

IMR: MT: RABBITS: A measurement toolkit for Reproducible Assessment of Broad-Band Internet Topology and Speed (PI: Ka Pui Mok, CAIDA/UC San Diego)

Overview

Web-based network throughput tests, also known as *speed tests*, are the most common (often the only) way that typical end users can measure their own broadband performance. As such, they have become important not just for end users, but also for regulators and researchers seeking to validate claims about broadband performance. Speed test services report basic metrics, including latency and average download/upload throughput of a bulk data transfer between users and one or more speed test servers. But the lack of standard approaches to test server deployment, implementation, configuration and even metrics measured might result in misalignment with users' measurement goals. Uneven server deployment and suboptimal server selection can distort measurement results. And while all speed test platforms use the same brute force approach to bandwidth estimation, they choose their own test parameters and other proprietary implementation details, making it impossible to effectively compare results across platforms.

We propose to develop a new measurement toolkit for Reproducible Assessment of BroadBand Internet Topology and Speed (RABBITS) enabling (1) comprehensive speed test infrastructure discovery and characterization, and (2) consistent test parameters across platforms. The goal of our RABBITS toolkit is to overcome the obstacles that have prevented these widely deployed global measurement infrastructures from supporting either rigorous scientific research or public policy needs for consistent and reusable tests of broadband performance and service availability.

The RABBITS toolkit consists of two modules: RABBITS-Mapper and RABBITS-Perf. RABBITS-Mapper discovers the locations of test servers hosted by multiple platforms. We will periodically run RABBITS-Mapper to collect geographical and topological information of speed test servers. Furthermore, we will design techniques to strategically deploy measurement vantage points (VPs) in CAIDA's Ark, RIPE Atlas, and the EdgeNet platform to characterize the paths and latency from the edge (users) to proximate speed test servers. We will leverage recent advances in mobile phone measurement methods to extend our coverage to cellular networks.

The RABBITS-Perf module will enable researchers to control parameters in speed tests, including the number and sizes of concurrent measurement flows, and selection of test servers. We will also design a standardized format to represent measurement data, facilitating reproducibility. These capabilities will significantly increase the flexibility of measuring Internet performance using various speed test infrastructures to achieve different measurement objectives.

Keywords: Speed test infrastructure, broadband measurement, Internet topology

Intellectual Merit

Consistent with the NSF's IMR Track 2's goals of *creating and deploying new tools to collect Internet measurement data*, this project will develop tools and share datasets that facilitate the use of speed test measurement infrastructure, an essential but under-studied Internet measurement tool for ISPs, regulators, and consumers. The outcome of this project will enable researchers to use speed tests more effectively for scientific research, and provide regulators with the tools they need to evaluate broadband deployment and performance.

Broader Impacts

Our toolkit will provide a lens to regulators and policy makers (e.g., FCC and NIST) to understand user-perceived network performance, identify underserved communities, and validate speeds advertised by ISPs. Our analysis will also help inform future speed test implementation and test server deployment. We plan to release data and tools resulting from this project, and provide accessible educational materials. We will incorporate our results into a CSE class on Internet Data Science for Cybersecurity. We will mentor students to use measurement data to comprehend digital divides within their communities.

1 Introduction

Accelerated adoption and use of online applications during the COVID-19 pandemic, such as video streaming and conferencing, have amplified concerns regarding how residential and mobile broadband performance affects the user’s quality of experience (QoE). Web-based network throughput tests, also known as *speed tests*, are the most common (often the only) way that typical end users can measure their own broadband performance. They have become important not just for end users, but also for regulators and researchers seeking to study and validate claims about broadband performance [1–12]. There are a variety of speed test providers today, including speed test-only services [13, 14], ISPs [15] and content providers [16, 17]. These services report basic metrics, including latency and average download/upload throughput of a bulk data transfer between users and one or more speed test servers [18, 19]. However, lack of standard approaches to test server deployment, implementation, configuration, and even metrics measured may result in misalignment with users’ measurement goals. Uneven server deployment and suboptimal server selection can distort measurement results. And while all speed test platforms use the same brute force approach to bandwidth estimation – flood the network to detect maximum TCP goodput – they choose their own test parameters and other proprietary implementation details, making it impossible to reliably compare results across platforms. For example, we have seen some platforms use a single TCP flow to measure throughput, and another use 18 concurrent TCP flows to measure the same metric.

We propose to tackle this challenge via an IMR Track 2 project. Specifically, we will create and deploy a measurement toolkit for *Reproducible Assessment of BroadBand Internet Topology and Speed* (RABBITS). The outcome will be a suite of new measurement datasets and tools clearing the two most significant barriers in speed measurement: (1) lack of comprehensive topology and performance information to facilitate the use of different speed test infrastructure, (2) inconsistent test parameters across platforms. Overcome these obstacles will transform the ability of these widely deployed global measurement infrastructures to support enforcing broadband consumer protection efforts as well as rigorous scientific research on performance and service availability.

The RABBITS toolkit consists of two modules: RABBITS-Mapper discovers the locations of test servers hosted by multiple platforms and selects vantage points (VPs) in CAIDA’s Ark [20], RIPE Atlas [21], and the EdgeNet [22] platform to measure the network paths and latency from the edge (end-users) to proximate speed test servers. To extend our coverage to cellular (LTE and 5G) networks, we will leverage our recent advances in mobile measurement methods by shipping running phones in packages. Our deployment will generate longitudinal datasets that will aid in understanding networks and regions that are underserved by speed test measurement infrastructure.

The second module, RABBITS-Perf, will facilitate execution of comparable speed tests across platforms. We will develop tools that allow experimenters to adjust test parameters, including the number and sizes of concurrent measurement flows, and selection of test servers. We will also design a standardized format to summarize measurement data while capturing relevant details. These capabilities will not only render speed test infrastructure more conducive to scientific use, but also promote the *reproducibility* of Internet performance measurements and analyses.

This research *creates and deploys new tools to collect Internet measurement data* that contributes scientific methods and curated datasets to *keep up with the importance and proliferation of wireless and fixed broadband access networks*, directly responsive to the NSF Internet Measurement Research Track 2’s program goals. We will design and implement tools to aid researchers to use the under-studied speed test measurement infrastructure for scientific research. The datasets that we will publish can assist ISPs, regulators and consumers to identify underserved communities.

2 Background and related work

Different organizations have built speed test platforms to enable users to gauge their Internet performance. The basic methodology for all platforms involves users executing a script on a speed test website from their web browser or dedicated mobile apps. The script establishes HTTP(S)/TCP connections from the user’s browser to test servers and conducts bulk data transfers over those connections to measure the available bandwidth of the network (at least the segment between the user and the server). Ookla, the most pervasively deployed speed test platform, has performed over 40 billion tests since its 2006 launch [13]. In this section, we discuss the state-of-the-art of two major components of speed test platforms—test servers (§2.1) and test methodologies (§2.2).

2.1 Overview of speed test infrastructure

Speed test platforms share a common goal of measuring performance from end users by deploying many geographical distributed test servers. But different platforms achieve different levels of server diversity and proximity to users, in terms of network topology and geography. The choice of test servers controls the network paths that measurement traffic traverses (Fig. 1) and thus significantly affects the measurement results [23].

ISPs deploy speed test servers within their own networks (*on-net*) to support their customers in diagnosing their last-mile performance [24–26], and verify their access bandwidth is as advertised. But ISPs do not guarantee a certain level of bandwidth to any given piece of content on the Internet, and thus do not architect speed test infrastructure to measure paths to remote sources of content. CDNs and content providers sometimes build their own test platforms to enable their subscribers to test performance of CDN services [27,28], although some CDNs (e.g., Netflix) offer speed tests from their server appliances inside the ISPs [27]. Cloud-based test platforms (e.g., [29]) enable consumers to monitor and benchmark their cloud services. In 2010, the U.S. FCC launched the Measure Broadband America (MBA) program, with the goal of establishing whether consumers were receiving advertised bandwidth, but this program does not support a general purpose speed test platform [1]. The FCC outsourced MBA’s measurement operations to a commercial company (SamKnows [30]), and the source code is proprietary. None of these platforms aim to support scientific research.¹




Figure 1: Placement of test servers (e.g., on-net, transit ISPs, CDNs) affects the measured network paths, and sometimes may not align with a user’s measurement goals. Lack of access to nearby servers (e.g., for User R) may lead to confusing test results. Measurement traffic sent to a remote on-net server in a different region, or to off-net test servers, will impede the accuracy of last-mile measurements. In contrast, user M with an on-net server in the same region will have low latency and interference with last-mile speed tests. High test server density will also enable users to cross-validate performance measurements using servers in different types of networks.

Table 1 summarizes seven popular speed test platforms based on measurements we performed in

¹In 2009 Google, launched its Measurement Lab platform of distributed servers connected with high bandwidth to support performance measurements. Scientific use of this infrastructure has been modest for reasons discussed in [7].

February 2022. Except for M-Lab [31], no platform provides detailed information (e.g., IP address and location) about their test servers. Each platform has a different scale in terms of number of servers, geographic distribution, networks, and use of infrastructure. In our survey, Ookla [32] had the largest number of servers, spanning 8,901 ASes, including educational networks and residential ISPs; 2,073 (10.9%) of them had IPv6 addresses. M-Lab [14] deployed servers to transit providers and Google Cloud Platform (GCP), and were IPv6-capable. Comcast’s Xfinity speed test servers were all in Comcast’s network (AS7922). CloudHarmony [29] set up virtual speed test servers in different cloud regions of three public cloud platforms—Amazon AWS, Microsoft Azure, and GCP.

Netflix hosts fast.com [16] on its Open Connect Appliance (OCA) to help customers diagnose video streaming performance problems. In addition to connecting to OCA caches through peering or transit networks, ISPs can install embedded appliances [33] within their networks to reduce the traffic load on the interdomain links. Fast.com uses DNS names to load balance clients to different caches, mainly based on the ISP and geolocation of the clients [34]. Speedof.me and Cloudflare provide speed test services using Edgio (formerly named as Edgecast) and Cloudflare’s CDN, respectively. Both CDNs use IP anycast to let BGP route client content requests to the topologically closest content caches.

Table 1: Summary of speed test platforms RABBITS will support in this project.

Platform	Servers		# of countries	# of ASes	Type	Public Server list	
	IPv4	IPv6					
Ookla [32]	18,997 [‡]	2,073	216	8,901	Unicast	No	
M-Lab [14]	181	181	36	43	Unicast	Yes	
Comcast [15]	139 [‡]	139	1	1	Unicast	No	
CloudHarmony [29]	-AWS	23	0	17	1	Unicast Cloud	No
	-Azure	36	0	17	1	Unicast Cloud	No
	-GCP	24	0	17	1	Unicast Cloud	No
Fast.com [16]	300 [‡]	Unknown	28	6	Unicast CDN	No	
Speedof.me [35]	127 [36]	127	32	1	Anycast CDN	No	
Cloudflare [17]	Unknown	Unknown	>100 [37]	1	Anycast CDN	No	
FCC [1]	133 [38]	133	1	15	Unicast	No	

Note: [‡] Based on our preliminary results.

There is no published analysis of speed test platform deployments, e.g., the geographical/network distance from test servers to end-users. Many studies have evaluated latency dynamics between CDN caches and end-users [39–45], as latency is a key performance metric to consider in the design of these infrastructures [46]. These studies employed different methods to measure latency from end-users to CDN caches, including embedding scripts on web pages [40], conducting traceroutes from RIPE Atlas and PlanetLab VPs [41, 43–45], and issuing HTTP requests from self-hosted VPs [42]. In a 2017 study, end users in Africa and Asia experienced higher latency to anycast DNS-root servers than other continents [41]. A 2020 study found that cloud platforms, which usually have sparser geographical footprints than Akamai CDNs, could only reach 62% of RIPE Atlas VPs within 20ms [45].

2.2 Speed test measurement methodologies

Although all speed test platforms use the same basic measurement approach, their implementations vary in terms of test server selection, number and size of concurrent measurement flows, and algorithms for computation of measurement results.

Server selection algorithm. Ookla selects 10 servers nearest to the user based on proprietary IP geolocation databases, and conducts ping tests to find the one in this set with the lowest

latency from the users [47]. Test server selection in Comcast and M-Lab solely is based on IP geolocation, such as Google AppEngine’s IP geolocation [48]. CDN-based platforms share the same client catchment (assigning clients to their nearest anycast nodes) algorithms as their host CDNs. Measurement traffic toward different test servers could traverse different network paths.

Measurement flows. The number and the size of measurement flows varies widely among platforms [18,19]. The number of concurrent TCP flows varies from 1 (speedof.me) to 18 (Comcast). In contrast to Ookla and Comcast, fast.com uses many small consecutive HTTP transactions rather than performing a few large-size transfers [19]. These characteristics can affect a test’s ability to saturate the bottlenecks, and resource utilization in the client, making different test incomparable.

Methods for computing speed estimates. Speed test platforms use proprietary methods to process measurement results and estimate network throughput. We can analyze the method’s used by M-Lab and Comcast’s speed tests, because they publish the associated source code. Ookla [49] and fast.com [50] provide only high-level descriptions. A common step in the methods is to remove samples collected during TCP slow start, removal of which mitigates underestimation of throughput. Recent work proposed FastBTS (Bandwidth Testing Service) [51] applies a statistical sampling framework to filter noise in throughput samples to reduce measurement overhead and test duration. However, FastBTS requires software installation on test servers, and thus cannot easily apply to existing widely deployed test infrastructure.

Other studies have investigated the challenges of building and using speed test platforms [18, 52–54]. Bauer *et al.* [18] studied five popular speed tests and found significant differences in measurement methods, including aforementioned factors, which led to discrepancies in results. High-speed residential broadband services (up to 1Gbps) impose new measurement challenges. Test platforms that rely on off-net servers may underestimate the last-mile capacity of high-speed access links [52]. Other methodological limitations in speed tests [54]. include interference on cross-traffic within home networks [54], and sparse server placement and inaccurate IP geolocation results [23]. To our knowledge, no analysis has either systemically evaluated the current architectures and methods of platforms or quantified the potential impacts to measurement results caused by server placement. Our research aims to provide tools and datasets to bridge this gap and overcome some limitations highlighted in [54] via scientific Internet measurements.

Recent developments put us in a unique position to address these research gaps. Last year as part of the NSF’s COVID-related Internet measurement research projects [55], PI Mok developed *WebTestKit* [19], a unified and configurable framework for facilitating automatic speed test execution and cross-layer analysis of test results. *WebTestKit* controls a headless Chromium browser to execute speed tests, and captures packets and metadata from the browsers. We applied this tool to analyze sources of inaccuracy in latency measurements [19] and study network performance of public cloud platforms [56, 57]. This experience trying using speed test platforms for scientific research generated more questions than answers. For example, different test parameters prevented us from directly comparing results collected using different test platforms. Furthermore, our attempt to conduct and handle the huge volume of raw data gathered in longitudinal experiments demonstrated the need for a compact and standardized format to facilitate sharing of such measurement data in support of reproducible and extensible research. We structure our research agenda to overcome both obstacles.

3 Broader Impacts

Speed tests are an essential tool to understand Internet performance. This research will provide scientific understanding of deployment and methodology of speed test platforms and enable

researchers to conduct reproducible experiments across platforms. The data and tools will also benefit consumers, industry, and regulators.

A number of recent policies and programs, including FCC Broadband Consumer Labels [58], the new FCC broadband maps [59], and NIST's Broadband Equity, Access, and Deployment (BEAD) [12], proposed to use network performance (throughput and latency) to regulate the broadband marketplace and distinguish unserved/under-served communities in broadband coverage.

Our tools and data will provide a lens to regulators and policy makers trying to understand actual network performance that end users perceive, conduct robust and defensible measurements to support policy decisions and examine the performance advertised by ISPs. ISPs will also benefit from our analysis which will inform considerations of future test server deployment.

We will broadly disseminate results of this project to the measurement and operational communities via publications at academic conferences, journals, and operational and policy meetings. We will leverage CAIDA's MSRI GMI workshops [60] and solicit community feedback on our design and deployment of the toolkit. Apart from releasing data and tools resulting from this project, we will provide accessible educational materials (e.g., slide sets, animations and videos) via CAIDA's resource catalog [61].

We plan to engage undergraduate students, particularly from underrepresented groups, through the Department of Computer Science and Engineering (CSE) and HalÁšcÁšoÄŸlu Data Science Institute (HDSI) at UCSD. We will incorporate our experience with Internet measurements into our recently developed CSE course on Internet Data Science for Cybersecurity [62].

References

- [1] FCC, “Measuring broadband america.” <https://www.fcc.gov/general/measuring-broadband-america>, 2021.
- [2] New York State Broadband Program Office, “Speed test.” <https://nysbroadband.ny.gov/speed-test>.
- [3] CIRA, “Internet performance test.” <https://www.cira.ca/improving-canadas-internet/initiatives/internet-performance-test>, 2021.
- [4] D. Riddlesden and A. D. Singleton, “Broadband speed equity: A new digital divide?,” *Applied Geography*, vol. 52, pp. 25–33, 2014.
- [5] F. Bronzino, N. Feamster, S. Liu, J. Saxon, and P. Schmit, “Mapping the digital divide: Before, during, and after COVID-19,” in *Proc. TPRC*, 2021.
- [6] I. Canadi, P. Barford, and J. Sommers, “Revisiting broadband performance,” in *Proc. ACM IMC*, 2012.
- [7] D. Clark and S. Wedeman, “Measurement, meaning and purpose: Exploring the M-lab NDT dataset,” in *Proc. TPRC*, 2021.
- [8] G. S. Ford, “Form 477, speed-tests, and the American broadband user’s experience,” *Phoenix center policy perspective*, vol. 20, no. 3, 2021. <http://www.phoenix-center.org/perspectives/Perspective21-03Final.pdf>.
- [9] F. Michel, M. Trevisan, D. Giordano, and O. Bonaventure, “A first look at starlink performance,” in *Proc. ACM IMC*, 2022.
- [10] A. Jain, D. Patra, P. Xu, J. Sherry, and P. Gill, “The ukrainian internet under attack,” in *Proc. ACM IMC*, 2022.
- [11] U. Paul, J. Liu, M. Gu, A. Gupta, and E. Belding, “The importance of contextualization of crowdsourced active speed test measurements,” in *Proc. ACM IMC*, 2022.
- [12] National Telecommunications and Information Administration, “Broadband equity, access, and deployment (BEAD) program.” <https://broadbandusa.ntia.doc.gov/taxonomy/term/158/broadband-equity-access-and-deployment-bead-program>, 2022.
- [13] Ookla, “About ookla.” <http://www.speedtest.net/en/about>.
- [14] M-Lab, “NDT (network diagnostic tool).” <https://www.measurementlab.net/tests/ndt/>.
- [15] Comcast, “Xfinity speed test.” <http://speedtest.xfinity.com>.
- [16] Fast.com, “Internet speed test.” <https://fast.com>.
- [17] CloudFlare, “Cloudflare speed test.” <https://speed.cloudflare.com>.
- [18] S. Bauer, D. Clark, and W. Lehr, “Understanding broadband speed measurements,” in *Proc. TPRC*, 2010.
- [19] R. Yang, R. Mok, S. Wu, X. Luo, H. Zou, and W. Li, “Design and implementation of web-based speed test analysis tool kit,” in *Proc. PAM*, 2022.
- [20] CAIDA, “Archipelago (Ark) measurement infrastructure.” <https://www.caida.org/projects/ark/>.
- [21] RIPE NCC, “Ripe atlas.” <https://atlas.ripe.net>.
- [22] “Edgenet.” <https://www.edge-net.org>.
- [23] K. MacMillan, T. Mangla, J. S. nd Nicole P Marwella, and N. Feamster, “A comparative analysis of Ookla speedtest and measurement labs network diagnostic test (NDT7),” in *Proc. ACM SIGMETRICS*, 2023.
- [24] AT&T, “Check your internet speed.” <https://www.att.com/support/article/u-verse-high-speed-internet/KM1010076/>.
- [25] Spectrum, “How to conduct a spectrum speed test.” <https://www.spectrum.net/support/internet/spectrum-speed-test>.

- [26] Comcast, “Check your computer’s performance.” <https://www.xfinity.com/support/articles/determine-slow-computer-speeds>.
- [27] Netflix, “Netflix help center: Internet connection speed recommendations.” <https://help.netflix.com/en/node/306>.
- [28] Hulu, “Hulu help center: Test your internet connection.” https://help.hulu.com/s/article/speed-test?language=en_US.
- [29] CloudProbe, “Cloudharmony.”
- [30] “SamKnows.” <https://www.samknows.com>.
- [31] M-Lab, “M-Lab sites.” <http://mlab-ns.appspot.com/>.
- [32] Ookla, “Speedtest.” <http://www.speedtest.net>.
- [33] Netflix, “Deploying embedded appliances.” <https://openconnect.netflix.com/en/deployment-guide/>, 2021.
- [34] T. V. Doan, V. Bajpai, and S. Crawford, “A longitudinal view of Netflix: Content delivery over IPv6 and content cache deployments,” in *Proc. IEEE INFOCOM*, 2020.
- [35] “Speedof.me.” <https://speedof.me>.
- [36] speedof.me, “How it works.” <https://speedof.me/howitworks.html>, 2021.
- [37] Cloudflare, “The Cloudflare global network.” <https://www.cloudflare.com/network/>, 2021.
- [38] FCC’s Office of Engineering and Technology, “Measuring Broadband America: Technical appendix to the eleventh MBA report.” <https://data.fcc.gov/download/measuring-broadband-america/2021/Technical-Appendix-fixed-2021.pdf>, 2021.
- [39] Y. Zhu, B. Helsley, J. Rexford, A. Siganporia, and S. Srinivasan, “LatLong: Diagnosing wide-area latency changes for CDNs,” *IEEE Transactions on Network and Service Management*, vol. 9, pp. 333–345, sep 2012.
- [40] M. Calder, A. Flavel, E. Katz-Bassett, R. Mahajan, and J. Padhye, “Analyzing the performance of an anycast CDN,” in *Proc. ACM IMC*, 2015.
- [41] R. de Oliveira Schmidt, J. Heidemann, and J. H. Kuipers, “Anycast latency: How many sites are enough?,” in *Proc. PAM*, 2017.
- [42] J. Xue, D. Choffnes, and J. Wang, “CDNs meet CN an empirical study of CDN deployments in china,” *IEEE Access*, vol. 5, pp. 5292–5305, 2017.
- [43] D. Cicalese and D. Rossi, “A longitudinal study of IP anycast,” *ACM CCR*, 2018.
- [44] S. McQuistin, S. P. Uppu, and M. Flores, “Taming anycast in the wild Internet,” in *Proc. ACM IMC*, 2019.
- [45] B. Charyyev, E. Arslan, and M. Gunes, “Latency comparison of cloud datacenters and edge servers,” in *Proc. IEEE Globalcom*, 2020.
- [46] H. Ballani, P. Francis, and S. Ratnasamy, “A measurement-based deployment proposal for IP anycast,” in *Proc. ACM IMC*, 2006.
- [47] C. Overturf, “How does speedtest select a server?.” <https://help.speedtest.net/hc/en-us/articles/360038679834-How-does-Speedtest-select-a-server->.
- [48] M-Lab, “M-Lab naming service.” https://github.com/m-lab/mlab-ns/blob/main/DESIGN_DOC.md.
- [49] T. Ookla, “What is the test flow and methodology for the speedtest?.” <https://support.ookla.com/hc/en-us/articles/234575828-What-is-the-test-flow-and-methodology-for-the-Speedtest->, Nov. 2016. Accessed on Jun 26, 2019.
- [50] S. Fedorov and E. Livengood, “Building fast.com.” <https://netflixtechblog.com/building-fast-com-4857fe0f8adb>, 2016.
- [51] X. Yang, X. Wang, Z. Li, Y. Liu, F. Qian, L. Gong, R. Miao, and T. Xu, “Fast and light bandwidth testing for internet users,” in *Proc. USENIX NSDI*, 2021.

- [52] S. Bauer, W. Lehr, and S. Hung, “Gigabit broadband, interconnection propositions, and the challenge of managing expectations,” in *Proc. TPRC 43*, 2015.
- [53] S. Bauer, W. Lehr, and M. Mou, “Improving the measurement and analysis of gigabit broadband networks,” tech. rep., Massachusetts Institute of Technology, 2016.
- [54] N. Feamster and J. Livingood, “Measuring internet speed: Current challenges and future recommendations,” *Communications of the ACM*, vol. 63, no. 12, pp. 72–80, 2020.
- [55] R. Mok and kc claffy, “Measuring critical infrastructure for coronavirus-related congestion.” https://www.caida.org/funding/rapid-speedtest/rapid-speedtest_proposal.xml, 2020.
- [56] R. Mok and kc claffy, “Measuring the impact of COVID-19 on cloud network performance,” in *COVID-19 Network Impacts Workshop*, Internet Architecture Board, 2020.
- [57] R. Mok, H. Zou, R. Yang, T. Koch, E. Katz-Bassett, and kc claffy, “Measuring the network performance of Google cloud platform,” in *Proc. ACM IMC*, 2021.
- [58] F. C. Commision, “Broadband consumer labels.” <https://www.fcc.gov/broadbandlabels>, Jan. 2023.
- [59] F. C. Commission, “FCC national broadband map.” <https://broadbandmap.fcc.gov/home>, Nov. 2022.
- [60] CAIDA, “CAIDA Global Measurmeent Infrastructure (GMI) Meetings.” <https://www.caida.org/workshops/>, 2023.
- [61] CAIDA, “CAIDA resource catalog.” <https://catalog.caida.org>.
- [62] UCSD CSE, “Internet data science for cybersecurity.” <https://cseweb.ucsd.edu/classes/wi23/cse291-e/syllabus.html>, 2023.