

# Bridging the digital divide between research and home networks

Caleb Smith-Salzberg\*, Fraida Fund†, Shivendra Panwar†

\*Stuyvesant High School, New York City, New York

†Department of Electrical and Computer Engineering, NYU Tandon School of Engineering

**Abstract**—Computer networking testbeds have made it easier for researchers to conduct realistic evaluations of new network protocols or services. However, it is challenging for users to configure testbeds or other experimental networks so that they are representative of typical home broadband links. To address this issue, we have developed a tool with which experimenters can configure links whose characteristics are drawn from a dataset of over 8,000 profiles of home Internet links in the United States, including fiber, cable, DSL, and satellite Internet connections from a range of speed and cost tiers. This paper describes our tool, and explains how it may be used with a variety of experimental platforms. We hope to make it easier for researchers to mimic networks that are representative of a variety of home users, so as to potentially increase the relevance of their experiments to this population.

## I. INTRODUCTION

In this work, we address the “digital divide” that currently exists between home networks and research networks. There is a great deal of variation in the quality and affordability of Internet service plans across the United States. Many households in the United States do not have high-speed Internet available in their area, or cannot afford to pay for high-speed Internet. Similar variation exists among the population of Internet-connected households in other countries. Meanwhile, researchers working at a university or industry research lab typically have high-quality Internet connections, which they may use in their evaluation of network services and protocols. This disparity is a “digital divide” between research networks and home networks. In this paper, we present a tool aimed at bridging this divide.

Internet researchers and developers who want to evaluate an implementation of a new network protocol or service that is intended for general use on the Internet often use a small prototype network in the lab, or a research testbed such as GENI [1] or PlanetLab [2]. However, a lab setup that operates on a university network will not be representative of a single home broadband user, and also does not represent a variety of home users. Meanwhile, testbed hosts on GENI or PlanetLab are typically connected to the Internet or to one another primarily through university networks, or through dedicated research and education networks, and thus can be similarly unrepresentative (as discussed by PlanetLab designers in [3]). For example, the default link speed between connected InstaGENI nodes on GENI (when the experimenter does not specify another link speed) is 100 Mbps; in contrast, the global average peak connection speed in 2016 was 37.2 Mbps, and only three

countries have an average peak connection speed above 100 Mbps [4]. Other network quality metrics, such as latency and packet loss, may be similarly unrepresentative in the *default* configurations of a university lab or testbed network. Seattle, a testbed which aims to be more representative of “regular” connectivity by using donated resources, still includes a mix of home Internet connections and university connections [5], and may not be suitable for experiments because of limits on what experimenters are permitted to do on end-user devices.

Simulation platforms such as QualNet and ns-3 [6], and emulation platforms such as Mininet [7] and Emulab [8], offer experimenters a great deal of control over network characteristics, including the potential to more accurately represent typical home broadband networks. Similarly, link characteristics such as downlink capacity, uplink capacity, latency, jitter, and packet loss, can be modified on testbed or in-house experiment networks using emulation tools such as netem [9]. However, in all of these, the *experimenter* must supply the specific characteristics of the network to emulate, and it is not always obvious how to select these. Due in part to the difficulty of identifying a large number of appropriate configurations and reporting them in a way that allows others to reproduce the work, many researchers select one or two representative network settings and focus exclusively on those, potentially missing important results that might occur if the evaluation included a wider range of link characteristics.

When a protocol or service is not evaluated on a wide range of realistic networks, it may not work as intended for some users. This lesson has been acknowledged in industry; for example, Facebook instituted a “2G Tuesdays” program where employees have the option of using low-quality Internet speeds similar to that of the developing world, in order to get a better understanding of how people with poor Internet quality experience their applications [10]. Using tools to support evaluation on realistic networks led the Facebook team to change the Messenger app work so that it would work better for users in emerging markets [11].

In other aspects of network experiments, data collected on real networks is used to make experiments more realistic. For example, tools such as Tmix [12] help researchers represent realistic network workload, based on real traffic captures, in experiments. Similarly, empirical user mobility traces (available in repositories such as Crawdad [13]) allow researchers to represent realistic mobility patterns in simulations. We use a similar data-based approach to create a tool for the selection

Selected household 5670 has the following characteristics:  
 Plan: 20.0/0.896 (Mbps down/up), CenturyLink (DSL), AZ  
 Estimated price per month: \$59.00

```
-----
Upload rate (kbps)      | 767
Download rate (kbps)   | 16627
Round-trip delay (ms)  | 29.914500
Uplink jitter (ms)     | 1.633500
Downlink jitter (ms)   | 0.913000
Packet loss (%)        | 0.000000
-----
```

Listing 1: A representative household profile.

of network characteristics for testbed experiments

The goal of this work is to make it easy for researchers to evaluate networked applications, services, and protocols, on networks that are representative of a wide range of home broadband users, so as to potentially increase their impact on real Internet users. Specifically, we suggest an approach similar to empirical research in many other fields, in which experimenters draw many samples of link characteristics from a sampled population that is representative of a target population: either all broadband links in the United States, or a specific subset of them. We contribute a tool for use with this methodology, including a dataset of link characteristics based on empirical data of home broadband connections in the United States. It produces several output formats that can be directly used to reserve resources on GENI [1], to apply `tc` and `netem` [9] rules on any Linux-based experiment setup, with Facebook’s Augmented Traffic Control tool [11], and with the `geni-lib` [14] Python library.

The rest of this paper is organized as follows. In Section II, we explain how we derive the data used in our “digital divide” tool from several public data sources. Next, in Section III, we discuss the usage of the tool and the output formats it produces. In Section IV we describe how we validate this tool and compare the requested and actual network characteristics for a series of household profiles realized on GENI. Finally, in Section VI we conclude with a brief discussion of some of the limitations of our approach and directions for future work.

## II. DATA SOURCES FOR THE “DIGITAL DIVIDE” TOOL

At the core of our “digital divide” tool is a dataset of approximately 8,000 household profiles. In this section, we describe how these profiles are derived.

Listing 1 shows a sample profile. It includes an identifier unique to that profile (which can be published to aid reproducibility), details of the Internet service plan to which the household subscribes (included estimated monthly cost of service), and a series of link quality metrics. An additional profile field which is not shown is the weight, which is used when selecting profiles from the sampled population. To generate the 8,000 profiles we use three open data sources: the Measuring Broadband America data [15] for link quality metrics and service plan details except cost of service, the Urban Rate Survey [16] for estimated cost of service, and the United States Census Current Population Survey Computer and Internet Use Supplement [17] for weighting samples. Fig. 1 shows a summary view of this workflow.

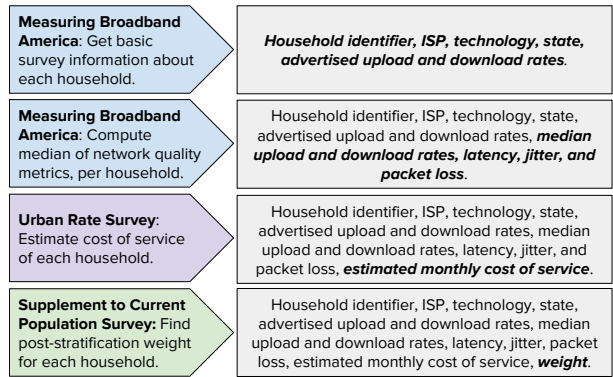


Fig. 1: Data processing workflow. The boxes on the right show how data is progressively added to each household profile to derive the final dataset, with the emphasized text indicating newly added fields at each stage. On the left, we show the data that is the source of the new fields, and the computations with which they are derived.

### A. Link quality metrics

The FCC’s Measuring Broadband America (MBA) program [15] is a study of home broadband networks in the United States. Its goal is to collect information about the empirical quality of Internet service at households in the United States, in different regions and with different Internet Service Providers (ISPs), so as to provide consumers with information regarding their broadband options. The measurements collected by this program are available as open data.

The MBA data is especially valuable in this context because unlike many other Internet performance datasets, it does not rely on user-initiated measurements. In datasets of user-initiated measurements, such as from speed test services [18], measurements may be skewed towards users who are experiencing problems with their Internet service (which motivates them to run the speed test). In the MBA program, volunteer panelists are given a wireless router, through which they connect to the Internet using their regular ISP. The router automatically runs tests at regular intervals and reports the results, avoiding the bias inherent in user-initiated measurements.

We selected key metrics of home Internet service quality, and used the MBA datasets released in 2015 and 2016 to generate observations of these metrics for each panelist. The data includes 8,119 unique panelists with complete information; for panelists that are present in both years, we use the data from the first year. Table I lists these metrics, along with the methodology by which they are measured for the MBA program [20] and the reason each was selected. For each metric, we compute its median value for each household during a one month interval (September of 2014 for the 2015 data, September of 2015 for the 2016 data). This subset of the data is “validated” to ensure the accuracy of the users’ reports of their Internet service plan (ISP, advertised upload rate, advertised download rate), as described in [15], [20].

### B. Monthly cost of service

Research on computer networks often includes an analysis of consumer utility, cost modeling, and related features; there-

Metric	Test methodology	Reason for selection
Download and upload speed	Measures download and upload throughput every two hours by performing three simultaneous HTTP GET and HTTP POST requests to a target test node outside of the user's own ISP's network. The test operates for 30 seconds, and the initial TCP "warm-up" period is excluded from measurement results.	The empirical download and upload speed of a link is often different from the advertised rates. We include observations of this metric as measured empirically, so as to more accurately represent "real" link capacity available to the end user.
Latency, jitter, and packet loss	Measures a bidirectional stream of UDP packets at 64 kbps (meant to resemble VoIP traffic), transmitted hourly between the router and a target test node outside of the user's own ISP's network. The test operates for 30 seconds.	For fast broadband connections, web performance is often bottlenecked by latency [19]. Jitter and packet loss are also important for video streaming and voice over IP applications. All of these metrics are also known to affect the congestion control behavior of TCP.

TABLE I: Metrics of network quality included in our household profile, and methodology by which they are measured.

fore, we are also interested in the estimated monthly cost of Internet service of each household. In the next stage of our data processing workflow (Fig. 1) we estimate cost using the FCC's Urban Rate Survey [16], which tells us for each ISP, the monthly cost of service for a plan with a given upload and download rate, in a particular urban area. Based on this data, we assign a dollar amount (monthly cost of service) to each household based on the following rules:

- 1) For 2665 of 8119 households (33%), there is at least one plan in the Urban Rate Survey offered by the same ISP in the same state, using the same technology, with the same advertised upload and download rates. If there are multiple matching plans (e.g., with different monthly data caps), we sample from them with uniform probability.
- 2) For another 1697 of 8119 households (21%), there is at least one plan in the Urban Rate Survey offered by the same ISP, using the same technology, with the same advertised upload and download rates, but not in the same state. If there are multiple matches (e.g., with different monthly data caps), we sample from them with uniform probability.
- 3) For 912 of 8119 households (11%), there is at least one plan in the Urban Rate Survey using the same technology, with the same advertised upload and download rates, but not necessarily in the same state or from the same ISP. If there are multiple matches (e.g., with different monthly data caps), we sample from them with uniform probability.
- 4) For 2626 of 8119 households (32%), there is no plan with the same technology and advertised upload and download rates. To estimate a cost for these, we find the three plans using the same technology (cable, fiber, or DSL) with the closest advertised upload and download rates, then sample from them with uniform probability. For each household in this group, Fig. 2 shows how much the plan whose cost was assigned deviates from the actual advertised upload and download speeds of that households' plan.
- 5) The remaining 219 households (3%) have satellite Internet service, which is not included in the Urban Rate Survey. Since only three plans (combination of ISP, download rate, and upload rate) are represented, we manually found the advertised price of each plan in 2014 and assign this known value to each household according to its plan. One satellite ISP offers multiple plans with the same download/upload rate with different data caps; for this ISP, we randomly select one of the plans (with uniform probability) and assigned its cost to the household.

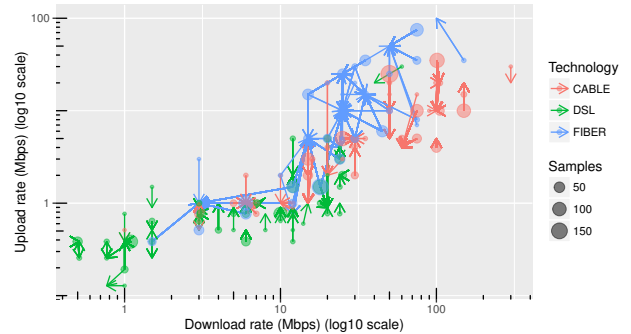


Fig. 2: Mapping of group 4 households in the MBA data (samples for which there is no plan with the same technology and advertised upload and download rates) to plans in the Urban Rate Survey. The size of the circle denotes the number of samples having that upload and download rate; the error vector shows the upload and download rate of the rate plan to which they are mapped.

### C. Post-stratification weight

MBA panelists are selected to have geographic diversity, diversity in type of broadband technology, and diversity in Internet service plans, so that the measurements provide information about the full range of Internet quality experienced by home users in the United States. Specifically, a target number of panelists was set for each of three speed tiers ( $< 3$  Mbps,  $3 - 10$  Mbps, and  $\geq 10$  Mbps) in each of the four census regions in the United States (Northeast, Midwest, South and West), and a secondary goal involved covering major ISPs in each of the broadband technologies (cable, DSL, fiber, satellite) [15]. However, the distribution of these and other characteristics among MBA panelists does not represent their distribution among the general population of Internet-connected households in the United States. Some groups are oversampled in the MBA program so as to support statistically valid inferences based on the gathered data.

Thus, in the final data processing stage we assign a post-stratification weight to each household, so that random weighted samples from the dataset will have the same distribution of broadband Internet technology and cost of monthly service as the population of U.S. households with broadband Internet. These weights are computed according to the prevalence of the technology (e.g., cable, fiber, DSL, or satellite) and monthly cost of service (in brackets: \$0, \$0-39.99, \$40-79.99, \$80+) in the U.S. population, as measured by responses to the 2013 supplement on "Computer and Internet Use" to the Current Population Survey of the United States Census [17].

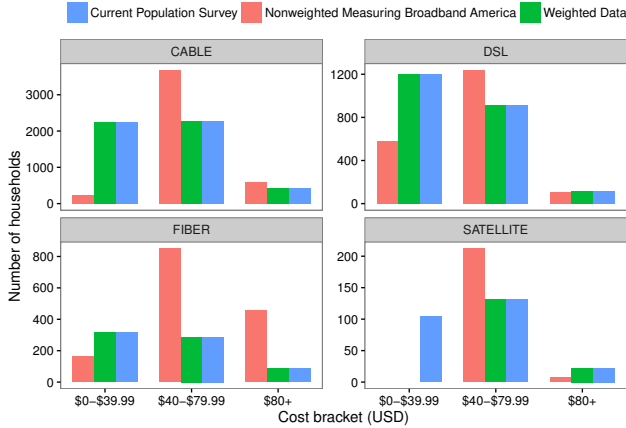


Fig. 3: Distribution of samples before and after weighting. The figure shows (1) expected number of households in each category according to the Computer and Internet Use Supplement to the Current Population Survey conducted by the United States Census Bureau, (2) number of households in each category in the validated Measuring Broadband America data, and (3) the weighted number of households in each category in our final re-weighted data set.

This supplement asks, among other questions, what kind of technology the respondent’s home Internet service uses, and how much the respondent pays for Internet service every month. (Those who purchase Internet service as part of a bundle including TV or other services are asked what the cost of the Internet part of the bundle is.) In computing the weights, we consider only the portion of the population with home Internet service from exactly one source, and only those whose Internet service is either cable, fiber, DSL, or satellite.

Fig. 3 shows the count of unweighted samples in each category, the count of samples that would be expected based on U.S. Census responses [17], and the sum of weighted samples in each category. For some categories, e.g., low-cost satellite service, we have no samples in the MBA data.

### III. THE “DIGITAL DIVIDE” TOOL FOR GENERATING REPRESENTATIVE NETWORKS

The “digital divide” tool is a Python script that uses the dataset produced as described in Section II to generate configuration files or other outputs that can be used directly in a testbed or other experimental setting.

When used without any configuration, the “digital divide” script selects a weighted sample from the dataset of  $\approx 8000$  households, and produces some output that represents the network characteristics of the selection. The intent is for users to run experiments many times, each time drawing a new random sample, so that the overall result is then more broadly representative of the U.S. population. However, we also provide some filters that experimenters may use to sample from a specific subset of the population, as well as a variety of output formats for use in different experiment frameworks.

#### A. Filters

The tool supports several command-line arguments that allow users to sample from a smaller subset of the population




Fig. 4: Uploading a Request RSpec generated by the “digital divide” tool to the GENI Portal will create a star topology, as shown in the representative example, with one “router” and one or more “households” (labeled with household identifiers).

with specific technical or demographic characteristics:

- **Technology:** Researchers may focus on one kind of broadband technology: cable, fiber, DSL, or satellite. This can be useful when evaluating a protocol that is designed for a certain technology (for example, TCP variants intended for satellite connections with extremely high latencies) or when comparing the behavior of a protocol across technologies.
- **State:** For investigating regional differences in Internet quality, e.g., for comparisons between populations that vary in the quality and quantity of local Internet service offerings, users may specify a state from which to draw samples.
- **Price range:** Researchers can also specify a price range, and limit the output to samples whose monthly estimated cost of service is in that range.


These may be applied individually or in combination; however, users should take care in combining filters, as they may then be sampling from a very small set meeting the criteria.

Experimenters can also specify the household identifier as a command-line argument, to get the profile of a specific household. This is intended to promote reproducibility. Researchers can record identifiers of the households that were randomly selected for their experiments, and report these with their results. Then, others can replicate their work with the same households, passing the household identifier as a command-line argument (similar to the way that a seed for a random number generator facilitates a replication of a simulation).

#### B. Output formats

The “digital divide” tool can produce a file or other output that may be used directly in some of the most popular experimental frameworks for computer networks:

- **Request RSpec:** GENI [1] and many other testbeds use Resource Specification (RSpec) files to describe and reserve resources. The Request RSpec generated by our tool may be submitted directly to the GENI Portal or used with other resource reservation tools that support the RSpec format. The RSpec describes a star topology, as shown in Fig. 4, with a “router” node and one or more “household” nodes (with the number of “household” nodes determined by another command-line argument). The capacity requested for each router-household link is the maximum of the uplink and downlink rates for the household. Postboot commands defined in the RSpec use `netem` to emulate the network characteristics of each household on those links.



(a) For the downlink, measured throughput values were 95%, 98%, and 113% of the expected value for the 5th, 50th, and 95th percentiles, respectively. For the uplink, these were 96%, 113%, and 118%.

(b) For the downlink, measured jitter values were 63%, 69%, and 86% of the expected value for the 5th, 50th, and 95th percentiles, respectively. For the uplink, these were 62%, 68%, and 74%.

(c) Measured times were 101%, 103%, and 108% of expected for 5th, 50th, and 95th percentiles.

Fig. 5: Network characteristics for 800 sampled households. Each point represents the mean of measured value at one household, and a vertical line marks one standard error above and below the mean. Points below the blue line the line indicate that the observed value of the metric was smaller than the target value. Captions give 5th, 50th, and 95th percentile deviation of the measured value from the target.

- **Augmented Traffic Control (ATC) profile:** Augmented Traffic Control [11] is a tool created by Facebook for measuring application performance over a variety of (emulated) network conditions. With ATC running on a router, experimenters can select from a list of installed “profiles” using a browser-based UI, and apply the network characteristics associated with that profile to all traffic between the router and a particular host. We have found this tool to be a useful alternative or complement to remote-access testbeds for evaluations that require the use of a graphical interface [21], [22]. Our “digital divide” tool generates ATC profile files that may be added directly to an ATC installation.
- **tc and netem commands:** This option prints as output a set of commands that install Linux traffic control (tc) options to mimic the link characteristics of the selected household. This format is intended for use in other Linux-based experiment settings, and may also be used directly within an experiment script (as we did in Section IV).

Users may also print the sampled households’ broadband characteristics in a table, as shown in Listing 1.

We also provide two Python classes that may be used in a `geni-lib` script: a `Household` class that is initialized with a row of data from the sample dataset and sets up attributes representing its network characteristics, and a `Star` class that represents a star topology with a router and one or more households (as in Fig. 4). `Star` objects can be used within a larger `geni-lib` script to create more complex topologies with multiple connected stars.

#### IV. VALIDATION

In this section, we investigate how accurately the network characteristics of a sampled household are realized on a testbed link. We ran the “digital divide” tool 800 times, each time generating `tc` and `netem` commands and applying them to a link between two virtual machines on GENI (with the 800 experiments spread across ten InstaGENI aggregates). For each network configuration, we ran the following, ten times each:

- `ping` for 100 ICMP messages and replies, at 0.2 second intervals, to measure latency,

- `iperf3` for 60 seconds with three parallel TCP flows, in both directions (sequentially), to measure throughput,
- `iperf` with 100 kbps UDP packets for 60 seconds in both directions (sequentially), to measure jitter and packet loss.

We use these for validation because they are similar to the measurement tools with which the MBA data is collected.

Fig. 5 shows the results of these tests. Most samples had a measured throughput close to the target value. At low or moderate throughputs, error was not biased in either direction, while for higher throughputs the measured value was usually lower than the target. (This may be due to the testing procedure, e.g., the three parallel TCP flows may not have saturated the link, or it may indicate insufficient ability of `tc` to handle traffic at a high data rate.) Measured jitter was almost always less than the configured value. However, because different utilities (include `iperf` and `netem`) measure jitter differently, it is difficult to compare jitter values across different tools. The jitter values as measured by `iperf` are consistently about 68% of what we set with `netem`. Almost all of the households had measured latencies slightly higher than the target (usually by less than one millisecond), which we attribute to the inherent latency of the link that is added to the emulated latency.

We do not include packet loss in Fig. 5 because the majority of households (about 700) were configured to have zero packet loss, and did. About ten households had some packet loss that were supposed to have none. Another 90 households were configured to have packet loss, and did, but usually a lower packet loss ratio than expected (50% of the configured value, for the median case). Because the ratio of packets lost was very small, even for these households, it is difficult to accurately measure mean packet loss without a prohibitively long experiment, and so there was wide variation in the measured packet loss across trials of the same household.

Our validation experiments included a small group of outliers in which the difference between the expected and realized network characteristics is much greater than for other experiments. We attribute these to occasional transient testbed problems (e.g., CPU or network congestion) or software issues related to the measurement tools (e.g., `iperf`).

A demonstration of the capabilities of this tool and its use in experiments is available at [22]. Also, a public `git` repository for this project [23] includes:

- Python source code for the “digital divide” tool and library. The code is released under the MIT License.
- A set of `bash` and `R` scripts that retrieves the raw data sources from the Internet and applies the procedure of Section II to recreate our dataset of  $\approx 8000$  households, and generate Fig. 2, and Fig. 3.
- The dataset, in CSV format.
- A set of `bash` and `R` scripts and instructions to run the experiment and generate the figures of Section IV.

## VI. CONCLUSION

In this paper, we present to the research community a new tool for setting up testbed experiments involving networks that are representative of real home broadband links.

We briefly discuss here some limitations of our approach. Our data is exclusively from the United States; while SamKnows (who conducts the Measuring Broadband America campaign) has similar programs in the United Kingdom, Canada, Singapore, Brazil, and the European Union, this data is not available for redistribution. There is also a time lag between the collection of measurements of home broadband speed, and their release. Because the quality and availability of Internet connectivity is changing rapidly, the most recent available data - which can be up to two years old - may not fully represent current networks. In the MBA data, only a few home network technologies are included. Also, the measurements are conducted on the router, and exclude the effects of the WiFi link through which many users connect to their home broadband. Cellular networks are increasingly used for home broadband access, but are not represented in our data. Finally, the specific data points available, and the way we use them, are only an approximation of the true experience of home broadband users. For example, the data only includes measurements for links between the home user and a small set of remote sites, and may not capture many problems in network performance associated with ISP policies (e.g., traffic shaping applied to certain classes of traffic) or problems that occur along the path to *other* remote sites (e.g., buffer bloat, congested exchange points). We also do not include time of day and day of week variation in network performance, instead using the median of all measurements, all the time.

As future work, we would like to increase the geographical and technological diversity of households represented in this tool by including open data from other sources.

## ACKNOWLEDGMENT

This work was supported by the NSF Graduate Research Fellowship Program, the New York State Center for Advanced Technology in Telecommunications, the NYU Tandon Center for K12 STEM Education, and the Pinkerton Foundation.

- [1] M. Berman, J. S. Chase, L. Landweber, A. Nakao, M. Ott, D. Raychaudhuri, R. Ricci, and I. Seskar, “GENI: A federated testbed for innovative network experiments,” *Computer Networks*, vol. 61, pp. 5–23, 2014.
- [2] B. Chun, D. Culler, T. Roscoe, A. Bavier, L. Peterson, M. Wawrzoniak, and M. Bowman, “Planetlab: an overlay testbed for broad-coverage services,” *ACM SIGCOMM Computer Communication Review*, vol. 33, no. 3, pp. 3–12, 2003.
- [3] N. Spring, L. Peterson, A. Bavier, and V. Pai, “Using PlanetLab for network research: myths, realities, and best practices,” *ACM SIGOPS Operating Systems Review*, vol. 40, no. 1, pp. 17–24, 2006.
- [4] A. Technologies, “Akamai state of the Internet report, Q3 2016,” <https://www.akamai.com/uk/en/our-thinking/state-of-the-internet-report/>, 2016.
- [5] Y. Zhuang, A. Rafetseder, and J. Cappos, “Experience with Seattle: A community platform for research and education,” in *Second GENI Research and Educational Experiment Workshop*, 2013, pp. 37–44.
- [6] G. F. Riley and T. R. Henderson, “The ns-3 network simulator,” in *Modeling and Tools for Network Simulation*, K. Wehrle, M. Güneş, and J. Gross, Eds., 2010, pp. 15–34.
- [7] B. Lantz, B. Heller, and N. McKeown, “A network in a laptop: rapid prototyping for software-defined networks,” in *Proceedings of the 9th ACM Workshop on Hot Topics in Networks*, 2010, p. 19.
- [8] M. Hibler, R. Ricci, L. Stoller, J. Duerig, S. Guruprasad, T. Stack, K. Webb, and J. Lepreau, “Large-scale virtualization in the Emulab network testbed,” in *USENIX Technical Conference*, 2008, pp. 113–128.
- [9] S. Hemminger, “Network emulation with netem,” in *Proceedings of the 2005 Linux Conference Australia, Canberra, Australia*, 2005.
- [10] C. Marra, “Building for emerging markets: The story behind 2G Tuesdays,” <https://code.facebook.com/posts/1556407321275493/building-for-emerging-markets-the-story-behind-2g-tuesdays/>, October 2015, Facebook Code Blog.
- [11] J. M. Manu Chandra, “Augmented Traffic Control: A tool to simulate network conditions,” <https://code.facebook.com/posts/1561127100804165/augmented-traffic-control-a-tool-to-simulate-network-conditions/>, March 2015, Facebook Code Blog.
- [12] M. C. Weigle, P. Adurthi, F. Hernández-Campos, K. Jeffay, and F. D. Smith, “Tmix: a tool for generating realistic TCP application workloads in ns-2,” *ACM SIGCOMM Computer Communication Review*, vol. 36, no. 3, pp. 65–76, 2006.
- [13] D. Kotz and T. Henderson, “Crawdad: A community resource for archiving wireless data at dartmouth,” *IEEE Pervasive Computing*, vol. 4, no. 4, pp. 12–14, 2005.
- [14] “geni-lib,” <https://geni-lib.readthedocs.io/>, 2014–2016.
- [15] United States Federal Communications Commission, “2015 measuring broadband America fixed broadband report,” <https://www.fcc.gov/reports-research/reports/measuring-broadband-america/measuring-broadband-america-2015>, December 2015.
- [16] —, “Urban rate survey data and resources,” <https://www.fcc.gov/general/urban-rate-survey-data-resources>, 2015.
- [17] United States Census Bureau, “Current population survey: July 2013 computer and Internet use supplement,” [http://thedataweb.rm.census.gov/ftp/cps\\_ftp.html#cpsupps](http://thedataweb.rm.census.gov/ftp/cps_ftp.html#cpsupps), July 2013.
- [18] C. Dovrolis, K. Gummedi, A. Kuzmanovic, and S. D. Meinath, “Measurement lab: Overview and an invitation to the research community,” *ACM SIGCOMM Computer Communication Review*, vol. 40, no. 3, pp. 53–56, 2010.
- [19] S. Sundaresan, N. Feamster, R. Teixeira, and N. Magharei, “Measuring and mitigating web performance bottlenecks in broadband access networks,” in *Proceedings of the 2013 Conference on Internet Measurement Conference (IMC '13)*, 2013, pp. 213–226.
- [20] SamKnows, “SamKnows test methodology: Methodology and technical information relating to the SamKnows testing platform,” <https://www.samknows.com/broadband/uploads/methodology/SQ301-005-EN-Test-Suite-Whitepaper-4.pdf>, February 2015.
- [21] F. Fund, “2G Tuesdays: emulating realistic network conditions in emerging markets,” <https://witestlab.poly.edu/blog/2g-tuesdays-emulating-realistic-network-conditions-in-emerging-markets/>, August 2016, “Run my experiment on GENI” blog post.
- [22] C. Smith-Salzberg, “Bridging the empathy gap between web developers and users,” <https://witestlab.poly.edu/blog/bridging-the-empathy-gap-between-web-developers-and-users/>, August 2016, “Run my experiment on GENI” blog post.
- [23] “Digital divide,” <https://github.com/csmithsalzberg/digitaldivide>.