

User-Centered Network Measurement

Mark Claypool, Robert Kinicki, Craig Wills
Computer Science Department at Worcester Polytechnic Institute
100 Institute Road, Worcester, MA, 01609, USA
{claypool,rek,cew}@cs.wpi.edu

ABSTRACT

With the dramatic growth in Internet access from residences and out in public, traditional network measurements increasingly exclude performance behavior seen by the majority of Internet users. Our position is that a new user-centered network measurement platform is needed to include the burgeoning numbers of largely invisible Internet users. This work defines a network measurement platform taxonomy to motivate a three-pronged network measurement platform that provides flexibility, user incentives and reduced impediments for user participation. This continuum of options, along with application-oriented experiments, is needed for an effective platform. Insight obtained from this platform is critical to inform future Internet design choices.

1. INTRODUCTION

A *network measurement platform* is a system of Internet nodes and software designed to gather network performance data. Traditionally, Internet measurement has been done from points in the network infrastructure or from research labs and universities. However, with the dramatic growth in Internet access from residences and out in public, often hidden behind NATs, the old measurement paradigm increasingly excludes the performance vantage point seen by the majority of Internet users. The size of this cadre of “invisible” Internet users will only increase as public wireless networking becomes more commonplace and home networking spreads further through the developing world. While this user-centered perspective is important for today’s Internet, our position is that providing user-centered network measurement capabilities is critical for next-generation network design and experimentation such as with GENI [4].

The need for network measurement approaches centered on where users live and their specific interactions with the Internet has already been recognized. Desirable outcomes from a recent NSF Computing Infrastructure session on testing for the new Internet [11] include better representation of the user population, running non-Linux and a “SETI@home” type mechanism for networking. Previous work [1] laments the widening gap between measurements for the visible and largely invisible portions of the Internet community.

While existing network measurement platforms have sev-

eral desirable features, they do not satisfy future needs. Platforms such as PlanetLab [10] and Archipelago [6] provide researchers flexibility in choosing metrics to collect, but their platform nodes are permanent, immobile and within a dedicated infrastructure. Alternative platforms such as NETI@home [13], DIMES [12] and DipZoom [17] allow measurements from any node in the Internet, but the scope of their measurements are limited with little incentive for the general populace to participate. Finally, Gomez [5] and a variety of “speedtest” services [3, 14] include incentives for user participation, but they are not designed for the network researcher. Due to these platform limitations, previous published network performance studies collecting data from the home or public have typically employed ad hoc “ask my friends” approaches to participation that yield measurements from a limited number of vantage points with little indication of whether these participants are representative of the population at large.

To highlight relative strengths and weaknesses of existing platforms, this paper introduces a network measurement platform taxonomy and then use this taxonomy to motivate an alternative approach for *flexible, user-centered* network measurements. Rather than apply a single interface for network measurement, the proposed platform encompasses three integrated measurement approaches: 1) The first features measurement flexibility from dedicated hardware via experiment modules that are dynamically downloaded and easily changed. 2) The second executes application-oriented performance tests on user hardware. The application-oriented nature of these tests provides performance data that can be uploaded for sharing among network researchers while providing valuable user feedback on network applications of interest. This feedback is a strong incentive for user participation. 3) The third executes within a user Web browser via Flash or Java to provide a low impediment choice to encourage wide user participation while still yielding useful Internet performance data to be shared.

The distinctions of this three-pronged network measurement platform is the realization that flexibility, user incentives and reduced impediments are critical properties if future measurement platforms are to include the burgeoning number of invisible Internet users. However, providing these

properties involve trade-offs such that including each within the same measurement vehicle is difficult. Hence, our position is that a continuum of options that includes application-oriented tests are necessary components of any network measurement platform intended to assess future design changes in the Internet.

The remainder of this paper is organized as follows: Section 2 defines a taxonomy to characterize and classify existing network measurement platforms; Section 3 presents the architecture of our network measurement platform within the context of the measurement platform taxonomy; Section 4 describes measurement experiments that are possible with this platform; and Section 5 concludes with a summary of our work.

2. NETWORK MEASUREMENT PLATFORM TAXONOMY

This section defines a taxonomy characterizing current platforms, then uses this taxonomy to examine three broad classes of existing network measurement platforms, with representative measurement platforms selected as examples. The taxonomy contains three entities classified along multiple dimensions. The entities and their dimensions are:

1. *Measurement Points (MPs)* are nodes from which measurements are taken. These nodes may be users' machines or computers dedicated to measurements. MPs are classified along a number of dimensions:
 - (a) *Source*. If any user machine connected to the Internet can be an MP then the platform source is *open*. If equipment or specialized software must be provided (with permission) by a user then this is a *restricted* source. Otherwise, it is *closed*.
 - (b) *Location*. If the MP is at a university or a corporation, the location is *commercial*. If the MP is at a residence then it is *home*.
 - (c) *Incentive*. User incentives can include the ability to *access* MPs as a Client (see below) or data as an Analyst (see below). Alternately, incentives for the users can be more direct such as *monetary* compensation or *feedback* via performance data. If the incentive is only voluntary then there is *none*.
 - (d) *Impediment*. User impediments for setting up an MP include contributing *equipment* or installing *software* locally. It is assumed all MPs impose some resource demands, typically during active measurements.
 - (e) *Metric*. The specific performance metrics measured by the platform may be *fixed* or the performance data gathered may be *flexible*.
2. *Clients* are users of the platform that can run experiments and take measurements. If any user can be a

client, the platform is *open*. If special permission is required or MP impediments must be met, it is *restricted*. If it is not possible for any user to become a client, the platform is *closed*.

3. *Analysts* are users of the data obtained from the measurements. If MP data is accessible by any user then the platform is *open* to analysts. If only part of the data is available or access is limited to the client taking the measurement, the platform is *restricted*. Otherwise the platform is *closed* to analysts.

2.1 Dedicated Infrastructure Platforms

The first two rows of Table 1 represent the class of platforms that operate from a dedicated infrastructure. PlanetLab [10], which follows NIMI [9], utilizes a permanent infrastructure designed to employ wide area network services. PlanetLab provides a flexible network measurement environment, but access to PlanetLab requires contribution of equipment that restricts Clients and Analysts. The permanent nature of these nodes means they are located in commercial settings.

Archipelago [6] provides open access to Internet latency and topology data collected from a restricted set of MPs. Archipelago plans to provide more data collection flexibility by adding a sandboxed platform for vetted experiments from collaborators using a model similar to Scriptroute [15].

The distinguishing characteristics of this measurement platform class are measurement flexibility and MP restrictions that severely limit the vantage points from which measurements can be obtained.

2.2 Open Measurement Platforms

Measurement platforms in this class are primarily characterized by the lack of limitations on which nodes can become Measurement Points. These open measurement platforms, shown in the middle three rows of Table 1, can be deployed in both commercial and home settings. All of these platforms require software installations on the participating MP. The platforms differ in performance data measured and the level of data sharing allowed. The NETI@home [13] platform collects TCP/UDP usage data, but without data sharing and the only user incentive is to "help make the Internet a better place." The DIMES [12] software collects Internet topology information, which is shared with Analysts, but again there are no specific incentives for user participation. DipZoom [17] takes a peer-to-peer approach to measurement with software that serves as both a Measurement Point and allows Client access to a fixed set of network measurements such as ping and nslookup. As with the other platforms in this category, there is no incentive for participation by the general populace.

2.3 Measurement Platforms with Incentives

The final group of measurement platforms specifically address the need for user incentives. These platforms, shown

Table 1: Taxonomy of Current Network Measurement Platforms

Name	Measurement Point					Client	Analyst
	Source	Location	Incentive	Impediment	Metric		
PlanetLab	restricted	commercial	access	equipment	flexible	restricted	restricted
Archipelago	restricted	commercial	access	equipment	fixed (latency, topo)	restricted	open
NETI@home	open	commercial & home	none	software	fixed (TCP/UDP)	closed	closed
DIMES	open	commercial & home	none	software	fixed (topology)	closed	open
DipZoom	open	commercial & home	access MPs	software	fixed (varied)	restricted	restricted
Gomez	open	commercial & home	monetary	software	fixed (Web)	closed	closed
Speedtest Services	open	commercial & home	feedback	Flash, Java	fixed (latency, t'put)	closed	restricted

in the last two rows of Table 1, provide clear reasons for user participation. Gomez [5] is a company employing monetary incentives for users to allow their machines to serve as MPs to collect Web-based metrics. However, Clients or Analysts are unable to access this data. A number of “speedtest” services [3, 14] allow users to obtain latency and throughput metrics from their own machines to Internet servers. These mechanisms are particularly appealing to users as they require no software installation, only execution of Flash or Java by users’ browsers. Users obtain feedback about their machines’ performance and most services provide a more regional/global performance perspective.

3. NETWORK MEASUREMENT PLATFORM FOR THE HOME

Our goal is a flexible platform for public and home networking measurements with low participation barriers and attractive incentives to encourage a broad-based user participation. To meet these objectives, we employ three techniques within the continuum of tradeoffs between encouraging participation and providing flexible measurements.

Table 2 highlights the distinguishing features of the three components of the proposed measurement platform. All three techniques are open to measurements initiated from any machine with Internet access. The first technique, *Wireless Home Internetworking with Mobility* (WHIM), expects wireless connectivity and mobility in the MP location. A WHIM MP is a dedicated machine, such as a laptop, running experiments developed by clients (researchers) of the platform. WHIM provides great flexibility in client measurement experiments, but with relatively high user participation barriers.

In contrast, the other two measurement techniques encourage user participation by lowering impediments and pro-

viding incentives. Accessed by users via the *HowMyNetwork.com* Web site, abbreviated as *HMN*, the intent is to attract users by providing feedback on their network performance for applications of specific interest to them, e.g. on-line computer games. After a HMN Executable (Exe) is downloaded and executed, the user selects applications for HMN to evaluate. HMN then performs the appropriate measurement tests and provides performance feedback specific to the user’s MP. Similarly, HMN Flash/Java is a Flash program or Java applet that executes user-centered tests within limitations imposed by executable environments. Although HMN Exe yields more information from the MP than HMN Flash/Java, it requires a higher user trust level that will lower user participation. With both HMN choices, not only is user feedback provided, but network measurement data is sent to a central Internet repository.

The architecture where these three techniques fit together is shown in Figure 1. The ovals depict the main entities of a measurement platform (see Section 2) and the dashed rectangle in the center encompasses the measurement platform. Clients (researchers) provide experiment modules, which are appropriately tested and stored in a queue of WHIM modules. When an WHIM MP executes from a user node, it obtains these experiments and serially runs them with results reported back to the central platform for later retrieval by Clients. In contrast, HMN Flash (Java) and HMN Executable MPs run one experiment containing the set of measurements appropriate for their platform and provide feedback pertinent to user applications. Data collected during experiments is sent from the MPs to the repository for long-term storage. Repository data can be viewed by an Analyst after the experiments are run.

An important aspect of the proposed measurement platform is that these separate measurement techniques are not

Table 2: Taxonomy of Measurement Platforms for the Public and Home User

Name	Measurement Point					Client	Analyst
	Source	Locations	Incentives	Impediments	Metrics		
WHIM	open	commercial and home	access MPs	equipment	flexible	restricted	restricted
HMN Executable	open	commercial and home	feedback	software	fixed (appl focus)	restricted	open
HMN Flash/Java	open	commercial and home	feedback	Flash, Java	fixed	restricted	open

performed in isolation. Although the HMN Flash measurements have the broadest distribution, similar measurements can also be performed by WHIM MPs and HMN Executable MPs to compare against the baseline measurements from the HMN Flash MPs. The cross-technique approach allows calibration of WHIM and HMN Executable MPs with the larger performance data set. The following provides more detail on each measurement technique in our platform.

3.1 Flexible Network Measurement

In general, WHIM allows clients to run arbitrary experiments on a dedicated measurement point in a location with wireless access. This capability is achieved by having a dedicated WHIM laptop for the duration of the experiments. The expectation is that clients registering experiments would themselves arrange for WHIM nodes to be executing not only their own experiments, but also the experiments of other researchers. This arrangement is similar to PlanetLab or DipZoom where researchers must provide MPs to use the platform.

Upon startup, a WHIM MP contacts the WHIM server to obtain the experiments to be run. The WHIM MP runs the experiments serially, not simultaneously. Any resource management required, such as limitations on the length of an experiment or the bandwidth consumed, is enforced by the WHIM MP. The use of a dedicated machine and vetting of experiments reduces security and resource protection concerns.

In particular, a virtual machine (such as used in PlanetLab) is an effective model for running concurrent tests and for allowing remote access to measurement points (that may be running experiments). However, in the case of WHIM, the MPs are indicative of a typical home client that is generally not up and accessible. Thus, remote access is not needed. Additionally, tests are run serially, not in parallel, so virtual machine time sharing is not needed. This reduces potential overhead caused by a virtual machine that can impact timing measurements in latency sensitive experiments.

Currently under development, the WHIM MP runs on Linux which facilitates using a modified MadWifi¹ driver to gather 802.11 performance characteristics. In addition to conduct-

¹<http://madwifi.org/>

ing active measurements, this configuration allows WHIM to passively collect home network data. By configuring a WHIM laptop as a wireless sniffer [7], passively obtained information about the surrounding wireless network environment can be obtained at the wireless MAC level.

This WHIM component has a number of desirable properties. Executing flexible researcher experiments similar to PlanetLab, WHIM MPs can gather both active and passive measurements as well as run experiments for longer periods than the HMN components in the measurement platform. However, WHIM operates within a smaller deployment base of trusting users than the HMN techniques and currently only runs Linux-based experiments.

3.2 Incentive-Based Network Measurement

The HMN techniques are designed to provide broad-based network measurements from anywhere that users access the Internet. Incentives for these users are in the form of feedback on network performance specific to the users' applications of interest.

The HMN Exe technique executes workloads of interest to the typical residential user. For example, a home workload might include representative DNS queries, downloads of select Web pages, pings to a few online game servers, or multimedia streams from popular Internet video sites. When HMN users choose their applications such as VoIP, online games, Web, Instant Messaging or file-sharing for performance evaluation by the network measurement platform, their incentive is receiving timely performance statistics tailored to the chosen application.

A key challenge for the HMN executable development is gathering a rich texture of performance information from stand-alone programs executing on unprivileged user MPs running from a variety of operating systems while the user is waiting. Data such as IP addresses and local DNS server information can be obtained by sending requests to known servers and observing the source. Information on the end-host network connection type, such as DSL, Cable, 802.11 variant or cell phone protocol, may be obtained directly from the operating system or deduced by using techniques such as examining statistics on the dispersion of packets sent from a server [16].

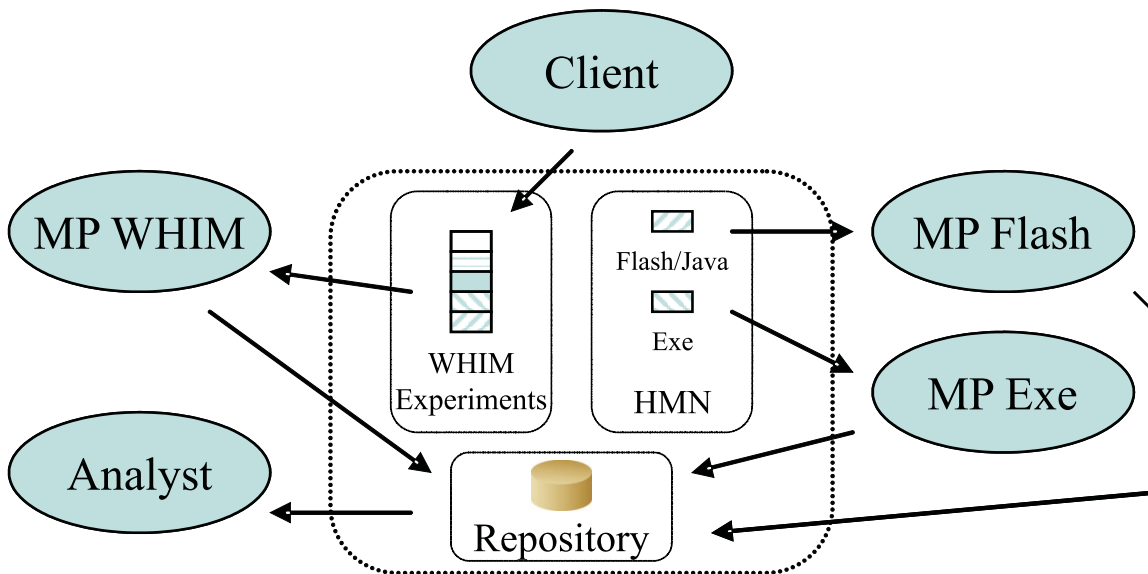


Figure 1: Architecture of Measurement Platform

The HMN Exe technique provides an intermediate point in the tradeoff between measurement flexibility and broadening participation by currently invisible Internet users. By supporting a variety of operating environments that specifically includes Windows, this method will yield more representative performance measurements than is available on existing measurement platforms. It is envisioned that this tool is most likely to be used when the user is experiencing degraded Internet performance which is also the scenario of most interest to network researchers. The trade-off with this approach is the barrier of having to download and install a HMN executable, which is a deterrent for some Internet users.

3.3 Easy Access Network Measurement

The third entity in the proposed measurement platform is designed to leverage the ease-of-use associated with running Flash or Java applications directly from the user's browser, which promises significantly higher user participation than the other two techniques. Similar to the HMN Exe scheme, the HMN Flash/Java approach gets preferences from individual users with respect to which applications to monitor prior to executing selective tests from the user's browser.

Focusing on ease of use associated with Flash and Java, this scheme operates under a more restrictive environment that limits the network performance characteristics that can be measured. Working from experience gathered by reviewing the set of currently available speedtest services, we have begun developing a HMN Flash/Java component that integrates effectively with the other two prongs of the proposed network measurement platform.

4. EXPERIMENTS

The key question regarding our multi-faceted network measurement platform is how well it works in supporting public- and home-based experiments for which it was designed. This section describes types of data from different network layers that can be obtained from our platform.

Our measurement platform supports a range of measurement opportunities. The WHIM nodes provide the most flexibility in allowing researchers to construct experiments to be executed on a dedicated platform. We see two types of such experiments for execution. The first are experiments to measure and characterize the nature of wireless network traffic from a variety of devices that represent typical residential user activity. There have been many characterization studies of wireless networks, but these have largely focused on campus or public wireless settings with little traffic data collected on wireless networks in the home. The NSF-supported repository CRAWDAD, a Community Resource for Archiving Wireless Data at Dartmouth [2], which has only one data set of three houses from a wireless home network and this data set includes only UDP and TCP throughput measurements [8]. Our WHIM platform allows the collection of such data from a number of vantage points.

The second type of experiments available for the WHIM portion of the platform are active measurement experiments that can now be executed in a representative user setting. Up until this point, these type of flexible experiments have been performed in testbeds such as PlanetLab, whose nodes are not representative of conditions experienced by home users, or performed in an ad hoc manner by friends of a researcher. The WHIM portion of the platform allows a diverse set of experiments to be constructed and through the

common queue for experiment modules they can be executed over a broader set of domains than can be attained by a single researcher.

The HMN portion of the platform does not provide the same level of flexibility, but affords a scale of participation that previous measurement platforms have not been able to provide. The HMN Executable technique allows a range of network measurements that can be associated with application-level performance as well as provide meaningful feedback to the user. Application-oriented measurements for games, Web, VoIP, DNS and streaming media also provide lower-layer measurements on latency, loss, throughput and available bandwidth. In addition, a key aspect of this technique is that it not only performs measurements, but is designed with an interface that can obtain input from users on applications and even specific content that is of most interest to them. This type of user input not only allows feedback to be customized, but serves as valuable data on what users are doing. The HMN Flash technique provides a similar purpose as the Executable technique, albeit with even less flexibility in what type of measurements can be performed, but with a lower barrier for use allowing for even wider participation.

5. SUMMARY

Current and future network research is hampered by a true understanding of what has been termed the invisible portion of the Internet. Other platforms have aspects of the flexibility, openness, incentives and lack of impediments that our platform embodies, but none provide a continuum of measurement tradeoffs.

This work proposes a multi-pronged network measurement platform that provides a basis to perform measurements that are needed and currently not possible. Our platform provides three primary advantages in comparison to existing measurement platforms:

1. The platform provides for representative *public* and *home* measurements. This environment typically involves machines running a variant of Windows over wireless LANs with broadband access to the Internet. Our work is novel in targeting this crucial environment for wide-scale network measurement.
2. The platform retains flexibility in the experiments that can be performed on dedicated hardware via a software structure that supports a wide range of researcher-developed experiments.
3. The platform provides incentives for user participation. Incentives are in the form of user-centric feedback on application performance users care about. A crucial missing element in previous measurement platforms targeting home users is the lack of incentives. Users will only participate on a large scale if they perceive benefit for themselves.

The three portions of our network measurement platform are currently being developed. Network measurement data and invitations to submit experiments will be made available to network researchers as portions of the platform become active.

6. REFERENCES

- [1] M. Casado, T. Garfinkel, W. Cui, V. Paxson, and S. Savage. Opportunistic Measurement: Extracting Insight from Spurious Traffic.
- [2] CRAWDAD: a community resource for archiving wireless data at Dartmouth.
<http://crawdad.cs.dartmouth.edu>.
- [3] Broadband reports.com speed test.
<http://www.dslreports.com/stest>.
- [4] Global Environment for Network Innovations.
<http://www.geni.net/>.
- [5] Gomez peer community. <http://www.porivo.com/>.
- [6] Y. Hyun. The Archipelago Measurement Infrastructure. In *Proceedings of the 7th CAIDA-WIDE Workshop*, November 2006.
- [7] F. Li, M. Li, R. Lu, H. Wu, M. Claypool, and R. Kinicki. Tools and Techniques for Measurement of IEEE 802.11 Wireless Networks. In *Proceedings of WiNMee*, Boston, MA, USA, Apr. 2006.
- [8] K. Papagiannaki, M. Yarvis, and S. Conner. Experimental Characterization of Home Wireless Networks and Design Implications. In *Proceedings of the IEEE Infocom*, Apr. 2006.
- [9] V. Paxson, A. Adams, and M. Mathis. Experiences with NIMI. In *Proceedings of Passive and Active Measurement (PAM)*, April 2000.
- [10] L. Peterson, A. Bavier, M. E. Fiuczynski, and S. Muir. Experiences Building PlanetLab.
- [11] H. Schulzrine. The New Internet – Goals, Testing and Infrastructural Needs. In *NSF Computing Research Infrastructure PI Mtg*, Boston, MA, June 2007.
- [12] Y. Shavitt and E. Shir. DIMES: Let the Internet Measure Itself, 2005.
<http://www.arxiv.org/abs/cs/0506099v1>.
- [13] C. R. Simpson, Jr., D. Reddy, and G. F. Riley. Empirical Models of TCP and UDP End-User Network Traffic From NETI@home Data Analysis.
- [14] Speedtest.net. <http://www.speedtest.net/>.
- [15] N. Spring, D. Wetherall, and T. Anderson. Scriptroute: A Public Internet Measurement Facility. In *USENIX Symposium on Internet Technology and Systems*, 2003.
- [16] W. Wei, B. Wang, C. Zhang, J. Kurose, and D. Towsley. Classification of Access Network Types: Ethernet, Wireless LAN, ADSL, Cable Modem or Dialup? In *Proceedings of IEEE Infocom*, Miami, FL, USA, 2005.
- [17] Z. Wen, S. Triukose, and M. Rabinovich. Facilitating Focused Internet Measurements.