M-Lab: User initiated Internet data for the research community

Phillipa Gill
Google, USA

Christophe Diot
Google, USA

Lai Yi Ohlsen
Code for Science and Society

Matt Mathis
Google, USA

Stephen Soltesz
Google, USA

Corresponding author: Phillipa Gill (phillipagill@google.com)

This article is an editorial note submitted to CCR. It has NOT been peer reviewed.
The authors take full responsibility for this article's technical content. Comments can be posted through CCR Online.

ABSTRACT

Measurement Lab (M-Lab) is an open, distributed server platform on which researchers have deployed measurement tools. Its mission is to measure the Internet, save the data and make it universally accessible and useful. This paper serves as an update on the M-Lab platform 10+ years after its initial introduction to the research community [5]. Here, we detail the current state of the M-Lab distributed platform, highlights existing measurements/data available on the platform, and describes opportunities for further engagement between the networking research community and the platform.

CCS CONCEPTS

• **Networks** → **Network measurement**;

KEYWORDS

Network measurement

1 INTRODUCTION

M-Lab was founded in 2009 by New America's Open Technology Institute, the PlanetLab Consortium, Google, and a group of academic researchers. Spurred by conversations with Internet researchers that identified a lack of widely-deployed, well-connected servers for active measurements and challenges with sharing large datasets, M-Lab was founded to address these challenges and promote large-scale open source measurement of the Internet [25]. In 2019, M-Lab became a fiscally sponsored project of Code for Science and Society, a US-based nonprofit supporting open collaboration in public interest technology [3].

Since its founding, M-Lab has grown into a large open source project with contributors from civil society organizations, educational institutions, and private sector companies. The project is dedicated to: (1) Providing an open, verifiable measurement platform for global network performance, (2) Hosting the largest open Internet performance dataset on the planet and (3) Creating visualizations and tools to help people make sense of Internet performance. All the data collected by the M-Lab platform is openly available and all measurement tools hosted on the platform are open source.

M-Lab offers two key services to the network measurement community: (1) Hosting server endpoints for different active measurement projects and (2) Serving open data from experiments running on the platform. In the following sections we describe the M-Lab infrastructure and how it can be used by active measurements (Section 2). We then discuss the data currently being collected by the platform, which consists of network performance measurements to hosts around the world collected starting in 2009 (Section 3). Finally, we wrap up by discussing some open research problems and opportunities for collaboration between M-Lab and the broader research community (Section 4).

2 M-LAB FOR ACTIVE MEASUREMENTS

2.1 The M-Lab Platform

M-Lab physical infrastructure. Over the last decade, M-Lab has grown from a total of 16 sites across 8 countries in 2010 to 210 sites across 47 countries in 2021. Each site (also referred to as "pods") consists of four servers and a switch, generally connected at 10 Gb/s to a transit provider [18]. Each of the servers is a server-class machine with three of the servers hosting active experiments and the fourth being used for experimental traffic (eg., testing OS/measurement server updates). M-Lab distinguishes itself from other shared measurement platforms like PlanetLab [28] and RIPE Atlas [8, 29] by using modern Linux kernels to enforce stricter namespace isolation, limiting the number of experiments that run on the platform, and monitoring switch discard information and node resource consumption to identify potential cases of resource contention. See [22] for more details about the M-Lab platform specifications.

M-Lab network connectivity. The M-Lab pods are connected to the Internet via existing transit providers with a focus on ISPs that do not have significant connectivity to Internet users (ie., avoiding "eyeball" networks). This connectivity strategy stems from M-Lab's initial goal of ensuring measurements from an end user will cross at least one interdomain link (as opposed to measuring within a single ISP's network). These pods are connected to 70 ASes in 44 different metropolitan areas. See [22] for more detail and specifics about the location of each pod. M-Lab provides geographic load balancing between servers with the locate service [17] directing clients towards the M-Lab pod that is geographically closest to them that has enough available resources. See [17] for more information about how M-Lab directs clients towards a server for testing. Additionally, M-Lab supports measurements that override this load balancing and direct measurements towards a specified metro area (subject to resource availability).

M-Lab software stack. Since 2010, M-Lab's software stack has also evolved. 2019 saw the transition of M-Lab from using a PlanetLab-compatible resource "slicing" approach, to a more modern approach for deploying and managing experiments. The M-Lab servers are administered and maintained as a managed Kubernetes cluster. Experiments are deployed to servers using Docker containers. This approach is similar to what is currently used by the Edgenet project [6] which may provide an opportunity for prototyping M-Lab experiments before deploying on the M-Lab platform. See [19] for details about building and configuring measurement services for M-Lab.

Supported experiments. Unlike more general shared network measurement infrastructures, M-Lab focuses on providing server side resources for client-initiated active measurements. This means that the servers will only measure/record data upon receipt of a connection from an external user. For measurements that require probing arbitrary endpoints on the Internet, existing platforms such as RIPE Atlas [29] may be better suited. However, for experiments needing server endpoints for bandwidth intensive client initiated tests (*e.g.*, network speed tests, tests of network neutrality) M-Lab offers a well provisioned and reliable platform. One requirement for experiments run on M-Lab is that measurement software and data are open source [16]. Indeed, many experiments are specified such that a M-Lab enables the transfer of data from M-Lab servers and provides hosting for datasets produced by experiments running on the platform.

2.2 Measurements currently running on M-Lab

M-Lab has hosted a variety of measurements over the years [5] and currently hosts the following four measurement services:

Network Diagnostic Tool (NDT) [26]. NDT is the most popular measurement service running on M-Lab at this point. It tests a user's network connectivity by downloading/uploading an object via a WebSocket over TLS. It aims to capture how fast a user's device can send/receive data using a single TCP connection (bulk transport as defined in IETF's RFC 3148 [12]). Earlier versions of NDT used TCP Reno or Cubic with the latest version (NDT7) using BBR if available. Details about NDT's evolution over the last 10+ years can be found in [27].

Neubot DASH [1]. This experiment aims to understand streaming performance using Dynamic Adaptive Streaming over HTTP (DASH) [32]. It accomplishes this by emulating a streaming video player. More details can be found in [1].

Reverse Traceroute [9] The Reverse Traceroute project uses M-Lab to perform measurements needed to enable Reverse Traceroute to measure the network path back to a user from a selected endpoint.

Wehe [10]. Wehe exchanges network traffic between a user's device and the M-Lab servers to measure whether networks perform throttling of specific Internet applications based on their content. It tests a range of popular applications using client software that can be run from the Linux command line or in iOS and Android apps.

3 DATA AVAILABLE FROM M-LAB

A key principle of M-Lab is enabling research by making measurement data available to the network measurement community. Raw data from each measurement on M-Lab is stored in Google Cloud Platform with some measurements also having data imported into BigQuery for ease of use.

Country	# of tests	% of tests
USA	22,720,053	26%
India	15,564,933	18%
Brazil	6,372,670	7%
UK	5,304,506	6%
Philippines	4,847,981	6%

Table 1: Top countries in terms of number of M-Lab tests run in November 2021.

3.1 NDT data

Running since 2009, NDT represents the longest-lived and largest dataset made available by the platform. Data from NDT tests is available in BigQuery to make working with the data easier for external researchers [20]. The raw data for each test is available in the BigQuery tables, as well as a filtered dataset which aims to include only tests that did not experience a problem/error (details of this filtering can be found in [20]). Each test is annotated with server and client geolocation and network bandwidth, loss and latency.

Currently the platform is seeing an average of 2.9 million NDT tests per day (average for Nov. 2021) from 239 countries (for all Nov. 2021) with the vast majority of tests (90%) coming from the Internet Speed Test which integrates NDT into Google's Internet Speed Test functionality. The top countries in terms of number of tests are summarized in Table 1.

3.2 Core services and platform data

In addition to data collected directly by measurement services running on M-Lab, the platform also collects metadata about each TCP connection (except in cases where measurements have opted out of this data collection). This metadata includes:

- **Packet headers [21]**: a capture of packet headers for all incoming TCP flows during a measurement published as a per-stream .pcap file.
- **TCP_INFO [23]:** Statistics about TCP connections on the platform using tcp-info.
- **Traceroute [24]**: M-Lab uses Scamper [11] to traceroute from the M-Lab server towards the client that initiated a measurement (ie., opened a TCP connection) on the server.

We also publish a variety of metadata including switch discard rates to help identify tests that occurred when the system was under high load [15], BGP prefix/AS info [30] and MaxMind geolocation data for IP addresses in our data [13].

4 WHAT'S NEXT FOR M-LAB?

M-Lab offers a large spectrum of research opportunities to use data science and machine learning to better diagnose network issues. Here we overview potential research directions and opportunities for collaboration with M-Lab.

4.1 Understanding M-Lab data bias

There has been a lot of debate about potential biases in M-Lab data (eg., [2]). As with all data, it is very difficult to generalize any observation and every detail about how the data is collected is crucial in the interpretation of the results.

Server bias. M-Lab servers are chosen to be located in transit ISPs to measure paths that contain at least one interdomain link. This can result in paths that are longer than those seen by users when they access popular content which may be located in edge-caches or locations that are generally closer to the user. It is an open question to understand how the paths measured by M-Lab compare to paths towards popular online content or cloud hosting locations and how our choice of transit providers impacts conclusions drawn from the data.

Client bias. There could be other biases coming from routing or network conditions. Notably, users tend to run speed tests when there is a problem with the network. The open nature of the NDT protocol also means that there are clients incorporated into a variety of tools (*e.g.,* some measurement tools, Google's Internet speed test). It's important to avoid bias towards clients that may be automated and run measurements on a regular basis (or to center analyses on these clients when it makes sense). Machine learning to identify features of different clients in terms of when they test and how often they test can help illuminate these different user populations to better understand how they impact data from the platform.

4.2 Diagnostics for network users and ISPs

Currently the bulk of M-Lab's data comes from users running the Network Diagnostic Tool (NDT). However, the tool simply reports the download and upload bandwidth observed in the tests. Users are often running speed tests when they are dissatisfied with their network performance and this information does not give them the tools to understand how or why their Internet is not performing well. Right now NDT collects a variety of metrics (packet captures, traceroutes, latency, loss) that could be used to help diagnose where the bottleneck on a user's connection may be. The NDT data set also presents a unique opportunity to correlate tests from users in a single ISP to determine if the user is testing from an ISP with a widespread issue vs. having a problem with their individual connection.

Combining these data signals to produce information that is useful to users is an open research challenge. Specifically users may want to know: what is wrong with the network? Is the network fast enough to do a video call? Is there a way to improve the network performance? One could imagine applying machine learning or statistical inference to model the relationship between given speed test results and specific application-level metrics to help answer these questions. There is also the opportunity to develop novel network diagnostic tools to run as new experiments on the platform to help end users analyze and debug their network connections.

4.3 Network neutrality and accessibility

Recent years have seen heated debates about users' ability to access the Internet. Specifically, there has been a lot of discussion about the concept of network neutrality (ie., that ISPs should not provide differing network service based on the content of a given network flow). M-Lab is actively engaged in gathering data with the Wehe measurement service [10]. While data is being gathered, there is opportunity to analyze data being reported back by this service to identify new trends in traffic differentiation around the globe or to develop interfaces that allow the use of this data by policy makers.

M-Lab also provides an opportunity to perform new measurements of online information controls or censorship. Specifically, our servers can be used for different types of connectivity tests that benefit from a controlled server to test against. For example, one could implement transparent proxy detection [31] by sending specially crafted HTTP requests towards our servers. There is also the opportunity to measure specific types of censorship against a controlled server (*e.g.,* fuzzing to detect specific keyword blocking [4], or altering the content/URI of the measurement service to understand what about a connection will trigger censorship).

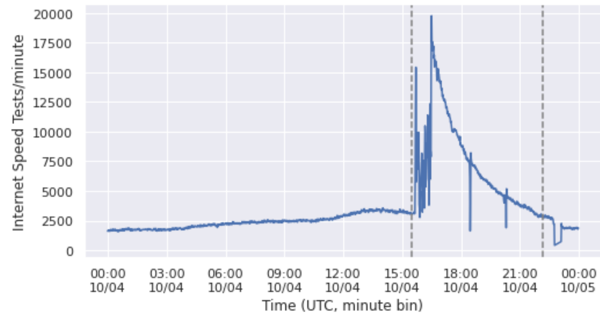

Figure 1: Internet Speed Tests per Minute During the 10/4 Facebook Outage

4.4 Network reliability

One interesting aspect of the NDT data collected by M-Lab is that users tend to test their Internet when something is wrong. This makes NDT test volume an interesting indirect signal of network problems. An example of this for the Facebook outage on October 4, 2021 is shown in Figure 1. Finding anomalies in the test volume of NDT and identifying how these anomalies relate to Internet outages is one interesting avenue for study. The NDT data itself may also be useful for identifying time periods where multiple clients or networks experience poor performance. By performing correlations and statistical inference it may be possible to develop an alerting system that publishes events of interest by mining the NDT data for times when multiple clients or networks experience poor performance.

4.5 Open governance

M-Lab was founded on principles of openness and transparency. These ideals are reflected in our data and code but are also meant to be embodied through the structure of our governance. Starting with the project's transition to Code for Science Society in 2019, M-Lab has been moving towards a governance structure that enables open and transparent decision making. In the next year and more we will build working groups, processes and decision making frameworks that enable the project to be guided by input from our vast, diverse community of stakeholders. Open governance is key to ensuring that M-Lab acts in the public interest and will help us validate that our output is useful, relevant and accessible.

4.6 Opportunities for collaboration

M-Lab is eager to engage academics and collaborate on solving network measurement problems using the platform and its data.

Funding. Fall 2021 has seen the first round of M-Lab fellowships to fund short term projects (6 months) related to M-Lab. The Google Research Scholar program [7] is another opportunity for researchers working on research related to M-Lab that is relevant to Google.

New measurement services. At the time of writing, M-Lab's has the capacity to host up to eight additional measurement services on our global platform. We welcome proposals for new open-source measurement methodologies which will extend the current suite's capabilities and help improve the Internet for end users.

Community engagement. M-Lab currently hosts regular community calls on the fourth Wednesday of the month [14]. These calls are an opportunity to discuss research and topics relevant to M-Lab and the community of researchers using the platform.

M-Lab blog. M-Lab is interested in hearing about what researchers are doing with M-Lab data! If you have used M-Lab data in your publication and would like to write a guest post summarizing your findings for the wider M-Lab community we are keen to hear from you.

ACKNOWLEDGMENTS

M-Lab – the platform, dataset, and community – is the culmination of over a decade of effort. We did not build it alone, and if not for the dedication of those before us, none of what we have described here had to be. Thank you to those we've work closely with over the years who brought the fight, prioritized kindness, kept things working, and provided endless supported, roughly chronologically, Tiziana Refice, Dominic Hamon, Josh King, Jordan McCarthy, Stephanie Alarcon, Meredith Whittaker, Josh Bailey, Michael Lynch, Mike Wawrzoniak, Peter Boothe, Ya Chang, Mike Millner. And, thank you to the Experiment Review Committee members for their experience and guidance curating the services that shape the project Jeff W. Boote, Georgia Bullen, Vint Cerf, Nick Feamster, Krishna Gummadi, Aleksandar Kuzmanovic, Sascha Meinrath, Larry Peterson, and Stephen Stuart. Finally, we thank all the supporting partners, client integrators, analysts, developers, and users who run measurements every day around the world. Without your help, curiosity, and proof, we could not fulfill our shared mission to measure the Internet, save the data, and make it universally accessible and useful.

REFERENCES

[1] S. Basso, A. Servetti, E. Masal, and J. C. D. Martin. Measuring DASH streaming performance from the end users perspective using neubot. In *MMSys: Proceedings of the 5th ACM Multimedia Systems Conference*, March 2014.

[2] D. D. Clark and S. Wedeman. Measurement, meaning and purpose: Exploring the M-Lab NDT dataset. In *TPRC: The research conference on communications, information and Internet policy*, 2021.

[3] Code for Science & Society. Code for science & society. https://codeforscience.org/.

[4] J. Crandall, D. Zinn, M. Byrd, E. Barr, and R. East. ConceptDoppler: A Weather Tracker for Internet Censorship. In *ACM Conference on Computer and Communications Security (CCS)*.

[5] C. Dovrolis, K. Gummadi, A. Kuzmanovic, and S. D. Meinrath. Measurement Lab: Overview and an Invitation to the Research Community. In *ACM SIGCOMM Computer Communication Review*, July 2010.

[6] EdgeNet. EdgeNet: The Globally Distributed Edge Cloud for Internet Researchers. https://www.edge-net.org/.

[7] Google Research. Research scholar program. https://research.google/outreach/research-scholar-program/.

[8] T. Holterbach, C. Pelsser, R. Bush, and L. Vanbever. Quantifying interference between measurements on the RIPE Atlas platform. In *ACM Internet Measurement Conference*, 2015.

[9] E. Katz-Bassett, H. V. Madhyastha, V. K. Adhikari, C. Scott, J. Sherry, P. van Wesep, T. Anderson, and A. Krishnamurthy. Reverse traceroute. In *USENIX NSDI*, 2010.

[10] F. Li, A. A. Niaki, D. Choffnes, P. Gill, and A. Mislove. A large-scale analysis of deployed traffic differentiation practices. In *ACM SIGCOMM*, August 2019.

[11] M. Luckie. Scamper: A scalable and extensible packet prober for active measurement of the internet. In *ACM SIGCOMM Conference on Internet Measurement*, 2010.

[12] M. Mathis. A framework for defining empirical bulk transfer capacity metrics. https://datatracker.ietf.org/doc/html/rfc3148.

[13] Maxmind. Maxmind GeoIP. https://www.maxmind.com.

[14] Measurement Lab. 2021 community call schedule. https://www.measurementlab.net/blog/community-call-schedule/#2021-community-call-schedule.

[15] Measurement Lab. DISCOv2. https://github.com/m-lab/disco.

[16] Measurement Lab. Experiment code & resource requirements. https://www.measurementlab.net/experimenter-mou/#experiment-code--resource-requirements.

[17] Measurement Lab. Locate API Usage. https://github.com/m-lab/locate/blob/master/USAGE.md.

[18] Measurement Lab. M-lab - conceptual & technical scope & policies. https://www.measurementlab.net/mlab-scope/.

[19] Measurement Lab. M-lab experiment developer responsibilities, requirements, and guidelines. https://www.measurementlab.net/experimenter-requirements-guidelines/.

[20] Measurement Lab. NDT (Network Diagnostic Tool). https://www.measurementlab.net/tests/ndt/.

[21] Measurement Lab. Packet Header Service (PCAP). https://www.measurementlab.net/tests/pcap/.

[22] Measurement Lab. siteinfo: M-lab public site information. https://github.com/m-lab/siteinfo.

[23] Measurement Lab. TCP INFO. https://www.measurementlab.net/tests/tcp-info/.

[24] Measurement Lab. Traceroute. https://www.measurementlab.net/tests/traceroute/.

[25] Measurement Lab. What is Measurement Lab? https://www.measurementlab.net/about/.

[26] Network Diagnostic Tool. Network diagnostic tool (ndt). https://github.com/ndt-project/ndt.

[27] L. Y. Ohlsen, M. Mathis, and S. Soltesz. Evolution of NDT. In *M-Lab Blog*.

[28] Planet Lab. Planet lab: An open platform for developing, deploying, and accessing planetary-scale services. https://planetlab.cs.princeton.edu/.

[29] RIPE. RIPE Atlas. https://atlas.ripe.net/.

[30] Routeviews. University of Oregon Route Views Project. http://www.routeviews.org/routeviews/.

[31] N. Weaver, C. Kreibich, M. Dam, and V. Paxson. Here be Web proxies. In *Passive and Active Measurement (PAM)*, 2014.

[32] Wikipedia. Dynamic adaptive streaming over http. https://en.wikipedia.org/wiki/Dynamic_Adaptive_Streaming_over_HTTP.