

Benchmarks or Equity? A New Approach to Measuring Internet Performance

Ranya Sharma, Tarun Mangla, James Saxon, Marc Richardson, Nick Feamster, Nicole P. Marwell
University of Chicago

Abstract

A longstanding approach to measuring Internet performance is to directly compare throughput against pre-defined benchmarks (e.g., 25 megabits per second downstream, 3 megabits per second upstream). In this paper, we advocate, develop, and demonstrate a different approach: rather than focusing on whether speeds meet a particular threshold, we develop techniques to determine whether a variety of Internet performance metrics (including throughput, latency, and loss rate) are *comparable* across geographies. We define these metrics and apply them across a longitudinal dataset of Internet performance measurements comprising approximately 30 neighborhoods across the City of Chicago. The metrics we define show some geographical disparities, indicating that such comparative metrics may be promising for studying questions of equitable Internet access across neighborhoods.

1 Introduction

The Federal Communications Commission (FCC), advocacy groups, scholars, and others interested in Internet performance have had wide-ranging discussions over the years on two related (but distinct) topics: (1) benchmarking Internet performance standards (“Am I getting what I’m paying for?”); and (2) equity (“digital divide” challenges, Internet access haves and have-nots). Various organizations, from regulators to public interest groups, have developed tools, methods, and datasets to try to answer these questions. Historically, these datasets and methods have been used to study questions of benchmarking, although there is also increasing talk of issues related to digital equity, including among regulators (e.g., the “rural” or “urban” digital divide, the “homework gap”). Unfortunately, policy arguments around benchmarking have been based on metrics such as throughput that are poor proxies for user experience, using samples that do not adequately represent the population. These samples, while convenient, ultimately have little utility for benchmarking—and they are even less suitable for asking questions about equity, where marginalized popula-

tions are often under-represented or entirely absent in existing datasets. In addition, there are many outstanding questions about what the distribution of physical Internet infrastructure (fiber, coax, etc.) actually looks like, and whether there might be systematic differences in infrastructure provisioning across geographic areas, especially when considering data at a more granular level, such as within cities. Questions about appropriate sampling for assessing Internet and application performance thus must take this issue into account as well.

In this paper, we consider the types of sampling methods and metrics that could be appropriate for benchmarking Internet service provider (ISP) (and application) performance across the geography of Internet infrastructure. We find that no existing sample—including the FCC’s Measuring Broadband America program—is appropriate. Further, we find that constructing such a sample with contemporary datasets is extremely challenging in practice. On the other hand, constructing a sample to answer questions concerning equity of Internet performance—specifically, comparing performance properties across pairs of small geographic areas, such as urban neighborhoods—is much more tractable with a sample that can be obtained in practice, through targeted recruiting. In addition to presenting the sampling rationale to support these findings, we also show preliminary data from a sample across Chicago that is specifically designed to answer questions of equity in Internet performance.

In a preliminary analysis of pairwise comparisons in Chicago, we measure how Internet performance in a specific high-income neighborhood compares with Internet performance in a specific low-income neighborhood. Our sampling process allows us to collect appropriate data for these pairwise comparisons, suggesting a promising approach for measuring Internet equity more broadly. This approach could also provide researchers and policymakers an opportunity to offer citizens a data-driven approach to address their core concerns about Internet equity. More generally, this initial study and findings demonstrate the possibilities of studying more focused questions around equity and “digital redlining”

with carefully constructed samples—entirely separate from regulatory questions concerning benchmarks.

2 Motivation

The idea of “digital redlining” has been in the public discourse for a number of years, but has been used rather loosely, without clear distinctions between possible definitions. Definitional slippage also is evident between the related concepts of “digital divide” and “digital redlining.” Both terms aim to call attention to the observed reality of broadband disparities between places (e.g., urban/rural) and groups (e.g., high-income/low-income, White/Black, etc.). Here, we offer a distinction between these two terms, as well as an empirical basis for considering each.

One way to understand broadband disparities is by comparing rates of household broadband connectivity across population groups or geographies of interest. Doing so reveals that these rates vary substantially. For populations, there are significant disparities between groups in connectivity rates. As just two examples, in the U.S., approximately 80 percent of Whites, 71 percent of Blacks, and 65 percent of Hispanics have a broadband connection at home; in addition, approximately 64 percent of people over 65 have a broadband connection at home, compared to 86 percent of people ages 30 to 49, and 79 percent of those ages 50 to 64. When considering geography, there are also clear disparities, including between rural and urban areas, and between neighborhoods within cities. For example, broadband connectivity rates in Chicago neighborhoods range from a low of 57 percent in the community area of Englewood to a high of 91 percent in the community area of Lakeview. The term “digital divide” is useful for pointing to these connectivity rate disparities, which have multiple causes, including availability, affordability, and other barriers to adoption.

In contrast, the idea of “digital redlining” suggests a different process at work. Here, the question is whether ISPs engage in practices that prevent certain geographic areas—usually those inhabited by low-income people and/or by Black people or other people of color—from connecting to quality broadband service at competitive prices. The concept of digital redlining thus is analogous to the well-known practice of real estate redlining.

Beginning in the 1930s, the federal Home Owners’ Loan Corporation (HOLC) produced reference maps that classified neighborhoods within the nation’s cities as having different levels of risk for home loans. Neighborhoods that were deemed the highest risk were marked by red lines around their perimeters, and banks declined to provide mortgages or loans for home improvements in those areas. Those neigh-

borhoods also were predominantly Black, as practices of residential racial segregation meant Black families were concentrated in a small number of neighborhoods. As has been documented by historians, segregation and redlining led directly to declines in the quality of housing stock in Black neighborhoods, as well as significantly lower rates of homeownership by Blacks [15]. This was the case even though it was far from clear that making loans in Black neighborhoods in fact represented higher risk to banks.

Applying the idea of real estate redlining to the realm of broadband, “digital redlining” concerns the quality of broadband infrastructure that ISPs have built in different places. Here, the question is not about whether broadband infrastructure exists at all, but rather about the quality of that infrastructure and what that means for broadband customers and their quality of experience. For example, a 2020 study from the National Digital Inclusion Alliance showed that AT&T has prioritized areas with higher median incomes for upgrades to its broadband infrastructure (i.e., installing fiber to the home) [2]. Studies examining ISP infrastructure practices that might be seen as digital redlining, however, are rare, often due to lack of data availability.

In this paper, we draw on new data collected from households to assess whether there are significant differences in internet performance by geography (i.e., between neighborhoods within a single city). We treat our internet performance measures as a proxy for internet infrastructure quality. This allows us to make an initial assessment of whether digital redlining might be occurring.

3 Related Work

Existing infrastructure for measuring broadband access networks has made remarkable progress, but existing platforms only allow for a limited amount of testing and in some cases are limited in footprint.

3.1 Existing Platforms

In this section, we provide a brief overview of existing measurement platforms, including both client-based tests and router-based (or device-based) tests.

3.1.1 Client-Based Measurement Platforms

Ookla Speedtest and Measurement Labs Network Diagnostic Test (NDT). Ookla Speedtest [10] and Measurement Lab’s Network Diagnostic Test (NDT) [8] are two examples of client-based network measurement tools that collect various types of performance measurements, including throughput and latency measurements. Client-based

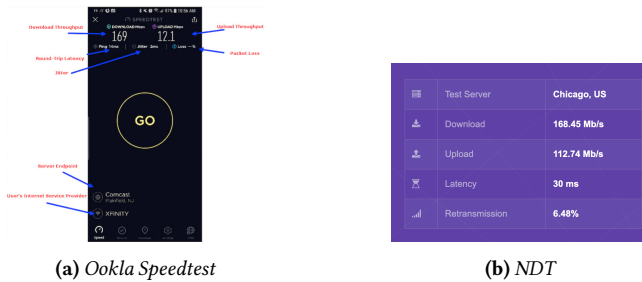


Figure 1: Example metrics from various client-based speed tests.

platforms collect end-to-end Internet performance metrics—specifically throughput (i.e., “speed”), latency, and packet loss. Figure 1a shows an example of an Ookla speed test from a mobile phone on a home WiFi network. It shows the results of a “native” speed test from the Ookla Android speed test application [10], a canonical Internet speed test. This native application reports the user’s ISP, the location of the test server destination, and metrics such as throughput, latency, jitter, and packet loss. Other client-based speed tests perform similarly. Figure 1b shows example output from running NDT.

Unfortunately, as we will describe below, although these client-based measurements offer the benefit of relatively large datasets in terms of absolute numbers of tests, these platforms face significant limitations, including: (1) the lack of a longitudinal sample; (2) the lack of a deployment footprint that represents any characteristic of interest; (3) the inability to deploy custom measurements. The present research aims to address many of these shortcomings.

3.1.2 Router & Device-Based Platforms

RIPE Atlas. Some platforms, notably RIPE Atlas, provide on-demand measurements, public data, and “real” network vantage points [12]. RIPE Atlas is a globally distributed testbed of thousands of devices that allow researchers to make fairly simple network measurements, such as ping and traceroute. The platform is one of the most successful network experimentation platforms, but the network tests it supports do not include capability for customization, and the sample—based on where RIPE Atlas probes are installed—is a convenience-based sample. As a result, the RIPE Atlas platform does not provide the opportunity to perform the types of controlled experiments that we aim to provide in the present research. That being said, RIPE Atlas is arguably the most successful Internet measurement platform to date, as it allows researchers to conduct custom measurement experiments, collect the data, and curate that data for other future researchers.

FCC’s Measuring Broadband America. Other existing deployments include the FCC’s Measuring Broadband America (MBA) program [6]. Each of these faces limitations: while the MBA program has the MBA-assisted research program (MARS) [7], the process for deploying a test on the platform is cumbersome (it can take almost a year to gather a limited amount of data), the platform itself is extremely limited (outdated home routers) which both slows development and limits the tests that can be deployed, and only one test can be deployed at a time. Due to the sampling approach taken by the FCC to construct the MBA participants (i.e., stratified by ISP), this platform suffers significantly from under-sampling of certain parts of the access network, especially in under-served regions. One need not perform complicated statistical analysis to reveal obvious shortcomings with the current FCC sample frame: The FCC MBA panel includes zero users in Chicago, let alone users in marginalized communities.

Broadband Internet Service Benchmark (BISmark).

Our previous work on router-based test suites partially addresses the shortcomings of other platforms. In particular, the BISmark (Broadband Internet Service Benchmark) platform [13, 14] is a suite of network performance tests including multiple types of throughput tests, latency, packet loss, and jitter to various destinations. It operates from a network endpoint and is commonly deployed on a single-board computer (e.g., Raspberry Pi, Odroid) which can then be installed on any network, including a user’s home network. The latest version of BISmark has been re-implemented from the ground up as **Netrics** [9], to include a wider variety of throughput tests, including Ookla’s Speedtest [10], iPerf [5], and Measurement Lab’s Network Diagnostic Test (NDT) [8]. In addition to the client software, BISmark includes a management suite that allows deployment of software and configuration updates to a distributed set of deployed measurement nodes.

4 Methods

In this section, we describe the measurement platform that we developed to conduct our measurements, as well as the sampling and deployment strategy that we used to collect the data.

4.1 Measurement Platform

The **Netrics** platform [9], which is used for the present research, has been underway for about one year. It currently provides a variety of open-source network measurement tests, with plans to enable the deployment of more tests. Netrics conducts measurements by distributing Rasp-

Community	# Devices	Coverage (%)
Focus Communities		
Logan Square	30	88.18
South Shore	26	78.96
Total (Focus)	56	
City-Wide Communities		
Albany Park	2	81.45
Ashburn	1	85.72
Avalon Park	1	70.02
Avondale	3	85.85
Belmont Cragin	1	76.69
Bridgeport	1	87.87
Dunning	1	78.71
East Side	1	84.31
Englewood	2	64.11
Hyde Park	1	88.40
Lake View	4	92.35
Lincoln Square	1	90.74
Loop	3	88.29
Lower West Side	1	85.85
Mount Greenwood	1	83.87
South Chicago	1	81.41
South Lawndale	1	74.89
Uptown	3	74.89
West Ridge	1	83.12
West Town	2	91.30
Total (City-Wide)	32	
Total (All)	88	

Table 1: Number of measurement devices per community area, and estimated coverage (percentages of households with broadband) per community area in Chicago.

berry Pi kits to participants in geographic areas of interest. Participants connect the Pis to routers or modems in their home networks; the Pis measure and archive metrics of interest. Participants were recruited in partnership with local community-based organizations and through social media channels; they were offered a monetary incentive of \$100 to install a device to an access point for one month and then offered an additional \$25 to continue hosting the device for another six months [11], an option accepted by a little less than 80% of participants.

We have devoted significant time to building awareness of the project across a wide range of stakeholders in the historically marginalized communities featured in the study. As communities that often have been excluded by major institutions, including government, banks, ISPs, and urban univer-

sities, the targeted communities harbor significant suspicion towards researchers. Building trust for the project through respected community representatives has been key to the success of our existing deployments.

Starting in 2021, 98 devices were deployed by those means in 31 distinct community areas [1] in Chicago to measure Internet speed, reliability, and equity in Internet performance across the city. The output of these measurements is a new dataset that combines survey data on demographics, Internet access equipment, and type and cost of Internet service with robust, granular measurements on Internet speed, reliability (including outages, latency, packet loss, etc.), local network bandwidth, and local network utilization (as indicated by the number of devices connected to a local network). Figure 2 shows the state of the deployment as of August 2022.

We have been collecting data from the deployment described above since June 2021 among our research team, and since October 2021 across a broader set of study participants across Chicago. Our pilot study and deployment has demonstrated that an observational instrument of this type is feasible, and blazed the trail for many of the deployment, logistical, and technical solutions that need to be put in place to deploy it.

Dashboard. Data is continually ingested from deployed measurement devices to an Influx database [4] with a Grafana front-end [3]. Figure 3 shows a sample of the data that is available from one device, as well as another dashboard that we constructed to facilitate exploratory comparisons across different participants in the pilot study.

4.2 Sampling and Deployment Strategy

Our sample consists of measurement devices deployed to 88 unique households across the city of Chicago. We chose two focus communities, one on the South Side of the city (South Shore), and the other on the North Side (Logan Square). This comparison is designed to capture effects of Chicago’s long history of racial residential segregation, where the South Side of the city concentrates its Black population, and the North Side concentrates its White population. The sociodemographics of our focus communities reflect those of the South and North Sides of the city, including a disparity in household broadband connectivity, as seen in Table 1. There are 26 devices deployed in South Shore, and 30 in Logan Square.

In addition to our two focus communities, we also deployed measurement devices more broadly across Chicago community areas (see Table 1). There are 32 devices deployed across 20 additional communities, with each community having between 1 and 4 devices. In our analysis, we group all




Figure 2: Current Netrics platform deployment in Chicago as of August 2022.



Figure 3: Real-time Netrics dashboard, showing a subset of the Internet performance metrics that are continuously collected from each participant.

of these devices into our "city-wide" category. This strategy was used to capture variation in Internet performance across the city, as a point of comparison with our focus neighborhoods.

Our analytic approach thus conducts pair-wise comparisons between three groups: our focus communities of South Shore and Logan Square, and our city-wide comparison group.

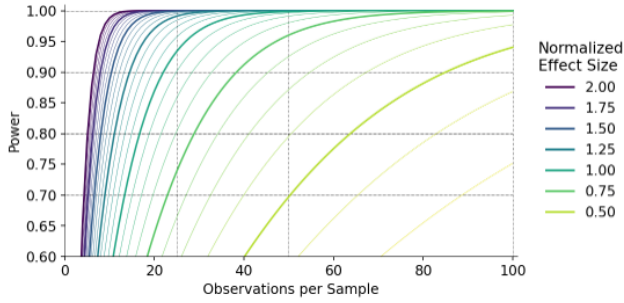




Figure 4: Sample vs. effect size.

When comparing performance statistics across population groups, it is naturally important to understand the number of observations needed for each variable that are required to claim a significant effect (i.e., one that would not be achieved by chance). Figure 4 shows the power one can achieve from a certain number of observations, given various effect sizes, assuming a t-test with two independent groups of samples, each with a certain number of observations. The figure shows that the larger the normalized effect size, the more observations one needs to make (x-axis) before significance can be claimed for a certain level of power (y-axis). Given 30 observations (i.e., households) in each population group, as we have assembled (Table 1), we can say that normalized effect sizes of more than 1.25 are significant to a power of near 1. Thus, when discussing the results in Section 5, we will keep a normalized effect factor of 1.25 as a general guideline for understanding the significance of the differences we see across neighborhoods.

Period of analysis. Due to fact that our measurements are continuous and ongoing—over more than a year—the data itself has various discontinuities. This is the result of many factors, such as when participants temporarily disconnect their devices, ISPs make unannounced speed tier upgrades, and so forth. To minimize the effects that these kinds of discontinuities could have on our analysis (e.g., averages over two distinct speed tiers, or over periods where data was missing or zero), we selected a one-month period that had minimal discontinuities from which we performed our analysis. After exploring the distribution of active devices across all months of the deployment, we chose the time period with the most continuously active devices, which was the month of May 2022. For this reason, the period of analysis in this paper reflects the time period of May 1–31, 2022.



(a) Ookla.



(b) NDT7.


Figure 5: Downstream throughput distributions across Chicago communities, measured with Ookla and NDT7 speed tests.

5 Results


In this section, we present the results of our measurements and analysis. We focus first on comparing both downstream and upstream throughput across neighborhoods (Section 5.1) before turning to latency (including latency under load) (Section 5.2) and loss rate or retransmission rate (Section 5.3).

5.1 Throughput

Metric: Normalized reduction. When comparing throughput across participants in this study, we face two



(a) Ookla.



(b) NDT7.

Figure 6: Upstream throughput distributions across Chicago communities, measured with Ookla and NDT7 speed tests.

challenges: (1) the *a priori* speed tier of one of our participants is generally unreliable data and can be very difficult to gather accurately; and (2) speeds can vary over time, for a variety of reasons. To account for these uncertainties and variabilities, we define a new, per-household variable, *normalized reduction* (in throughput), which is defined as $(p_{90} - p_{10})/p_{90}$, where p_i is the i th percentile measurement from that household. Intuitively, normalized reduction captures how a throughput measurement varies (i.e., degrades) over the course of longitudinal measurements. A normalized reduction of 0 suggests that the 10th and




Figure 7: Idle latency.

90th percentiles of throughput measurements from that household are equal—indicating no substantial degradation. A normalized reduction of 1 would suggest a p_{10} of 0—a particularly severe degradation at the 10th percentile. Thus, larger values of normalized reduction indicate more severe degradations—and overall a worse user experience. Each result, which shows the 10th, median and 90th percentiles per neighborhood, represents a percentile of *households* within that neighborhood, so one could read the 50th percentile of normalized reduction for a neighborhood as the median normalized reduction that a household would see for that neighborhood; 90th percentile indicates closer to a worst-case scenario for the reduction that a household in that neighborhood might see.

Downstream throughput. Figure 5 shows a consistent trend, where the normalized reduction in downstream throughput in South Shore tends to be more severe than city-wide communities; in contrast, households in Logan Square see much less normalized reduction in throughput. Households with the worst performance in Logan Square have a reduction of 0.27, which is significantly lower than the 0.62 reduction for users with the worst performance in South Shore. Thus, the reduction for South Shore tends to be twice that of Logan Square for households, across all service tiers—an effect size that is significant, given the size of our sample. These trends hold regardless of speed test, as well (i.e., for both Ookla Speedtest and NDT7).


Upstream throughput. Figure 6 shows that the top 10% of households with the best performance across all communities tend to have very low reductions, and the median user still has small, yet larger reductions. For this groups of users, South Shore and city-wide communities have higher reductions than Logan Square. However, the Ookla speed test shows that for the 10% of users with the worst performance, Logan Square has a normalized reduction of 0.64, which is much higher than South Shore’s normalized reduction of 0.16. The NDT speed test also demonstrates that Logan Square has a significantly greater reduction than that of South Shore. Again, these trends hold regardless of speed test.

5.2 Latency


Metric: Absolute increase. We present latency results in terms of an *absolute increase* between the 10th and 90th percentile latency values for each household. As opposed to throughput, we present results in terms of an increase, since larger values of latency (as opposed to smaller values of throughput) represent worse performance. We do not, however, normalize the reduction, since latency values are well-known and do not necessarily correlate with speed tier—giving us the ability to both measure a baseline and construct groups that are large enough to compare without having to first normalize. As with throughput, each result which shows the 10th, median and 90th percentiles per neighborhood represents a percentile of *households* within that neighborhood, so one could read the 50th percentile of absolute increase for a neighborhood as the median increase that a household would see for that neighborhood.

Idle latency. Figure 7 shows a similar trend as throughput: South Shore exhibits far higher absolute increase in latency than Logan Square and city-wide neighborhoods for all speed tiers. The increase is particularly large for the 10% of households that see the largest absolute increases, with a difference of nearly 30% (from about 30 ms to 40 ms in the 90th percentile). These measures of latency underscore the importance of looking beyond conventional notions of “speed” when considering equity—where discrepancies across neighborhoods may appear in conventional metrics, such as throughput, they can and are often equally (or more) pronounced in metrics such as latency.

Latency under load. We next explore increases in latency under load, measured with both TCP and ICMP ping measurements concurrent with download and upload traffic, respectively. Figure 9 shows Logan Square participants experiencing significantly less absolute increases across the distri-



(a) ICMP.



(b) TCP.


Figure 8: Downstream latency under load distributions across Chicago communities, measured through ICMP and TCP.

bution of households in that neighborhood, versus the city-wide and South Shore households.

Figure 8 shows a trend for the 10th and 50th percentiles in which South Shore exhibits poorer performance than Logan Square for both TCP and ICMP. However, the absolute numbers for the 90th percentile of devices do not show a clear trend.


5.3 Retransmission (Loss) Rate

Metric: Retransmission (loss) rate. The Netrics platform performs continual measurements of retransmission



(a) TCP.

Figure 9: Upstream latency under load distributions across Chicago communities, measured through ICMP and TCP.



(a) ICMP.

rates for each household, giving us a distribution of retransmission rates for each household. Our results then show a distribution across households; in other words, the plots show loss rates for the 10th, median, and 90th percentile households for each neighborhood, across the period of study. Figure 11 shows that the best 10% of households do not experience significant packet loss rates, regardless of neighborhood. The loss rates in these cases are less than 0.1%, which could be considered generally reliable. The median loss rate for South Shore is 0.54%, while that of Logan Square is 0.012% and city-wide is 0.021%. The discrepancy between the South Shore and the other neighborhoods is much larger. In all




Figure 11: Retransmissions.

groups, the worst 10 percent of households experience loss rates that exceed 1%; notably, however, the median household in South Shore *also* experiences a loss rate approaching 1%, far higher than loss rates in other neighborhoods. This neighborhood-level discrepancy again underscores the importance of exploring equity of Internet experience across multiple dimensions.

6 Conclusion

Our study offers an initial, empirically-based approach for assessing whether Internet users in different geographic locations experience equitable Internet performance. While much research has been devoted to examining whether ISPs reliably provide service at a pre-determined benchmark, we also need to take seriously the question of whether ISP service is equitable across users. The idea of "digital redlining" raises questions about whether Internet infrastructure provides equity to users in different geographic locations, or if the stratified patterns long observed in housing, public education, transportation, environmental toxins and more might also be observed in Internet performance.

As an initial approach to answering this question, we deployed the Netrics measurement platform to nearly 100 unique households in Chicago. Our sampling approach aimed to be able to conduct pairwise comparisons between concentrated device deployments in two focus communities (South Shore on the city's South Side, Logan Square on the city's North Side) and a similarly-sized group of devices deployed more widely across the city.

Results show that there are important differences between Internet measurements taken from households in these three groups. The South Shore neighborhood, on the city’s historically Black South Side, shows consistently worse performance than the Logan Square neighborhood, on the city’s North Side.

We recognize that there are some important limitations to this study. First, although we offer a rationale for our sample size and distribution within our two focus communities, we also recognize that the question of how to properly sample Internet infrastructure is very much an open question. Without knowledge of the underlying distribution of this infrastructure, we cannot apply standard statistical sampling procedures to create a representative sample. Further research is needed to make progress on this question, but we believe this study offers an important first step towards answering it. Second, we recognize that our Internet performance measurements may not be driven entirely by the quality of Internet infrastructure. Other factors, such as the number of users and devices, or the quality of the modem/router, may be affecting our measurements. However, we attempt to address this challenge through our strategy of gathering measurements that are both a) longitudinal and b) from multiple households within specific community areas.

Acknowledgments

This work was funded in part by a data.org inclusive growth and recovery grant, as well as NSF Awards CNS-1704077, CNS-2028145, CNS-2223610, and CNS-2213821. We are grateful to many members of the Internet Equity Initiative at the University of Chicago who have helped enable this research including Grace Chu, Guilherme Martins, and Kyle MacMillan. We thank Dan Black and Dan Nicolae for discussions concerning the statistical analysis we have developed in this paper.

References

- [1] *Chicago Ward, Community Area and Neighborhood Maps*. https://www.chicago.gov/city/en/depts/dgs/supp_info/citywide_maps.html.
- [2] “AT&T’s Digital Redlining: Leaving Communities Behind for Profit.” In: (). https://www.digitalinclusion.org/wp-content/uploads/dlm_uploads/2020/10/ATTS-Digital-Redlining-Leaving-Communities-Behind-for-Profit.pdf.
- [3] *Grafana: The Open Observability Platform*. <https://grafana.com/>. URL retrieved March 2022.
- [4] *InfluxData*. <https://www.influxdata.com/>.
- [5] *iPerf - The ultimate speed test tool for TCP, UDP and SCTP*. <https://iperf.fr/>.
- [6] *FCC Measuring Broadband America: Fixed Broadband*. <https://www.fcc.gov/general/measuring-broadband-america-measuring-fixed-broadband>. URL retrieved March 2022.
- [7] *FCC MARS: MBA Assisted Research Studies*. <https://www.fcc.gov/general/mba-assisted-research-studies>. URL retrieved March 2022.
- [8] *ndt7 Protocol - NDT (Network Diagnostic Tool)*. <https://www.measurementlab.net/tests/ndt/ndt7/>.
- [9] *Netrics*. <https://github.com/chicago-cdac/nm-exp-active-netrics>.
- [10] Ookla. *Speedtest by Ookla - The Global Broadband Speed Test*. URL: <https://www.speedtest.net/>.
- [11] *Participate in a Study on Internet Performance & Access*. <http://datascience.uchicago.edu/broadband/participate/>.
- [12] *RIPE Atlas*. <https://atlas.ripe.net/>.
- [13] Srikanth Sundaresan, Sam Burnett, Nick Feamster, and Walter De Donato. “BISmark: A Testbed for Deploying Measurements and Applications in Broadband Access Networks.” In: *USENIX Annual Technical Conference*. Philadelphia, PA, June 2014, pp. 383–394. URL: <https://www.usenix.org/conference/atc14/technical-sessions/presentation/sundaresan>.
- [14] Srikanth Sundaresan, Walter de Donato, Nick Feamster, Renata Teixeira, Sam Crawford, and Antonio Pescape. “Broadband Internet Performance: A View From the Gateway.” In: *ACMSIGCOMM*. Toronto, Ontario, Canada, Aug. 2011, pp. 134–145. URL: <https://dl.acm.org/doi/10.1145/2018436.2018452?cid=81548029281>.
- [15] Keeanga-Yamahtta Taylor. *Race for profit: How banks and the real estate industry undermined black homeownership*. UNC Press Books, 2019.