internet encountered by the Internet traffic of home users. Examples are in-home Wi-Fi, transit networks, and multiple peering interconnects [22]. This also means that application performance measured by user-initiated, browser-based tools will likely be lower than the measurements collected by MBA.

Other factors that may explain differing performance results include variations in the capabilities of measurement platforms; the available Wi-Fi bandwidth; and differences in the measurement approach, for instance, whether TCP is used as the protocol for the test application; how many flows are initiated in parallel, the duration for which they are sustained, and the portions, if any, of the TCP start-up phase ignored in the throughput calculations.

User-initiated measurements may suffer from confirmation bias, as users may only initiate measurements when they perceive problems. Alternatively, the demographic of subscribers initiating measurements may not represent the overall ISP customer base. We discuss these issues in Section 3.5.

Design features that implement the MBA's focus on the access network include:

Characteristics of the in-home MBA test agent: MBA test agents run on standalone hardware. The versions currently deployed on the MBA panel (referred to as 'Whiteboxes' due to their chassis color) are based either on an off-the-shelf or custom-built consumer router, with custom firmware developed for conducting measurements, using a modified version of OpenWrt. The test agent is configured to operate as an Ethernet bridge and is connected to the home gateway, whether that home gateway is provided by the subscriber or the ISP. Panelists are asked to connect all of their devices, including Wi-Fi extenders if any, to the Whitebox. This enables the Whitebox (which is the measurement agent) to check for user traffic before launching its tests. This removes potential measurement distortion due to concurrent household Internet traffic. Newer versions of the test agent can generate test traffic loads of up to approximately 940 Mbps. Since the nominal broadband speed of participants is known, the test agents are upgraded if the current version cannot support the maximum expected service speed subscribed to in a particular panelist's home.The different test agent models now in deployment are matched to the tiers they are deployed for.

Placement and specifics of test servers: Most MBA tests run between a pair of endpoints: The MBA test agent and an MBA test server. Exceptions are the measurements to evaluate web, CDN, and DNS performance. Test agents for fixed broadband measurements are located within the homes of consumer volunteers ('panelists') who have elected to participate in the MBA program. MBA relies on two categories

[6]https://atlas.ripe.net/docs/credits/
[7]https://atlas.ripe.net/results/maps/network-coverage/
[8]https://speed.measurementlab.net/

of test servers: off-net and on-net servers. Off-net servers are located external to the ISP networks under test and are hosted by a commercial cloud provider (currently, ten server locations operated by Stackpath) and required to conform to FCC-specified program requirements on server performance, access speed, availability, and reliability. Off-net servers are located as close (in terms of network hops) as possible to the ISP network's topological borders, to reduce the impact of transit networks beyond the control of the ISP. With such server placements, test traffic traverses at most one transit network and two interconnects. MBA uses the same set of off-net test servers to evaluate all ISPs that participate in the MBA program. They are sited at well-connected interconnection points distributed across the country.

On-net servers, where available, are located within the ISP network, operated by that ISP, and used as an additional test target if the ISP configures them to conform to MBA test server requirements. Not all MBA ISPs provide on-net servers, but where they exist, they can provide useful comparisons to gauge whether peering, external connectivity, or other factors outside the ISP network impact measured performance. Similarly, comparing results for two off-net locations can also identify bottlenecks other than the access network. MBA reports only draw on measurements from tests run to off-net servers, but the raw data released by the program includes all test results.

3.4 Sampling

To obtain the MBA measurements, the FCC relies on consumer volunteers served by domestic residential broadband providers who have agreed to participate in the program. Due to the cost of the hardware test agents and the limited program budget, the program can only deploy a relatively small number of whiteboxes. Despite this, the sample size is still larger than those in most other projects — in 2021 the number deployed were 5951, of which the data collection requirements for the corresponding report used 2488. MBA thus includes only subscribers of the most popular service tiers. These tiers are picked based on the subscriber counts available to the FCC via the provider's Form 477 [21]. Since it takes several months between the collection of the Form 477 data and when the FCC has analyzed it, participating ISPs can also confidentially convey subscriber counts to the MBA contractor.

As noted earlier, unlike other measurement efforts, the MBA report also compares advertised to measured download and upload throughput and thus needs to ensure that the data collection has up-to-date information on the current service tier of each panelist, even as subscribers move, switch providers, change service tiers or providers adjust the speed of existing subscriber service plans. As offered tiers and subscriber counts change continuously, it takes manual effort, covering several months, to determine the tiers to be sampled, ensure that each tier has a sufficient number of measurement agents and that the existing panelists are mapped to the correct service tier. Instead of trying to continuously update the panel composition and nominal speed tier information for each panelist, the MBA designates a "reporting period" as the basis for the annual reports, defined as a contiguous thirty-day interval sometime during the months of September and October. The "reporting period" is bracketed by two phases of sample adjustment. First, before the reporting period, the participating carriers and FCC MBA team identify the most popular service tiers and recruit additional volunteers if needed, to ensure that each such tier is adequately represented in the MBA panel. After collecting the reporting period's data, if necessary, the MBA contractor ensures that tiers are correctly identified in the data sample. At this stage, a decision may be made to have the data analysis process exclude time periods affected by natural disasters since the MBA aims to present typical performance, not impairments caused by widespread network outages. In addition, only panelists contributing at least five days worth of data during the reporting period are included.

The FCC publishes the validated data from the reporting period, i.e., the data used to create the annual report. The FCC also makes the raw data available for the whole year, published in batches of three months and with a time lag of approximately three months.

3.5 Reliable Longitudinal Data Collection

MBA measurements are designed to offer a longitudinal and verifiable picture of broadband access, allowing the FCC, researchers and the public to see how broadband access evolves and to draw sound inferences across providers. MBA makes a number of choices that favor consistency, repeatability, transparency and panelist privacy, even though these choices impose costs or reduce the number of samples collected.

Controlled and consistent test platform: MBA, as noted, collects measurements using a dedicated hardware platform; other test approaches such as [18, 46] implement the test agent as JavaScript in browsers on a user's computer or as apps on smartphones. The use of a dedicated device as the MBA test agent eliminates the impact of variations in platform processor capabilities, operating systems, CPU utilization, or memory. Variability is reduced at the test servers by ensuring they conform to detailed specifications. Bauer et al. [6] discuss some of the pitfalls of using infrastructure that is less well-specified. For example, performance strongly depends on the TCP receive window, which may change across server or agent operating systems and versions.

Active measurements: Active measurements introduce controlled test traffic into a network and characterize the network based on observing how it performs. In contrast, passive measurements rely on observations on user traffic. To maximize consistency in measurement, and minimize the introduction of other variables that could depend on user traffic patterns or application-provider design choices, MBA uses active measurements. Active measurements come with trade-offs: they add to the network load, to avoid interfering with user traffic, are best not run when the subscriber is using the network, may draw down subscriber data usage quotas, and could receive preferential QoS treatment from the ISP. On the other hand, it is unlikely that passive measurements could determine whether degraded performance was caused by the ISP access network, the consumer's Wi-Fi, some other network component in the home or in the networks beyond the edge of the service provider network. While passive measurement analyses can strip any personally-identifiable information, such as destination IP addresses, from the measurement data, once it is collected, active measurements by definition neither use nor require the storing or processing of user network traffic.

Deferring measurements in the presence of user traffic: As described above, MBA tests involve active measurements and thus generate test traffic. MBA test agents in panelist homes defer measurements if they detect active use of the broadband connection, i.e., operate in idle-only mode. First, this avoids annoying the volunteer household by potentially interfering with its traffic and reducing consumer-perceived network performance. This reduces participant churn and potential complaints to the ISP caused by competition between measurement and user traffic. Secondly, the throughput measurements obtained would not accurately reflect the access capacity or latency if they are taken when the link into the home is shared with other traffic. Varying cross-traffic would also prevent obtaining consistent measurements. Browser-based measurement tools cannot detect competing traffic in the home.

The MBA test agent is configured as an Ethernet bridge, installed between the customer premise equipment (CPE), also called residential gateway (see Section 3.3) and the remainder of the home network. It also features a built-in Wi-Fi access point, and so can detect wireless traffic, including that destined to or from the residential gateway. If installed correctly, the MBA test agent senses all home network traffic destined to or from the Internet, and can defer measurements when other devices are active. Since these test agents are installed by the subscriber, cabling errors cannot be completely prevented. If unusually poor performance suggest a wiring mistake, the MBA contractor can contact the consumer and review the in-home setup.

Avoiding temporal bias: MBA measurements are automatic and randomized over six time periods that a 24 hour day is divided into, rather than user-initiated. Since broadband usage patterns and network congestion vary throughout the day and year, MBA designs its test schedules to measure the access network over the course of a day, divided into the aforementioned six consecutive four-hour periods that cover the entire twenty-four hour period, including the evening busy hours. If the agent detects household broadband activity at the beginning of a test attempt, the measurement device backs off and retries a few times. If the retries fail, it will defer further attempts to the next four-hour window. The agent conducts measurements every day of the year.

Ensuring privacy: As noted, active measurements avoid many of the privacy challenges that arise with passive measurements, as in the first place, there is no need to observe consumer-generated network traffic. Despite this the prospect of potentially unintentionally revealing the identity of the panelists may still raise privacy concerns, as the measurement pattern may allow deducing, after the fact, whether a panelist was using their broadband connection and the amount of traffic generated. Thus, the MBA program only reveals the census block or, if fewer than 1,000 people live in the census block, the census tract. The FCC does not know the names and addresses of the participants. The MBA contractor does need to know the address of the panelist in order to ship the whitebox. Every panelist has to explicitly consent to the privacy policy and ISP stakeholders have to agree, through the code of conduct, "not publish or make use of any test data or testing infrastructure in a manner that would significantly reduce the anonymity of collected data, compromise panelists privacy, or compromise the MBA privacy policy governing collection and analysis of data" [37].

3.6 Transparency and Reproducibility

Along with each report generated by MBA, the FCC publishes a technical appendix that describes the methodology, measurement techniques, statistical and analytical processes used for that report (e.g., [44] for the 11th report). It also releases the validated and raw data, SQL queries, R scripts, and Excel graphs. As described below, all aspects of the MBA, including the analyses and data in the report, are subject to discussion among stakeholders. FCC staff are also available to discuss sample composition, technical and data analysis details.

As all MBA test data is publicly released, researchers can use the data to analyze aspects of home broadband performance not covered by the FCC report. However, since the report does not indicate the relative popularity and thus weighting of tiers, they cannot reproduce the graphs in the report showing the average performance of each provider as a fraction of advertised speed. (The report does not break out individual tiers.) The data discloses the number of panelists for each service tier, but the panel size does not scale with the number of subscribers nor does the panel reflect the geographic distribution of the subscribers for the tier. Inclusion of the tier in the validated data does indicate that the tier is among the popular tiers for the provider. The number of agents for each tier is chosen to be sufficient to ensure statistical validity, with a minimum of 45 agents per tier. The report does not provide the subscription numbers to protect the commercial interests of the providers and encourage their participation in the program.

3.7 Program Organization and Governance

The MBA program is independent of other programmatic activities at the FCC. Specifically, MBA is not directly involved in broadband mapping or broadband subsidy programs. Keeping the program separate from rulemakings and subsidy programs helps retain necessary focus on the technical aspects of measurements and analyses.

Service provider participation in the MBA program is voluntary. For the December 2021 report, "[t]he evaluated ISPs collectively represent over 70% of U.S. residential broadband Internet connections." [43]. Since the start of the MBA program, providers have joined and left the program. For example, the satellite provider sector was represented in the program from the Third report (released 2013) through the Eighth Report (released 2018) [38]. Participating in the MBA allows providers to benefit from first-hand exposure to measurement techniques, practices and system architectures from across the technical community. Additionally, there are other potential positives to joining MBA, from a provider's perspective. For example, ISPs that choose to participate in the program may disclose their results as a sufficient representation of the performance their customers can expect to experience, which may be used to satisfy the FCC transparency rules for ISPs (Code of Federal Regulations (CFR) Title 47, Section 8.1[9]). ISPs may also use MBA results to meet the performance requirements for recipients of subsidies in two FCC-administered rural broadband efforts, the Connect America Fund (CAF) [26] and the Rural Digital Opportunity Fund (RDOF) [63].

Consumers, like providers, volunteer to be part of the MBA effort. Each consumer member of the test panel (panelist) hosts an MBA measurement agent in their home. Panelist recruitment could be effected by one or more of several methods-ISP outreach (to populate the sample for specific tiers), word of mouth, and an informational website [50]maintained by the MBA contractor. From the volunteer pool, a subset is chosen, to represent service tier and subscription demographics across the country. As a token of appreciation for participation, the consumer is emailed a monthly performance report reflecting a summary of the speed measurements taken as well as access to a personal dashboard for more detailed ongoing information about the measurements on their service. Both ISPs and consumers can cease MBA participation at will.

MBA is distinguished by the level of dialog and consensus-building that influences how the program operates. Stakeholder discussions occur at monthly meetings, referred to as "MBA Collaboratives," organized by the FCC. Beyond FCC and contractor staff, these meetings typically include staff from the ISPs whose networks participate in MBA measurements, academics, public interest groups, other organizations, technical personnel from equipment vendors, non-participating ISPs, and members of the public. During these meetings, the FCC communicates progress notes, including MBA program schedules and current topics of interest to the program. Attendees ask questions of the FCC MBA team as to measurements and program operations in general. Examples of the types of interchange include discussions on measurement infrastructure concerns and introduction of new or improved metrics or methodology. The FCC distributes all meeting materials and minutes through its electronic comment filing system (ECFS). We are not aware of any other network measurement program that maintains this level of interaction among stakeholders.

A number of stakeholder suggestions have been implemented. For example, Bauer et al. [7] suggested subdividing the throughput test interval to help assess the effect of PowerBoost™ technology deployed by cable-based broadband providers. This was implemented until tests showed that the technology use was being deprecated, at which point the speed test durations were reduced sufficiently to maintain their fidelity but reduce data consumption. In another example, in response to ISP requests, the FCC established a technical working group to evaluate methodologies to measure performance at speeds near and beyond a gigabit per second and engaged the engineering community in discussions as to experimental evaluations and recommendations.

A further unique aspect of the MBA program is that, starting in 2013, it has encouraged and supported use of the MBA infrastructure by US academics, allowing researchers to experimentally explore Internet architectures and measurements, including topics of interest to the non-academic broadband community. In early 2015, the FCC formalized this activity as "MBA Assisted Research Studies" (MARS) [45], to simplify the use of the MBA platform as an Internet observatory offering perspectives on US residential internet access. The diversity of networks and technology, as well as the number of consumer vantage points offered by the MBA measurement infrastructure, have made it particularly valuable in understanding evolving network characteristics. Researchers at a number of U.S. universities [10, 20, 37, 48, 52, 55–58] have leveraged MARS or MBA data to conduct experimental research, to try out new measurement algorithms and techniques and evaluate emerging Internet phenomena on production networks, at a scale not available elsewhere. MARS experiments generally follow the process described in [45], where after the FCC staff evaluates feasibility and fit, a project is presented to the Collaborative to gauge community interest, before proceeding to deploy it.

4 MEASUREMENT AND INFRASTRUCTURE EVOLUTION

In the first year the MBA program issued a report (2011), 25 Mb/s was considered blazingly fast. In 2021, more than 10% of subscribers were provisioned for 1 Gb/s or more [47]. Both MBA test agents and measurement servers have evolved accordingly. The test servers now in use are subject to commercial service level agreements and provisioned for the speeds they are required to measure and the test traffic capacity they must serve. The test agent models in homes are calibrated and deployed according to the measurement capacity each needs to support, and new models are introduced to keep up with access technology as it progresses. Metrics such as data usage have evolved, as the earliest deployed models of test agents did double duty as both home router and measurement device, but the models introduced since 2013 function as bridges inline to the home router, making it necessary to change how the metric is implemented [49].

The throughput metric offers another example of how MBA measurements have evolved, with the number of simultaneous flows used in the measurement process increasing from three TCP streams to eight, while the test duration has been reduced from 30 seconds to 10. To respond to concern about test data consumption, and potentially reaching data caps imposed by some providers for some tiers, the entire test schedule was modified in 2020 to reduce the frequency of tests while still maintaining coverage across the 24x7x365 period [39].

The MBA reports and data also document the evolution in the average service speeds on offer in the United States. As per the 11th MBA report [43], in September-October 2020 the weighted average advertised download speed was 193.9 Mb/s among the measured ISPs, which represents a 33% increase compared to the average, in September-October 2019, of 146.1 Mb/s (as in the Tenth MBA Report [40]) and a 166% increase compared to the 2017 metric of 72.9 Mb/s shown in the Eighth MBA Report [38]. The measured-to-advertised throughput ratios, while not the main program focus, have also increased over time. In a 2011 analysis [33], the ratio across all ISPs was 82%. In the report issued in July 2012, this changed to 87% [35] and to 96-97% in 2013 [36].

[9] https://www.law.cornell.edu/cfr/text/47/8.1

5 LESSONS LEARNED

The MBA program, as an exercise in characterizing national broadband access performance, has both met many of the program objectives and faced its share of challenges and criticisms.

5.1 What Has Worked Well

For over a decade, the MBA program has met its objective of conducting ongoing measurements and compiling and making the resulting data publicly available, thus contributing significantly to an understanding of performance and service trends in U.S. fixed broadband access [4, 7, 11, 39]. The FCC has accomplished this despite MBA's significant budgetary, staffing, and other constraints. The MBA program has also remained flexible and open enough to incorporate suggestions for change and improvements, such as those listed in Sections 3.7 and 4, with other examples found on the public docket maintained to document MBA Collaborative meetings.

5.1.1 Program continuity and ensuing benefits for longitudinal research and inferences. Broadband data is particularly valuable for longitudinal research (see [53]). Archived data is often as important as the most recent data set. Many interesting research questions explore how metrics have changed, not just their current state. Purely research-focused network measurement efforts tend to be short-lived, as grant agency funding cycles, lack of suitable staff, and publication pressures make it difficult to maintain consistent longitudinally-focused initiatives within an academic setting. (Measurement-focused organizations such as CAIDA may last longer.) The small number of long-lived active measurement projects such as RIPE Atlas, established in 2010, and CAIDA Ark,[10] begun in 2007, are hosted within non-academic research organizations.

It is difficult to imagine an effort such as MBA being conducted successfully, for more than a decade, without either the involvement of the national telecommunications regulatory body or a degree of cooperation from broadband service providers. The MBA program has successfully integrated both of these. In contrast to research infrastructure projects, which are largely tracked only by academics, the MBA's operation and evolution is followed by a more diverse group of interested parties, as the attendee lists of the Collaborative meetings illustrate.

5.1.2 Collaboration between industry, researchers and regulator. Individual researchers often struggle to work with large companies, while having a government agency and its contractor work with ISPs has proven itself to be sustainable. Recruiting participants at scale and replacing them as needed, particularly if panelists are to include those outside the technical community and the data is to consist of more than just convenience samples, is challenging; the participation of ISPs, combined with the public-interest mandate, makes this much easier for MBA than for measurement projects run by commercial or academic entities.

5.1.3 The program as an enabler of substantial and diverse research efforts. The raw data generated and made available by the program has been the basis of a number of research papers, cited earlier.

The ability to reproduce analyses is increasingly seen as a key concern for computer systems research [3]. The ready availability of MBA data allows other researchers to reproduce MBA-derived analysis results, even if the data itself cannot be readily replicated. Unlike data that is offered by industry or individual research groups, the regularly released FCC MBA data comes with few restrictions and is available to all researchers.

The networking community has significant participation in MARS,[11] including a joint FCC-NSF workshop on measuring QoE [16] held in 2015. The MARS effort has also yielded well-received research results, such as studies on narrowing down whether congestion exists inside or outside the home, detecting carrier-grade NATs, and a performance study comparing emerging DNS privacy-enhancing measures completed in mid-2020 [29, 45].

Both the raw data repository and the MARS activity are currently unique among regulator-managed network measurement efforts.

5.2 Critiques and Limitations

MBA, like all such programs, has had to make trade-offs as necessitated by its design approach, finite resources, its reliance on voluntary participation by both providers and consumer panelists, and the practicalities associated with measurements that require the installation of test agents in panelists' homes.Some of the issues that have been pointed out due to these factors include the following.

5.2.1 Do the metrics reflect user experience? Some researchers have noted that the test agents or 'Whiteboxes' only conduct measurements when consumers are not accessing the Internet, and the measurements therefore do not capture the actual consumer network experience [54]. As explained in Section 3.5, active idle-only measurements reduce measurement noise, avoid degrading the consumer experience and reduce panelist churn. If active measurements were to be conducted when the household is actively using the Internet, the measured throughput would indicate the bandwidth *share* available to the measurement traffic, in the presence of that household's traffic, rather than the access network's performance; and measured latency might reflect the load imposed by the household's application traffic, rather than the ISP's propagation and queuing delays. Note that one of the MBA metrics is delay-under-load, which if allowed to run in the absence of household Internet traffic, will reflect buffering in the home access router, including any buffer bloat.

Multi-stream TCP measurements reflecting steady-state throughput, such as those used for the MBA throughput metrics, also may not reflect actual application performance. For example, many interactive audio-video applications adjust their throughput using proprietary algorithms and thus may yield a different quality-of-experience that may not be proportional to the TCP results. Passive measurements can capture this performance; a better understanding of the relationship between active metrics such as throughput or less and application performance would be helpful. However, such applications are often rate-limited by the server or the ability

[10] https://www.caida.org/projects/ark/

[11] https://www.fcc.gov/general/mba-research-code-conduct#publications-mars-mba

of the receiving side to process or render the data stream, so distinguishing network impairments from application-level controls remains challenging.

5.2.2 Documentation and reproducibility of MBA metrics. While the technical appendix describes the measurement metrics in general terms, the metrics are not traceable to published standards such as the performance measurement standards published by the IETF IPPM working group.[12] In the program's earliest years, source code for the measurements was available via a GPL license[13] but this is no longer the case. The MBA contractor may at one point have made the code available to academics under a non-disclosure agreement as was the experience of one of the authors of this paper, but whether such arrangements are still supported is unclear.The MBA is not resourced to support open-source efforts from elsewhere.

5.2.3 Delay in data availability. Nationwide network measurement systems such as MBA can be useful for purposes other than characterizing the performance of consumer network services. For example, the data capturing monthly download and upload values have been used to track anonymized and aggregate changes in household behavior, such as work-from-home, during the early days of the COVID-19 pandemic [20]. But the long lag times sometimes involved before MBA releases its 'raw' data sets reduces the timeliness and accuracy of such uses.

5.2.4 Limited geographic coverage. Partly due to the program design, and partly due to budget constraints, the measurement coverage afforded by the MBA system is non-uniform in both time and space. However, this objective remains elusive to other measurement efforts as well. As described in Section 6, each tier is represented by no more than a few hundred measurement agents, each of which measures performance only intermittently. Since both consumer and ISP participation are voluntary, and the measurement agents are expensive, it is also difficult to make the distribution of agents reflect the geographic distribution of subscribers. For many tiers, a whole state may be represented by a single measurement agent. For example, the data distribution reflecting performance results for a single tier can vary widely across test agents [51] as sections of a single service provider's network may suffer from different levels of congestion. Metrics like the 80/80-consistent speed attempt to capture this spatial and temporal variability, where the 80/80-consistent speed "refers to the minimum speed that was experienced by at least 80% of panelists for at least 80% of the time during the peak periods." [44] Given the significant spatial variability among service territories for large carriers, it may be helpful to ensure that each such area hosts several measurement agents. Limited resources for deploying measurement agents may make this difficult; thus, as noted below, alternative approaches that scale better should be pursued.

5.2.5 Limited coverage of broadband subscriber population. The budgetary and policy constraints of the MBA program impact how geographic areas, consumer categories, speed tiers, and ISPs are represented in its sample. Validating speed tiers is among the unique strengths of the MBA, and supports one of its original objectives, to compare advertised to actual throughput performance. This is not always straightforward: the technical appendix to the first report notes that 4% of volunteers misidentified their ISPs and that 9% were re-allocated to a different tier [34]. However, the validation process also requires active participation by ISPs, even though the test agent is provided by the FCC contractor and installed by the subscriber. Smaller ISPs may not see the value of participation, given the effort required. Participation by ISPs, as mentioned, is voluntary, and several large and mid-sized providers as well as whole categories, including Wireless Internet Service Providers (WISPs) and satellite providers, currently are not represented in the program at all. Since rural areas are more likely to be served by smaller ISPs, such as WISPs, and satellite providers, rural consumers are likely to be underrepresented in the MBA sample. In addition, some larger ISPs have dropped out of the program, expressing dissatisfaction with how their performance is represented.

It is difficult to estimate the precise fraction of broadband subscriptions represented by MBA panelists. For the Eleventh MBA Report (p. 22), "The evaluated ISPs collectively represent over 70% of U.S. residential broadband Internet connections." A rough calculation drawing on the Leichtman Research Group (LRG) 4Q2020 subscriber data [30] has the ISPs from that MBA report represent 78% of the total broadband subscriber market, but since the MBA panels only represent tiers that cover at least the top 80% of the subscriber based for each ISP, the overall percentage covered is likely less, depending on whether these popular tiers cover significantly more than 80% of the subscriber base or not.

5.2.6 The potential for skewed measurement results. There have been concerns that it would be easy for ISPs to skew the results by prioritizing measurement traffic, given that the measurement servers' IP addresses are released, and that the ISPs help verify the subscriber speed tier, i.e., are aware which households hosts test agents. Although it is not a technical deterrent, participating ISPs sign a code of conduct that requires them to "not act, nor fail to act, if the intended consequence of such act or omission is to enhance, degrade, or tamper with the results of any test for any individual panelist or broadband provider." [26] In addition, the MBA contractor deploys a control sample of agents and test servers, that are not part of the MBA program [50]. These operate as a check on MBA-observed performance.

5.2.7 Impact of federal guidelines on publications involving MARS. The utilization by academics of the MBA program's MARS component (Section 3.7) through which they can run custom experiments on MBA infrastructure, has proven useful, judging by the number of publications that have resulted. However, there have been frustrations expressed by some who have availed of this facility, stemming from perceived restrictions on publishing the results of their research, in particular regarding identification of the various providers on whose products and services these experiments were conducted. The related MARS provisos result from Federal law [25], not the MBA program, and the associated MARS Code of Conduct [42] expressly allows full publication of results to be explored, through independent agreements between the academics and the concerned service providers.

[12] https://datatracker.ietf.org/wg/ippm/about/
[13] https://www.fcc.gov/general/open-source-testing-software-source-code

5.2.8 Resource Challenges. Competing priorities and the need to balance available resources across many demands within the Commission inevitably limited what could be accomplished.

6 WHAT NEXT?

The future of the MBA program will in part depend on the utility of the program as seen by public interest organizations, researchers, carriers and federal broadband efforts. Federal agencies will likely perceive the program as more valuable if it is seen as supporting other programmatic objectives, such as universal broadband deployment and broadband consumer labels.[14] We believe that the MBA's purely technical focus offers significant value and accordingly it is critical that this be retained.

6.1 Broadening Participation

To make the MBA program more representative of the total United States broadband subscriber base, the organizers could allow direct consumer volunteer participation, even if the speed tier information may not be as reliable. This would be best facilitated by making available a range of test agents embedded in consumer routers, once the software has been validated to accurately perform measurement functions while also operating as residential gateways [27]. Indeed, the FCC contractor offers firmware upgrades for CPE,[15] but these would have to be installed by the provider, not the customer.

To further scale up MBA and to reduce the cost of the program, network measurement capabilities should be integrated into consumer network devices in an interoperable manner, e.g., using standards-based efforts such as the IETF LMAP specifications [15]. (LMAP is now also part of the Broadband Forum User Services Platform (TR-369, a successor to TR-069) [14].) Using measurement capabilities built into home routers or the network edge device avoids having to deploy standalone measurement agents and, in particular, the complexities of getting the wiring and placement details for the measurement device correct. It would also be helpful if CPE could query the provider network to obtain the configured service speed and technology, e.g., via a DHCP option or a TR-069 parameter.

6.2 Measuring Network Reliability

As access network speeds increase and consumer networks become more critical to work-from-home activities, the user community is likely to value network reliability as much as performance. Indeed, even though not an original design goal, MBA data has been used to characterize network reliability [9]. Since networks have life-safety functions, the FCC has long required voice service providers to report large scale outages [23], activating special reporting systems during natural disasters.[16] The MBA test servers could "see" larger-scale outages instantaneously, but are currently not set up to convey this information to the FCC, other agencies, or the public in real time. Besides requiring more promptly available data, a broadband outage monitor would likely also require more test agents, particularly if the goal is to track the impact of natural disasters and post-disaster recovery. For example, for the 11th report, 2,488 Whiteboxes reported data, i.e., on average less than one per county. (At the time of writing, for reasons mentioned above, it has not been possible to ensure that each county has at least one Whitebox.) To improve reliability monitoring, Whiteboxes could continue to measure, but label, DNS resolution and ICMP ping even when the home network is in use as these low-bandwidth and low-usage measurements do not impair the home network experience.

6.3 Program Evolution

We believe that the MBA program is most likely to be sustainable and continue to meet the need for basic residential network access performance statistics if it is funded consistently, is focused on measurement technology development and evolution, rather than program-specific measurements related to deployment subsidies, and maintains close ties to the academic research community. It should make data available more quickly, improve network coverage, and better integrate with consumer and possibly network edge elements. The raw data stores should be organized to improve access and allow visualization. As research interests in broadband performance, adoption and network reliability increase, the program should adjust the metrics it collects.

Instead of treating each measurement program as self-contained, it may be better to see if results can be correlated and the mission and objectives of the different efforts can be coordinated even as they pursue complementary goals. This may also improve the odd of securing long-term funding support. Additional private-public partnerships may be worth exploring.

7 SETTING UP MEASUREMENT PROGRAMS

The MBA experience may offer lessons for other organizations interested in setting up a measurement program evaluating national or regional communications infrastructures. Some of the questions that need to be addressed include:

- Beyond upload and download speeds, what other metrics should be gathered, analyzed and reported? What measurement standards should be adopted?
- How should the system be designed, set up and managed to ensure consistent technical soundness even as the measurement and measured infrastructures change? Should an external entity, such as an academic or government research lab, evaluate whether the measurement technology is technically sound?
- What is a representative panel? How variable are system components by provider, geography and time?
- How long will it take to recruit volunteers and replenish panelist pools where coverage is lacking?
- How would the logistics of delivering measurement devices, setting up supporting equipment and answering questions by the panel volunteers be handled?
- How should the data repository be set up and what services would best facilitate its use, both for more casual users and in-depth research?
- How often should measurements be conducted? Should measurements combine continuous data gathering and periods of higher validation?

[14] https://www.fcc.gov/broadbandlabels
[15] https://samknows.com/technology/agents
[16] https://www.fcc.gov/general/disaster-information-reporting-system-dirs-0

- What level of funding is sufficient to deploy and manage the infrastructure and an ongoing measurement program?
- Who is going to manage the program, validate data, analyze the data sets and draft reports?
- What governance and feedback mechanisms will be used to adapt the program to evolving technologies and measurement needs? How can the research and policy community be engaged, both during the initial concept phase and on a continuing, long-term basis?

8 SUMMARY

Government agencies have long played a critical role in collecting "boring" data continuously, across all sectors of the economy. Indeed, the FCC has collected telecommunications data for decades, albeit mostly based on industry filling out forms. Recently, the Open Data efforts and corresponding legal obligations [19] have further made government-collected data a helpful source for researchers. The MBA program offers an example, now backed by more than a decade of operational experience, of how a federal program grounded on empirical foundations can be established, sustained and managed. The collaborative governance model has been able to facilitate the operation and evolution of the program. It also illustrates some of the challenges of programs measuring the performance of commercial entities by the regulator on a voluntary basis and with a fairly limited budget. Given that network performance and reliability will remain key attributes of this critical infrastructure, building measurement capabilities into the infrastructure and CPE, rather than having to add it with significant effort, may well be the next challenge.

ACKNOWLEDGMENTS

The authors wish to acknowledge the dedication and perseverance of the FCC staff and management who have worked on this program, as well as the Commission's leadership for their support of it.

REFERENCES

[1] V.K. Adhikari, Yang Guo, Fang Hao, Volker Hilt, Zhi-Li Xhang, Matteo Varvello, and Mortiz Steiner. 2015. Measurement Study of Netflix, Hulu, and a Tale of Three CDNs. *IEEE/ACM Transactions on Networking* 23, 6 (2015), 1984–1997.

[2] A.C Babaoglu and R Dutta. 2014. A GENI Meso-Scale Experiment of a Verification Service. In *2014 Third GENI Research and Educational Experiment Workshop*. GENI program, 65–68.

[3] Vaibhav Bajpai, Mirja Kühlewind, Jörg Ott, Jürgen Schönwälder, Anna Sperotto, and Brian Trammell. 2017. Challenges with Reproducibility. In *Proceedings of the Reproducibility Workshop*. Association for Computing Machinery, Los Angeles, CA, USA, 1–4. https://doi.org/10.1145/3097766.3097767

[4] Vaibhav Bajpai and Jürgen Schönwälder. 2015. A Survey on Internet Performance Measurement Platforms and Related Standardization Efforts. *IEEE Communications Surveys & Tutorials* 17, 3 (2015), 1313–1341. https://doi.org/10.1109/COMST.2015.2418435

[5] Ilya Baldin, Anita Nikolich, James Griffioen, Indermohan Inder S. Monga, Kuang-Ching Wang, Tom Lehman, and Paul Ruth. 2019. FABRIC: A National-Scale Programmable Experimental Network Infrastructure. *IEEE Internet Computing* 23, 6 (2019), 38–47. https://doi.org/10.1109/MIC.2019.2958545

[6] Steven Bauer, David Clark, and William Lehr. 2010. Understanding Broadband Measurements. In *TPRC 2010*. TPRC, Arlington, VA, 38 pages. https://ssrn.com/abstract=1988332

[7] Steven Bauer, David Clark, and William Lehr. 2011. PowerBoost. In *Proceedings of the 2nd ACM SIGCOMM Workshop on Home Networks*. Association for Computing Machinery, Toronto, Ontario, Canada, 7–12. https://doi.org/10.1145/2018567.2018570

[8] David Belson. 2010. Akamai State of the Internet Report, Q4 2009. *SIGOPS Oper. Syst. Rev.* 44, 3 (Aug. 2010), 27–37. https://doi.org/10.1145/1842733.1842738

[9] Zachary Bischof, Fabian Bustamante, and Nick Feamster. 2018. Characterizing and Improving the Reliability of Broadband Internet Access. In *TPRC 46: The 46th Research Conference on Communication, Information and Internet Policy 2018*. SSRN, Washington, DC, 15 pages. https://doi.org/10.2139/ssrn.3121942

[10] Zachary S. Bischof and Fabián E. Bustamante. 2014. A Time for Reliability: The Growing Importance of Being Always On. In *Proceedings of the 2014 ACM Conference on SIGCOMM (poster)*. Association for Computing Machinery, Chicago, Illinois, USA, 131–132. https://doi.org/10.1145/2619239.2631460

[11] Zachary S. Bischof, Fabian E. Bustamante, and Rade Stanojevic. 2014. Need, Want, Can Afford: Broadband Markets and the Behavior of Users. In *Proceedings of the 2014 Conference on Internet Measurement Conference*. Association for Computing Machinery, Vancouver, BC, Canada, 73–86. https://doi.org/10.1145/2663716.2663753

[12] Ethan Blanton, Sarbajit Chatterjee, Sriharsha Gangam, Sumit Kala, Deepti Sharma, Sonia Fahmy, and Puneet Sharma. 2012. Design and evaluation of the S3 monitor network measurement service on GENI. In *2012 Fourth International Conference on Communication Systems and Networks*. IEEE, 1–10. https://doi.org/10.1109/COMSNETS.2012.6151327

[13] Peter Bowen and Shawn Hoy. 2009. *Broadband Performance (OBI Technical Paper 4)*. Technical Report. Federal Communications Commission. Retrieved January 9, 2022 from https://www.fcc.gov/document/obi-technical-paper-no-4-broadband-performance

[14] Broadband Forum. 2019. User Services Platform (USP). https://www.broadband-forum.org/download/TR-369.pdf

[15] Trevor Burbridge, Philip Eardley, Marcelo Bagnulo, and Jürgen Schönwälder. 2017. Information Model for Large-Scale Measurement Platforms (LMAPs). RFC 8193. https://doi.org/10.17487/RFC8193

[16] Fabian E. Bustamante, David Clark, and Nick Feamster. 2017. Workshop on Tracking Quality of Experience in the Internet: Summary and Outcomes. *SIGCOMM Comput. Commun. Rev.* 47, 1 (Jan. 2017), 55–60. https://doi.org/10.1145/3041027.3041035

[17] Brent Chun, David Culler, Timothy Roscoe, Andy Bavier, Larry Peterson, Mike Wawrzoniak, and Mic Bowman. 2003. PlanetLab: An Overlay Testbed for Broad-Coverage Services. *SIGCOMM Comput. Commun. Rev.* 33, 3 (July 2003), 3–12. https://doi.org/10.1145/956993.956995

[18] Comcast Xfinity. 2022. Xfinity Speed Test – Let's check your Internet Speed. Retrieved January 10, 2022 from https://speedtest.xfinity.com/

[19] U.S. Congress. 2019. Open, Public, Electronic, and Necessary Government Data Act (OPEN Government Data Act). Retrieved January 14, 2022 from https://www.congress.gov/115/plaws/publ435/PLAW-115publ435.pdf

[20] Jessica De Oliveira Moreira, Amey Pasarkar, Wenjun Chen, Wenkai Hu, Jan Janak, and Henning Schulzrinne. 2020. Social Distancing and the Internet: What Can Network Performance Measurements Tell Us?. In *TPRC48: The 48th Research Conference on Communication, Information and Internet Policy*. SSRN, online, 12 pages. https://doi.org/10.2139/ssrn.3748153

[21] FCC Office of Economics and Analysis. 2021. *Form 477 Orders and Releases*. Federal Communications Commission. Retrieved January 10, 2022 from https://www.fcc.gov/general/form-477-orders-and-releases

[22] Nick Feamster and Jason Livingood. 2020. Measuring Internet Speed: Current Challenges and Future Recommendations. *Commun. ACM* 63, 12 (Nov. 2020), 72–80. https://doi.org/10.1145/3372135

[23] Federal Communications Commission. 2004. 47 CFR Part 4: Disruptions to Communications. Retrieved January 11, 2022 from https://www.ecfr.gov/current/title-47/chapter-I/subchapter-A/part-4

[24] Federal Communications Commission. 2010. Comment sought on residential fixed broadband services testing and measurement solution. Retrieved January 9, 2022 from https://docs.fcc.gov/public/attachments/DA-10-670A1.pdf

[25] Federal Communications Commission. 2017. 47 CFR 0.457: Records not routinely available for public inspection. Retrieved January 10, 2022 from https://www.ecfr.gov/current/title-47/chapter-I/subchapter-A/part-0/subpart-C/subject-group-ECFR6eb707181ed88ec/section-0.457

[26] Federal Communications Commission. 2018. *WCB, WTB, and OET Adopt Performance Measures for CAF Recipients*. Order DA-18-710 (33 FCC Rcd 6509 (10)). FCC. https://www.fcc.gov/document/wcb-wtb-and-oet-adopt-performance-measures-caf-recipients

[27] Allen Flickinger, Carl Klatsky, Atahualpa Ledesma, Jason Livingood, and Sebnem Ozer. 2021. Improving Latency with Active Queue Management (AQM) During COVID-19. *CoRR* abs/2107.13968 (2021), 19 pages. arXiv:2107.13968 https://arxiv.org/abs/2107.13968

[28] James Griffioen, Zongming Fei, Hussamuddin Nasir, Xiongqi Wu, Jeremy Reed, and Charles Carpenter. 2014. Measuring experiments in GENI. *Computer Networks* 63 (2014), 17–32.

[29] Austin Hounsel, Paul Schmitt, Kevin Borgolte, and Nick Feamster. 2021. Can Encrypted DNS Be Fast?. In *Passive and Active Measurement*, Oliver Hohlfeld, Andra Lutu, and Dave Levin (Eds.). Springer International Publishing, online, 444–459. https://doi.org/10.1007/978-3-030-72582-2_26

[30] Leichtman Research Group (LRG) 2021. *4Q2020 Research Notes*. Leichtman Research Group (LRG). Retrieved July 9, 2022 from https://www.leichtmanresearch.com/wp-content/uploads/2021/01/LRG-Research-Notes-4Q-2020.pdf

[31] OECD. 2011. *National Broadband Plans.* OECD Digital Economy Papers 181. OECD Publishing, Paris. https://doi.org/10.1787/5kg9sr5fmqwd-en
[32] Ofcom Market Intelligence Team. 2009. *UK Broadband Speeds 2009: Consumers' experience of fixed-line broadband performance.* Research Report. Ofcom, London, UK. https://www.ofcom.org.uk/__data/assets/pdf_file/0013/40324/broadbandspeeds.pdf
[33] Office of Engineering and Technology. 2011. *Measuring Broadband America - August 2011.* Federal Communications Commission. Retrieved January 10, 2022 from https://www.fcc.gov/reports-research/reports/measuring-broadband-america/measuring-broadband-america-august-2011
[34] Office of Engineering and Technology. 2011. *Technical Appendix to the Eleventh Fixed MBA report.* FCC. https://data.fcc.gov/download/measuring-broadband-america/2021/Technical-Appendix-fixed-2021.pdf
[35] Office of Engineering and Technology. 2012. *Measuring Broadband America - July 2012.* Federal Communications Commission. Retrieved January 10, 2022 from https://www.fcc.gov/reports-research/reports/measuring-broadband-america/measuring-broadband-america-july-2012
[36] Office of Engineering and Technology. 2013. *Measuring Broadband America - February 2013.* Federal Communications Commission. Retrieved January 10, 2022 from https://www.fcc.gov/reports-research/reports/measuring-broadband-america/measuring-broadband-america-february-2013
[37] Office of Engineering and Technology. 2018. *MBA Research Code of Conduct and Publications that have utilized longitudinal data or MARS.* Federal Communications Commission. Retrieved January 10, 2022 from https://www.fcc.gov/general/mba-research-code-conduct#publications-mars-mba
[38] Office of Engineering and Technology. 2018. *Measuring Fixed Broadband - Eighth Report.* Federal Communications Commission. https://data.fcc.gov/download/measuring-broadband-america/2018/2018-Fixed-Measuring-Broadband-America-Report.pdf
[39] Office of Engineering and Technology. 2020. *Measuring Fixed Broadband - Ninth Report.* Federal Communications Commission. https://www.fcc.gov/reports-research/reports/measuring-broadband-america/measuring-fixed-broadband-ninth-report
[40] Office of Engineering and Technology. 2020. *Measuring Fixed Broadband - Tenth Report.* Federal Communications Commission. https://data.fcc.gov/download/measuring-broadband-america/2020/2020-Fixed-Measuring-Broadband-America-Report.pdf
[41] Office of Engineering and Technology. 2020. *Technical Appendix to the Tenth Fixed MBA report.* FCC. http://data.fcc.gov/download/measuring-broadband-america/2020/Technical-Appendix-fixed-2020.pdf
[42] Office of Engineering and Technology. 2021. *FCC Measuring Broadband America Program: Research and Special Studies Stakeholders Code of Conduct.* Federal Communications Commission. https://www.fcc.gov/sites/default/files/fcc_mars_code_of_conduct_2021_sample_forviewing.pdf
[43] Office of Engineering and Technology. 2021. *Measuring Fixed Broadband - Eleventh Report.* Federal Communications Commission. https://www.fcc.gov/reports-research/reports/measuring-broadband-america/measuring-fixed-broadband-eleventh-report
[44] Office of Engineering and Technology. 2021. *Technical Appendix to the Eleventh Fixed MBA report.* FCC. https://data.fcc.gov/download/measuring-broadband-america/2021/Technical-Appendix-fixed-2021.pdf
[45] Office of Engineering and Technology, Federal Communications Commission. 2021. *MBA Assisted Research Studies.* Office of Engineering and Technology, Federal Communications Commission. Retrieved January 9, 2022 from https://www.fcc.gov/general/mba-assisted-research-studies
[46] Ookla. 2022. Speedtest by Ookla - The Global Broadband Speed Test. Retrieved January 10, 2022 from https://www.speedtest.net/
[47] OpenVault 2021. *OVBI: Subscriber Speed Increases Driving Higher Median Broadband Usage.* OpenVault. Retrieved January 9, 2022 from https://openvault.com/ovbi-subscriber-speed-increases-driving-higher-median-broadband-usage/
[48] John P. Rula, Zachary S. Bischof, and Fabian E. Bustamante. 2015. Second Chance: Understanding Diversity in Broadband Access Network Performance. In *Proceedings of the 2015 ACM SIGCOMM Workshop on Crowdsourcing and Crowdsharing of Big (Internet) Data (C2B(1)D '15).* Association for Computing Machinery, London, United Kingdom, 9–14. https://doi.org/10.1145/2787394.2787400
[49] SamKnows. 2018. *Technical Methodology: Data Usage.* FCC. Retrieved April 28, 2022 from https://transition.fcc.gov/oet/mba/Data-Usage-Technical-Methodology_2018-08-24_Final-v1.3.pdf
[50] SamKnows. 2022. *How do you make sure ISPs aren't optimizing individual customers?* FCC. https://www.measuringbroadbandamerica.com/program-faqs#how-do%20you%20make%20sure%20isps%20aren%E2%80%99t%20optimizing%20individual%20customers?
[51] Henning Schulzrinne, Walter Johnston, and Andreas Freund. 2022. *Notice of Ex-Parte Meeting for Docket CG 22-2.* Columbia University. Retrieved July 8, 2022 from https://www.fcc.gov/ecfs/search/search-filings/filing/1041895940875
[52] Henning Schulzrinne, Walter Johnston, and Andreas Carlos Freund. 2022. Your Mileage May Vary or Performance You Can Count On: What Should Broadband Consumer Labels Measure? *SSRN (TPRC 2022)* (Aug. 2022), 25 pages. https://doi.org/10.2139/ssrn.4178758
[53] Henning Schulzrinne and Marie-José Montpetit. 2021. NSF Broadband Research 2020 Report. *CoRR* abs/2111.03779 (2021), 34 pages. arXiv:2111.03779 https://arxiv.org/abs/2111.03779
[54] Shalini Ramachandran, Lillian Rizzo, and Drew FitzGerald. 2019. *Your Internet Provider Likely Juiced Its Official Speed Scores.* The Wall Street Journal. Retrieved April 28, 2022 from https://www.wsj.com/articles/its-hard-to-trust-the-numbers-internet-providers-inflate-official-speed-results-11576171855
[55] Srikanth Sundaresan, Walter de Donato, Nick Feamster, Renata Teixeira, Sam Crawford, and Antonio Pescapè. 2011. Broadband Internet Performance: A View from the Gateway. In *Proceedings of the ACM SIGCOMM 2011 Conference.* Association for Computing Machinery, Toronto, Ontario, Canada, 134–145. https://doi.org/10.1145/2018436.2018452
[56] Srikanth Sundaresan, Nick Feamster, and Renata Teixeira. 2016. Home Network or Access Link? Locating Last-Mile Downstream Throughput Bottlenecks. In *Passive and Active Measurement*, Thomas Karagiannis and Xenofontas Dimitropoulos (Eds.). Springer International Publishing, Heraklion, Greece, 111–123. https://doi.org/10.1007/978-3-319-30505-9_9
[57] Srikanth Sundaresan, Nazanin Magharei, Nick Feamster, and Renata Teixeira. 2013. *Characterizing and Mitigating Web Performance Bottlenecks in Broadband Access Networks.* Technical Report GT-CS-13-02. Georgia Tech. http://hdl.handle.net/1853/46161
[58] Srikanth Sundaresan, Nazanin Magharei, Nick Feamster, Renata Teixeira, and Sam Crawford. 2013. Web Performance Bottlenecks in Broadband Access Networks. In *Proceedings of the ACM SIGMETRICS/International Conference on Measurement and Modeling of Computer Systems.* Association for Computing Machinery, Pittsburgh, PA, USA, 383–384. https://doi.org/10.1145/2465529.2465745
[59] Charles Thomas, Joel Sommers, Paul Barford, Dongchan Kim, Ananya Das, Roberto Segebre, and Mark Crovella. 2012. A Passive Measurement System for Network Testbeds. *Lecture Notes of the Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering* 44 (2012). https://doi.org/44.10.1007/978-3-642-35576-9_14
[60] U.S. Congress. 1996. Telecommunications Act of 1996: Advanced Telecommunications Incentives (47 USC 1302). Retrieved January 11, 2022 from https://www.govinfo.gov/content/pkg/USCODE-2019-title47/html/USCODE-2019-title47-chap12-sec1302.htm
[61] U.S. Congress. 2008. Broadband Data Improvement Act. Retrieved January 11, 2022 from https://www.govinfo.gov/content/pkg/PLAW-110publ385/pdf/PLAW-110publ385.pdf
[62] U.S. Congress. 2010. America COMPETES Reauthorization Act of 2010. https://www.congress.gov/bill/111th-congress/house-bill/5116
[63] Wireline Competition Bureau. 2020. *FCC Launches $20 Billion Rural Digital Opportunity Fund.* Federal Communications Commission. Retrieved January 10, 2022 from https://www.fcc.gov/document/fcc-launches-20-billion-rural-digital-opportunity-fund-0