Crowdsourced Active Speed Test Measurements and the Importance of Contextualization

Elizabeth Belding
ebelding@ucsb.edu
Dept. of Computer Science
University of California, Santa Barbara, USA

ABSTRACT

Crowdsourced speed test measurements, such as those by Ookla® and Measurement Lab (M-Lab), offer a critical view of network access and performance from the user's perspective. However, we argue that taking these measurements at surface value is problematic. It is essential to contextualize these measurements to understand better what the attained upload and download speeds truly measure. To this end, we develop a novel Broadband Subscription Tier (BST) methodology that associates a speed test data point with a residential broadband subscription plan. Our evaluation of this methodology with the U.S. Federal Communication Commission's (FCC) Measuring Broadband America dataset shows over 96% accuracy. We augment approximately 1.5M Ookla and M-Lab speed test measurements from four major U.S. cities with the BST methodology. We show that many low-speed data points are attributable to lower-tier Internet subscriptions and not necessarily poor access. Then, for a subset of the measurement sample (80k data points), we quantify the impact of access link type (WiFi or wired), WiFi spectrum band and signal strength, and device memory on speed test performance. Based on our results, we put forward a set of recommendations for both speed test vendors and the policymakers to contextualize speed test data points and correctly interpret measured performance.

KEYWORDS

broadband mapping, network measurement, Internet access

ACM Reference Format:
Elizabeth Belding. 2024. Crowdsourced Active Speed Test Measurements and the Importance of Contextualization. In *Proceedings of ACM WomENcourage: Responsible Computing for Gender Equality (womEnNcourage'24).* ACM, New York, NY, USA, 2 pages. https://doi.org/XXXXXXX.XXXXXXX

1 INTRODUCTION

The challenge of mapping fixed broadband Internet access was brought to the forefront during the stay-at-home orders of the Covid-19 pandemic. Suddenly, individuals without high-quality Internet access could not participate in the remote schooling, work, and telehealth that these orders required. While government grants are often available to help Internet service providers (ISPs) build infrastructure in more difficult-to-reach areas, a key challenge is knowing exactly where high-quality Internet access is lacking. Crowdsourced network measurements have emerged as a powerful tool to map fixed broadband access more accurately. These "speed tests" provide a critical snapshot of the network state from the vantage point of the end users. Popular network speed test platforms, such as Ookla's speedtest.net and Netmeter.eu, are utilized by Internet users worldwide. Because of the inherent benefits, governmental initiatives often rely on speed test data to map broadband access and discern where to make the economic investment in infrastructure to address digital inequality.

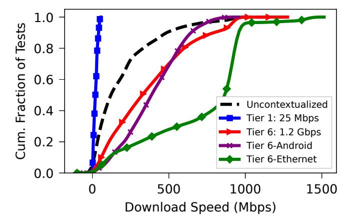

Figure 1: Comparison of raw speed test download speed distributions in a major U.S. city.

However, despite the broad use of these measurements, the data generated through speed tests suffers from several key limitations that must be addressed before drawing meaningful conclusions about fixed Internet performance. *We argue that speed test measurements must be contextualized to accurately interpret the measured performance.* The challenge is understanding *what a speed test measures* and *how it compares to expected speed values.* For example, many fixed broadband plans offer rates as high as 1 Gbps download and 35 Mbps upload. If a speed test measures performance significantly less than these values, is it because the access network is under-performing, the user has purchased a lower-tier (slower) plan, or the user's home WiFi network is misconfigured or experiencing interference? If the under-performance is due to issues in the access network, then the problem could be reported to the ISP to challenge coverage claims in an area. In contrast, if the under-performance is due to local factors, such as channel interference or poor signal quality, the user can address it directly. If the user purchased a lower-tier plan, then perhaps the speed test is measuring the paid-for speed. Finally, the methodology of the test itself can impact performance results, adding additional complexity.

As our poster will illustrate, the lack of context for these measurements prevents proper interpretation of results. As an example, figure 1 presents the distributions of 214k Ookla download speed test results in a major US city disaggregated by subscription plan tiers, access speed or link type, and measurement device type. The "uncontextualized" line represents the original data without context

womEnNcourage'24, June 26-28, 2024, Madrid, Spain

ACM ISBN 978-1-4503-XXXX-X/18/06
https://doi.org/XXXXXXX.XXXXXXX

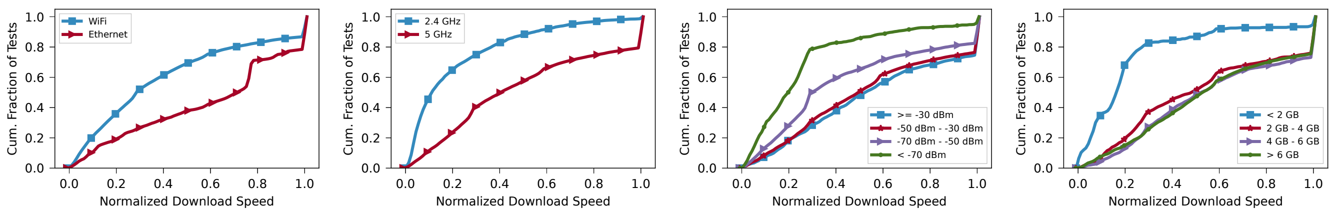

(a) Access Type **(b) WiFi Band** **(c) RSSI** **(d) Available Memory**

Figure 2: Impact of WiFi characteristics and available memory on speed test performance.

applied. The figure shows that the median download speed of the lowest (slowest) subscription tier (Tier 1, with a maximum download speed of 25 Mbps) is 19.22 Mbps, almost six times as slow as the overall median download speed. The median download speed, on the other hand, is nearly four times less than the premium ISP subscription tier (Tier 6: 1.2 Gbps) and almost seven times less than that recorded by test takers on Tier 6 Ethernet connections (Tier 6: Ethernet). Similarly, for speed tests that do not experience local bottlenecks from WiFi, the median download speed of the highest subscription tier for this group of speed tests (Tier 6: Android) is almost four times more than the median download speed. Based on these significant differences, our work proposes mechanisms for contextualizing crowdsourced speed test measurements so that the performance of the access network can be properly interpreted.

2 CONTEXTUALIZATION ANALYSIS

To contextualize speed test data, we first address the challenge of determining the subscription tier of the user conducting the speed test. We develop a novel Broadband Subscription Tier (BST) methodology that associates a speed test data point with a residential broadband subscription plan. Our evaluation of this methodology with a ground truth dataset shows over 96% accuracy. Then, we augment approximately 1.5M Ookla and M-Lab speed test measurements from four major U.S. cities with the BST methodology. We show that many low-speed measurements are attributable to lower-tier subscriptions and not necessarily poor access network performance. Then, for a subset of the measurement sample (80k data points), we quantify the impact of access link type (WiFi or wired), WiFi spectrum band and signal strength, and device memory on speed test performance, as shown in figure 2.

Given the challenges and complexities of WiFi communication, in figure 2(a) we examine the speed test results conducted over WiFi with those from desktop computers connected to the home network via Ethernet. The figure shows that the difference in the normalized download speed distributions of WiFi and Ethernet access links is significant. For speed tests conducted over a WiFi network, the median normalized download speed is 0.28. This value is almost three times less than the median normalized download speed of 0.71 for Ethernet speed tests. Without proper contextualization, the lower download speeds from tests conducted over WiFi could be misconstrued to be under-performance of the provider network.

Next, we examine WiFi speed test performance and the impact of the WiFi spectrum band. Figure 2(b) shows the distribution of the normalized download speed separated by the WiFi band. The figure shows a striking difference between the performance of tests in the two bands. While the median normalized download speed is just 0.11 for 2.4 GHz speed tests, it is 0.4 for 5 GHz tests.

We next analyse the impact of WiFi signal strength (RSSI) on speed test performance. Because our analysis showed that 2.4 GHz tests under-perform compared to 5 GHz tests, here we only consider 5 GHz WiFi band tests here. We bin the tests into four bins and we calculate the distance between the measured and subscribed performance for each test. Figure 2(c) shows the distribution of the normalized download speed achieved by speed tests for each RSSI bin. The performance difference varies by over a factor of two between the lowest and highest RSSI bins. This difference increases to more than five when considering speed tests in the fastest tier.

Finally, we study the memory available to the Android device kernel during the speed tests to understand its role in performance. Figure 2(d) presents the CDFs of the distance between subscribed and achieved speed test performance grouped by available kernel memory. The distance increases as less memory is available to the kernel during the speed test, showing that speed test performance can be greatly impacted by available memory and is therefore another important piece of context for speed test measurements.

More comprehensive results will be shown in the final poster, including a performance comparison of aggregate performance of tests with multiple local performance inhibitors (WiFi, poor RSSI, limited kernel memory).

3 CONCLUSION

In this work, we showcase the importance of contextualizing crowdsourced speed test datasets, analyzing and quantifying the impact of a variety factors that can degrade speed test performance. The impacts we uncover, which at times differentiate performance more than seven-fold, underlines the need for meaningful contextualization of crowdsourced speed test measurements prior to drawing generalizable conclusions about regional broadband access and quality. This is particularly important for policymakers prior to basing funding decisions on this data. We believe that the need for accurate broadband mapping has never been greater, and that crowdsourced speed test measurements will provide an invaluable part of the data needed to generate these maps. We hope that our work contributes to the advancement of this critical mapping effort.

4 ACKNOWLEDGMENTS

We thank Ookla and Zillow for sharing their datasets with us. This work was funded through National Science Foundation Rapid Response Research (RAPID) award NSF-2033946. This work is joint with Udit Paul (Ookla); and Jiamo Liu, Mengyang Gu and Arpit Gupta (UC Santa Barbara).

Crowdsourced Active Speed Test Measurements and the Importance of Contextualization

Elizabeth Belding, University of California, Santa Barbara ebelding@ucsb.edu

acm's women in computing
acm-w
Europe

Motivation

- Challenge: measuring fixed broadband access and quality
- Crowdsourced speed test measurements: an important view of network access and performance from the user's perspective
- What do speed tests actually measure? If a measured speed is low, is it because:
  a) the access network is underperforming;
  b) the user has purchased a lower tier (i.e. slow) access plan;
  c) the local WiFi network is misconfigured or experiencing interference;
  d) the measurement device has limited resources?
- Contextualizing the measurement, to know what it is actually measuring, is essential for identifying areas that are underserved!

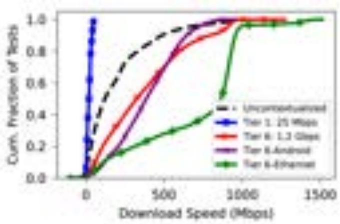

Tier 1-3: 5 Mbps upload speed; 25-200 Mbps download speed
Tier 4: 10 Mbps upload speed; 400 Mbps download speed
Tier 5: 15 Mbps upload speed; 800 Mbps download speed
Tier 6: 35 Mbps upload speed; 1200 Mbps download sped

Dataset

Speedtest Data:

- Ookla Speedtest: an internet speed testing application that can be used on web browsers, desktop computers, iOS, and Android smartphones. We obtain the proprietary dataset through a Data Use Agreement.
- Measurement-Lab's Speed Test: utilizes the Network Diagnostic Tool (NDT)
- Measuring Broadband America (MBA): sponsored by US FCC, 4,000 units in US households; groundtruth for broadband plan subscriptions
- The data contains measurements crowdsourced from Ookla, M-Lab and MBA between January 1 and December 31, 2022, for four US cities.

City/State	ISP	Ookla	M-Lab	MBA
A	1	214k	113k	25.9k
B	2	205k	376k	14.9k
C	3	128k	64k	10.9k
D	4	198k	166k	8.9k

Results

Impact of WiFi Characteristics and Available Memory on Speed Test Performance

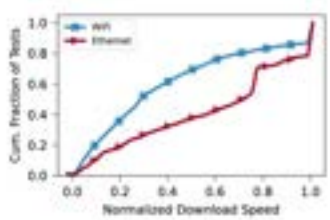

Access Type

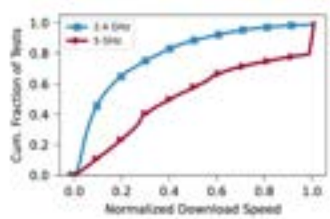

WiFi Band

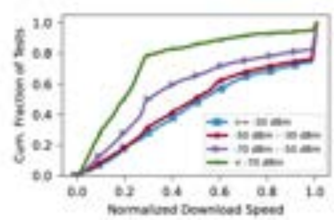

RSSI

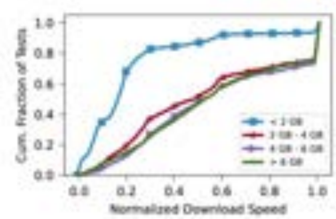

Available Memory

Access Link: WiFi median normalized download speed is 0.28. Ethernet median normalized download speed is 0.71.

WiFi Band: 0.11 median normalized download speed for 2.4 GHz speed tests. 0.4 for 5 GHz tests. Higher subscription tiers amplifies this difference. For example, in Tier 6, median normalized download speed for 5 GHz tests is 0.25, more than six times that of 2.4 GHz measurements (0.04).

WiFi RSSI: only consider 5 GHz tests. median normalized download speed for < -70 dBm is 0.2; 0.3 for tests with -70 dBm to -50 dBm; 0.49 for tests with -50 dBm to -30 dBm; and 0.52 for tests with > -30 dBm.

Available Kernel Memory: only consider 5 GHz tests with RSSI > -50 dBm. Median normalized download speed for devices with < 2GB memory is significantly worse than the other groups. Speed test performance can be greatly impacted by when available memory is not great enough.

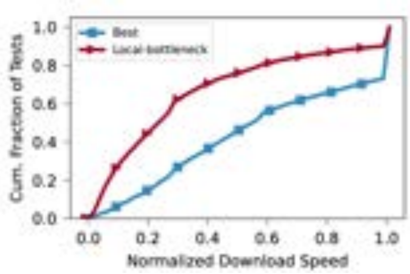

Platform	Type	Tier 1-3		Tier 5		Tier 6	
		#Measurements	Mean	#Measurements	Mean	#Measurements	Mean
Ookla	Android-App	8,890	5.25	2,810	17.04	5,512	40.23
	iOS-App	33,265	5.3	9,530	16.71	19,480	39.82
	Desktop WiFi-App	4,551	5.54	3,638	16.82	1,750	39.92
	Desktop Ethernet-App	1,031	5.69	1,400	16.95	2,098	40.13
	Net-Web	43,833	5.72	29,157	16.69	15,797	40.06
M-Lab	NDT-Web	70,789	5.32	16,417	16.71	9,490	39.94

Conclusions:

1. Subscription tier, device and network characteristics significantly impact Speed test performance, sometimes as much as seven-fold.
2. Speed test data must include meta-data context so that results can be properly interpreted.
3. Speed test vendors such design tests that maximize throughput.

This work was funded through National Science Foundation RAPID award NSF-2033946.

womENcourage™
Responsible computing for gender equality
MADRID - 2024

11th ACM Celebration of Women in Computing
womEncourageTM
Madrid, Spain,
26-28 June, 2024

acm