

Community-Oriented Network Measurement Infrastructure (CONMI) Workshop Report

KC Claffy
San Diego Supercomputer
Center
University of California, San
Diego
La Jolla, CA, USA
kc@caida.org

Mark Crovella
Department of Computer
Science
Boston University
Boston, MA, USA
crovella@cs.bu.edu

Timur Friedman
Networks and Performance
Analysis
Laboratoire LiP6-CNRS
Paris, France
timur.friedman@lip6.fr

Colleen Shannon
San Diego Supercomputer
Center
University of California, San
Diego
La Jolla, CA, USA
cshannon@caida.org

Neil Spring
Department of Computer
Science
University of Maryland
College Park, MD, USA
nspring@cs.umd.edu

Abstract

This report summarizes issues discussed at the first CONMI workshop held on 30 March 2005 in Boston, Massachusetts. Sponsored by the National Science Foundation's Office of Cyberinfrastructure (OCI-0532233), the workshop was intended to begin a discussion regarding the viability and utility of a community-oriented network measurement infrastructure. This report was published 20 December 2005 online at: <http://www.caida.org/workshops/conmi/>.

Executive Summary

Up-to-date, relevant Internet research requires comprehensive network measurement [9, 5], but conducting and supporting Internet measurement raises several daunting challenges for the research community and funding agencies. Researchers need current data to progress in many areas, including Internet topology structure, routing dynamics, security, and workload trends. Given the inherent diversity of the Internet, collection of data requires a large-scale, distributed network measurement infrastructure. However, several challenges must be solved to enable large-scale measurement: funding of collection infrastructure, preserving the privacy of ISPs and users, resolving legal and proprietary ownership concerns, and prohibiting experiments that might cause harm.

The Community-Oriented Network Measurement Infrastructure (CONMI) workshop brought together key members of the Internet measurement research community to discuss whether a community-oriented approach could address current and near-future challenges in large scale measurement. Our inspiration came from the astronomy and high-energy physics communities which have self-organized to build, operate, and allocate the use of large, unique, and expensive measurement platforms. The objective of this workshop was to explore whether this cooperative model would benefit the Internet measurement community.

Internet measurements must respect the privacy of both users and of network providers. We explored the privacy implications of network measurement, with particular attention to facilities that would accept experiments to be run on infrastructure deployed on actively

used networks. Each experiment could be examined in advance through community mechanisms, e.g., a review panel, to ensure that the collection process was not harmful and that the results, if released, would not raise privacy concerns.

We also explored whether a fundamentally new, community-oriented model for passive measurement could enable a dramatically more powerful set of measurement experiments. The workshop raised more questions than it provided answers regarding how to best use passive measurement infrastructure funding to support the needs of the network research community, and we expect this discussion to continue as progress is made in critical areas, especially the difficulties in funding evolving measurement equipment and balancing privacy and security concerns with collection of necessary, relevant trace data.

In contrast, we felt the development of a community-based *active* measurement infrastructure seemed currently feasible. The primary considerations include how to coordinate measurement requests from a large community of researchers, how to ensure responsible use, and how to ensure integrity of the data if gathered by an unknown party. In both cases, some community-oriented program is likely to be necessary to accommodate as many needs of the community as possible as cost-effectively as possible. Given the limited funding available to invest in this kind of measurement infrastructure, an objective cost-benefit analysis of the payoffs is essential.

Categories and Subject Descriptors

C.4 [Performance of systems]: Measurement techniques. C.2.3 [Network Operations]: Network Monitoring. H.3.5 [Online Information Services]: Data sharing.

General Terms

measurement, experimentation

Keywords

community, data sharing, measurement, monitoring, Internet, meta-data, annotations, topology, workload, performance, routing

1. COMMUNITY-ORIENTED ACTIVE MEASUREMENT

Community-oriented active measurement is an attempt to collaborate in maintaining infrastructure, in reusing measurements, and in choosing lowest-possible cost, low-risk measurement designs. Yet to collaborate, the community must agree on which infrastructures to maintain, which measurements to collect, and what costs and risks are acceptable.

Cooperation is essential to active topology measurement because comprehensive topology measurement requires widespread infrastructure. Such an infrastructure, in turn, requires maintenance of hardware, software, and most importantly of continuity in trust relationships between researchers and (often, but not always, commercial) organizations that allow measurement servers to be hosted in interesting places.

Several infrastructures have supported or are currently supporting active measurement for research: hardware fully dedicated to one project, such as Skitter [8]; hardware shared with certain Acceptable Use Policies (AUPs) but more easily accessed, such as PlanetLab [20] or RON [4]; and single-project software on fully decentralized, multipurpose hardware, such as NETI@home [23] and DIMES [22]. Other infrastructures that previously supported limited active measurement are no longer funded: Surveyor [13], AMP [2], and NIMI [19]. The Skitter project also runs out of funding this year. Each of these projects has different costs, advantages, and weaknesses. There is no consensus on a single correct model for supporting active network measurement, although integration of some platforms as a substrate for a more comprehensive and uniform platform development to support a variety of measurement projects seems promising. Each can contribute to the *representative* aspect of measurements: the goal that our measurements of the network accurately reflect network properties, despite the limited size of the platform.

Obstacles prevent open collaboration in active measurement. First, cost is a significant factor. The expenditure of time and infrastructure maintenance to support one project often exhausts the capabilities of the group doing the measurement, leaving no time and engineering support for coordinating measurements with others or sharing the infrastructure. Second, network measurement data may be abused to harm the network because accurate network measurement data may help the unscrupulous to attack infrastructure.

Third, network probes themselves are unsolicited and may be seen as malicious. Active measurements, especially at a large scale, can cause harm to the network by consuming precious router processing time and by appearing to be malicious. Experiments involving many thousands of active probing hosts are being proposed, and today, can be conducted without oversight. We explored community mechanisms for overseeing large-scale distributed active measurement, because without community oversight, these experiments could easily go awry and cause undesirable results for the Internet as a whole. Because no barriers prevent such experiments, we discussed ways to facilitate running them safely, via both community review in advance and monitoring during execution.

Fourth, network measurement data can be seen as exposing private information about network architectures. Also, network measurements are often works-in-progress that can mischaracterize the structure of ISP networks. Using network measurements to generate bad press for ISPs or difficulty for ISP operators may lead quickly to countermeasures that deceive or block network measurement. Fifth, agreements allowing one researcher to run a specific experiment on a machine deployed at a remote site may not allow other researchers to run the same experiment or any researcher to run a new experi-

ment. Iteratively updating agreements to add a new person or project can be complex, time consuming, and difficult enough that the platform host, motivated only by altruism, loses interest. Finally, practical differences in platform hardware, operating system, and software can preclude measurements from diverse locations.

A coordinated, cooperative active measurement project could overcome these obstacles. Explicit funding and design can reduce the cost of supporting a wide variety of active measurements. Coordination offers the potential to be more efficient in packets sent, since results can be reused. Greater efficiency can lead to greater accuracy, either through expansion of the portion of the Internet that is measured or through the more frequent measurement of the same network. Directly addressing security concerns centrally, with an actively maintained do-not-probe list and prompt response to questions and complaints will reduce the likelihood that networks will close themselves to measurement. Standardized agreements and data distribution policies can limit the malicious use of the data and reassure ISPs and organizations hosting measurement platforms. Central management of active measurement infrastructure can result in standard access and configurations to simplify running measurements from many nodes. Thus a centrally coordinated active measurement platform could be a significant benefit to Internet research by increasing diversity and depth of measurements and by allowing significantly more researchers to perform active measurement studies.

1.1 Measurement Research on PlanetLab

PlanetLab is a widely-deployed network testbed designed and operated to support computer systems research. It allows development and deployment of new networked technologies in a controlled environment, incorporating realistic topologies and behavior. PlanetLab is also capable of supporting limited active network measurements, and must be considered in any discussion of development of community-oriented measurement infrastructure.

Although some researchers have successfully performed active measurement experiments on PlanetLab, others had trouble using PlanetLab because of CPU load, its academic bias, its limited resources, and its Acceptable Use Policy [24]. PlanetLab currently allocates processor time by slice (user) rather than by thread, a method friendlier to low-CPU-usage measurement experiments. However, other experiments running on PlanetLab nodes can interfere with active measurement projects; PlanetLab does not claim to be a substitute for dedicated resources. That PlanetLab sites are primarily academic raises concerns for many researchers that the connectivity to those sites is not representative of the commercial Internet. A recent paper by Banerjee et al. [6] describes how and how not to use an academic testbed like PlanetLab. Typical paths between PlanetLab nodes typically traverse research networks, while many active measurement projects seek to explore paths through commercial backbone links. Resources available to PlanetLab nodes at each site are limited; often bandwidth is capped, processor time is limited, and storage can be exhausted. While some active measurement studies are able to work within these constraints, others cannot, e.g., bandwidth estimation and spectroscopy studies. Finally, PlanetLab's Acceptable Use Policy can be a significant obstacle to active measurement research, as it explicitly forbids both systematic and random network scanning [21]:

- Do not do systematic or random port or address block scans.
- Do not spoof or sniff traffic.

While humans make exceptions and some experiments that violate this rule have occurred on PlanetLab, there is no infrastructure explicitly dedicated to supporting responsible, well-conceived active network measurement.

PlanetLab serves as a model of a centralized, shared infrastructure successfully promoting systems and applications research. PlanetLab's methodology for administration of machines, interactions with hosting sites, abuse reports, and support for user code execution on hundreds of machines around the world provides valuable operational expertise and a starting point for development of a community-oriented active measurement infrastructure. Indeed, such an infrastructure would complement PlanetLab and the use of both would allow novel research projects not otherwise possible.

1.2 Client-side Software Infrastructure

Distributed computation projects, inspired by SETI@home, have set out to use otherwise idle compute cycles of home machines to solve interesting scientific problems. One approach to increasing the number and representativeness of vantage points available for measurement is to use home machines as a platform: providing a downloadable tool that reports back information about network performance, topology, and workload. There is a strong justification that such a massive increase the number of vantage points is required by the Internet's current size—it is no longer conceivable to perform measurement from a few vantage points probing the network in the same way.

At the workshop, we briefly discussed the challenges faced by three nascent projects in this area: `traceroute@home` [3], `DIMES` [22], and `NETI@home`.

The `traceroute@home` project does not, itself, produce an artifact, but identifies and addresses research challenges of @home-style distributed network measurement. The motivation is clear: the diversity of vantage points made possible by a tool that can run on tens of thousands of Internet hosts may improve representativeness [7] and limit topology sampling bias [14, 1]. We discuss some of the challenges below.

First, what guidelines for responsible deployment would ensure that active measurement tools do not harm the infrastructure? Problems experienced by shared and dedicated infrastructures recur in this domain, but are more pressing. Intrusion detection alarms may fire in response to traffic, but abuse mail will likely go to a user's ISP rather than to the network operators of a research network, possibly leading the ISP to disconnect the user. High-rate traffic from thousands of sources may appear as a DDoS attack, already a problem for PlanetLab-hosted measurements and overlay applications. With tens of thousands of hosts, experiments could look too similar to attacks, without traceability and thus no point of contact for opting out of measurement as is possible with dedicated infrastructures. The high rate of traffic could generate significant costs for users who pay for bandwidth by usage. Finally, measurement traffic from different hosts, if uncoordinated, may interfere; while not harming the network, uncoordinated measurement harms the integrity of any research using the data.

One method for limiting the number of probes that reach a destination relies on network routing being deterministic and destination-based. If so, paths from different sources to the same destination, once they converge, will never separate: routers make next hop decisions independently of the source. The result is a tree of paths that converge as they near a destination. This assumption changes topology discovery from an exhaustive all-sources to all-destinations process to discovering the edges of each tree of paths rooted at every destination. By probing paths only until they intersect an already-probed part of the tree, the number of probes that reach the destination is minimized. This method is used by both Doubletree [10] and Scriptroute's [25] reverse path tree tool. But the community has

not yet solved the general problem of how to safely scale up active measurement techniques to thousands of nodes.

Second, @home-style measurement must verify the integrity of data collected from untrusted sources. Verifying results is a general problem, because honeypots might masquerade an entirely fictitious network for other infrastructures to measure. The problem is aggravated in @home-style measurement, because a source can invent erroneous data and exist behind many different interfering middle-boxes: transparent proxies, firewalls, or exceptional routing. When there is an incentive to manipulate the measurements, for example to skew AS coverage of a global Internet map, this risk becomes significant.

One approach is cross-validation with other, more trusted measurements from controlled infrastructure, an approach enabled by collaboration, or by trusting longer-lived measurement hosts more than new hosts.

Third, distributed measurement using client-side software raises intractable security and liability concerns. Researchers must ensure that the distribution sites are secure so that users do not download software that has been tampered with. Released code must be thoroughly vetted for security vulnerabilities to ensure that users' computers will not be compromised via measurement project software. Care must be taken to ensure that safeguards prevent the measurement software from being hijacked and used to perpetrate denial-of-service attacks and other malfeasance; the set of edge hosts performing an active measurement bear a remarkable similarity to a botnet. These concerns are particularly relevant to any projects that have released their source code, and the code itself is available for public scrutiny. Management and distribution of bug fixes and software updates is also a significant challenge for client-side measurement projects.

Finally, how might data sharing for client-side measurement results be encouraged? Standards for data queries and requests for remote measurement would help unify different projects and make the results more accessible to researchers interested in analysis more than data collection. Some standard data formats have been proposed, including in the IETF IPPM working group, but the challenge of developing a compact, extensible, easily manipulated representation of network measurement data remains.

`DIMES` [22] is an active measurement infrastructure that applies @home-style measurement. As a measurement platform, it is a collection of machines on which users have installed a freely-downloaded Java program from the NetDimes website. To minimize network impact, `DIMES` restricts probing bandwidth to 1 kB/s. `DIMES` has not reported on the trustworthiness of the data collected. It presents a useful starting point for exploring the practical issues of @home-style distributed network measurement.

`NETI@home` is a passive measurement infrastructure that uses an @home-style approach, i.e., software running on end user volunteered machines, to collect network performance and workload statistics from hosts. The software sends the resulting data to a server at the Georgia Institute of Technology (Georgia Tech), where they are aggregated to respect privacy and then made publicly available. This approach can give researchers much needed data on the end-to-end performance of the Internet as measured by end users. `NETI@home` users select a privacy level that determines what types of data will be collected. `NETI@home` is designed to run quietly in the background using few resources, with little or no intervention by the user. `NETI@home` faces all the problems that `DIMES` does, with additional privacy concerns due to the use of passively-collected packets.

Other client-side infrastructures are supporting different measurements. A positive outcome of the workshop could be to keep as

many of these platforms available as possible, and make their data available to researchers who use a centrally managed measurement platform.

Even with a community-oriented active measurement project, significant challenges to network measurement remain. Active measurements attempt to infer properties of an opaque Internet. Simply keeping pace with infrastructure deployments that impede measurement as a side effect remains a significant challenge as techniques like MPLS, VPNs, and tunnels obscure the underlying network structure. Solving the measurement infrastructure deployment and access problems frees researchers to work on more significant and neoteric problems.

2. COMMUNITY-ORIENTED PASSIVE MEASUREMENT

While there are many one-off passive measurements performed on questionably representative edge-of-network links, there are only a few larger projects (NLNR, CAIDA, Internet2) that perform systematic measurements over a long period of time and make the data available for Internet research. The volume of data involved in measurements of core network links presents a significant challenge to passive measurement projects. The cost of measurement platforms (particularly accurate network monitoring cards) and the complexity and time involved with building trust relationships to get access to relevant collection points make monitoring a network link, particularly a core link carrying traffic from many enterprises, quite difficult. The commitment of time, capital, and other resources that such projects require are out of the scope of the usual foci of researchers (published papers and theses), so few individuals or organizations attempt to collect passive Internet measurements.

Research infrastructure that is too difficult and expensive for most organizations to maintain, and yet provides a great benefit to large groups of researchers, seems like the ideal environment for widely deployed community-accessible infrastructure. However the CONMI Workshop discussion generated more questions than answers about what community-oriented passive measurement platform would be feasible.

2.1 Privacy

By far the two largest concerns in passive measurement are the privacy of the data and cost of collecting data. For years, the primary impediment to granting researchers access to data from Internet backbone links has been privacy: the privacy of users is a paramount concern. Despite widespread interest in performing measurements while ensuring individual privacy, there lacks a clear definition of what portions of network packets are in essence private. There is a dearth of information about the legality of various types of network data collections, as most relevant legislation and court precedent involves telephone networks, which are substantially different from the Internet. With the lack of information about what information they are obligated to protect, what constitutes sufficient measures for data protection, and what the potential risks of providing data to researchers are, large ISPs are reluctant to authorize official data collections in their networks. The scale of community-oriented passive measurement infrastructure would necessitate official consent, so unanswered questions of legality and privacy remain a significant barrier to development of such a measurement system.

Several strategies might resolve some of the privacy concerns that currently inhibit Internet measurement. Network measurement is not the only science in which data with significant privacy implications is collected and studied: medical science successfully collects and studies data about living human beings. For many large studies,

an independent organization collects and aggregates the data before releasing it to researchers for study. The methodologies and funding models that support this research model could prove helpful to similar efforts in network science. Indeed, many (if not most) Internet research studies are concerned with aggregate characterization of traffic, not with specific details about packet contents or communicant identity. Exploring ways to pre-process and aggregate data while preserving its research utility could mitigate privacy concerns. The Intel CoMo project [11] provides one model of allowing passive measurements that meet a privacy level pre-defined by an ISP.

2.2 Cost

Two significant costs restrict passive measurement efforts. First, passive measurement of core Internet links requires a large amount of time to build trust with ISPs to gain permission to collect data, to negotiate types of data to collect, and to secure donations of time and space necessary to deploy measurement platforms onto the Internet. Time is a scarce resource for any researcher, and attempts at establishing measurement infrastructure compete directly with research and analysis efforts, to the detriment of the scientific utility of the resulting less-than-representative datasets. Because the means to collect any core Internet data are so far beyond most researchers, the major forms of professional cachet, publications, have low standards for data used to produce research results. Thus there is little motivation to deploy significant infrastructure to get reliable data, since there is little payoff in incurring the high time cost of data collection.

Even given complete commitment to putting in the time to develop passive infrastructure, the monetary cost of developing and maintaining passive measurement collection infrastructure remains a significant barrier. Unlike active monitors, which require comparatively less CPU and disk space resources during collection, ever-increasing network speeds require significant resources at passive monitoring points. While commercial NICs are sufficient for many (but not all) active measurement platforms, robust passive measurement requires the use of network cards specialized for packet collection. The limited market for such hardware results in high prices for these cards (in sharp contrast with most other computational hardware trends towards increased functionality for decreased cost). Finally, a passive measurement platform has a short useful lifetime before it must be replaced by new, better-performing equipment. Unlike active measurement, in which the time-to-failure of the hardware components determines the lifetime of a platform, passive measurement hardware is regularly made obsolete by upgrades to the network paths being monitored. Because measurement hardware development lags significantly behind network core infrastructure (routers and such) development, passive measurement infrastructure remains locked in a vicious cycle of: traffic collection:

1. network upgrade
2. wait for new measurement hardware to be available at more than double the cost of existing infrastructure
3. attempt to get scarce infrastructure funding to cover the cost of upgrading (once the price is known)
4. finally purchase and deploy new hardware
5. traffic collection
6. network upgrade...

This cycle has occurred at least four times in the past decade, and as a result, there is at the time of this writing no current publicly

available data from an Internet backbone link. Explicitly recognizing the community value of passive Internet measurement datasets, and committing aggregated resources to maintaining a passive network infrastructure would result in more diverse and useful data for research.

2.3 Summary and Open Questions

We must resolve the question of what a community-oriented passive measurement infrastructure would collect and provide to researchers. What data would the system collect? Would the system collect the same data over time, or would different collections run at different times? How do you balance the desire to look closely at current hot topics and emerging trends with the value of consistent data for longitudinal analysis? Do you distribute entire datasets to researchers, or do you give researchers the ability to run code or otherwise query a dataset *in situ*? What are the costs and benefits of each approach with respect to privacy, security, and administration complexity? Would funding data mining to develop a repository of intermediate results for further processing and research use be a more viable and cost-effective strategy? How widely deployed should passive measurement infrastructure be? What are the trade-offs between breadth and depth of monitor coverage? What (if any) sampling should be performed on data either during collection or during analysis? How do you develop datasets that are user-friendly even to non-measurement-experts? Who is responsible for curating data?

A community-oriented passive measurement infrastructure could be a highly useful and successful endeavor, as concentration of available resources would help to solve the high cost of deploying and maintaining such a system. Unfortunately, too many unsolved problems and unanswered questions remain for such a system to be viable in the immediate future. The research need for Internet data is high, so significant resources must be put towards finding the solutions necessary to make widely-available distributed passive data collection a reality.

3. LARGE-SCALE MEASUREMENT CHALLENGES

Both active and passive measurement efforts share logistic challenges in the areas of information custody and infrastructure deployment. Problems solved and expertise gained in these areas eases collection and distribution of both types of data.

3.1 Information Sharing Complications

3.1.1 Security Concerns

Network measurement data, if abused, may provide a hit-list of potentially-vulnerable networks and hosts. Although recent publicized attacks on the Internet show that widespread disruption can be caused without such careful target selection, a shared picture of the Internet that would be valuable for researchers and operators may also have value for attackers. The possibility of enabling new attacks through accurate and comprehensive measurement is a danger that deserves study. Many active measurement projects aim to characterize the network itself, rather than the properties of edge hosts, which are the common focus of security and privacy concerns. However an attack on a core node could have a widespread and devastating impact. At the same time, we must not limit the network measurement community to studying all aspects of the Internet *except* its vulnerabilities: this is precisely where measurement has the most value. Further, as long as we depend economically on the Internet,

we are all vulnerable to widespread failure due to our inability to assess the health of networked systems. Few major outages have been caused by malicious activity; simple mistakes and acts of nature have been much more damaging—configuration errors, cable cuts, fires, unforeseen policy interactions at network borders. Comprehensive network data could improve on network stability and function by helping to identify and eliminate vulnerabilities that lead to widespread failures. This data source could be incorporated into an “Internet Center for Disease Control” as described in [26].

3.1.2 Privacy Concerns

Another social problem that inhibits information sharing is the widespread notion that network topologies and the business relationships that lead to them are proprietary. Even though significant topology data can be extracted from the public Route Views infrastructure [16], it may be sensitive information because ISPs might use such data to court customers away from competitors. The more detailed the information, the more sensitive ISPs will consider it.

Passive measurements containing packet header data, especially those also including packet payloads, are particularly sensitive—a communication channel with potentially private information is monitored and recorded. It is often technically impossible to obtain the consent of individuals whose communication is intercepted, particularly because traffic for a single session can traverse many different paths depending on network conditions and configuration. Yet significant research, including such basic questions as “What are people and organizations using the Internet for?” require inspection of packet payloads to answer. This information can provide critical input to current social, legal, and public policy questions. For example, there is a shortage of current, accurate, well-documented information on the extent of file sharing of copyrighted material.

The challenge of preserving privacy while answering questions in the public interest is not one that the research community is well-equipped to navigate. Current technologies are not ideal. Data can be anonymized as it is collected [17][12], but this can significantly inhibit extension of datasets, meaningful repetition of experiments, and the ultimate utility of the data. The success of anonymization methods depends on variability in the measured system, but communication patterns and network configurations are not random, and the underlying structures can be exploited by those intent on decrypting anonymized data. Results can be anonymized before they are published, which may protect providers and end users, but results may identify the provider involved to those with outside knowledge. More fundamentally, this model involves full disclosure to the researcher, which could be considered a significant invasion of privacy. Due to privacy and security concerns, single-organization infrastructures have been the only viable model for collecting passive (header or full packet capture) data from commercial Internet links.

There are other methods of obfuscating private information. Research labs associated with ISPs, such as AT&T, use techniques to present results without scale: percentages and fractions are presented instead of raw traffic volumes. Raw results are typically published with caveats.

Finally, researchers must actively seek to prevent the propagation of incorrect inferences. For example, a study of topology must emphasize that many links, especially backup and layer-2 links, may remain undiscovered, and that the topology must not be used to estimate the resilience of the network to the loss of a router or link, or used to assert that one network is more or less reliable than another.

3.2 Infrastructure Deployment Challenges

As mentioned in section 1, the Internet research community has used several models for deployed measurement infrastructure. Cur-

rent efforts can be classified into single-owner infrastructures, which are deployed, administered, and used by a single organization, and shared infrastructures, which are deployed, administered, and used by many organizations.

3.2.1 Maintenance

Deploying and maintaining measurement infrastructure is a significant challenge. Measurement platforms must be purchased, have operating systems and measurement software installed, and be physically installed and connected to the network in their designated location. Once measurement is begun, data must be organized, permanently stored, documented, and delivered to researchers. If data is stored or aggregated at a site remote from the measurement platform, maintaining data integrity through data transfer can be a significant challenge, particularly if the volume of data collected is large or the transfer medium is not reliably available. Data must be distributed to researchers; if the same researcher performs the collection and uses the results, this process is trivial, but providing and maintaining access to data for a community of researchers requires dedicated infrastructure and provides complications and challenges independent of data collection. Finally, researchers must be assisted with using the data as well as understanding the accuracy and applicability of the data to a particular scientific inquiry and the corresponding sources of error in the measurement [18].

Recruiting sites for measurement platforms and maintaining contact with those sites can be difficult. Timezones, language barriers, and lack of free time make coordination with local maintainers of a measurement platform difficult. Network connectivity problems, changes in local configuration (changing the IP address, installing a firewall, etc.), and lack of physical access make remote maintenance of a measurement platform difficult.

Most significantly, funding to support such mundane aspects of research as infrastructure deployment and maintenance is difficult to find.

3.2.1.1 Maintenance of Shared Infrastructure.

While the above complications affect both single-user and shared infrastructure, there are additional benefits and challenges for shared infrastructures. Shared infrastructure can have an advantage in deployment, since the benefit of using the system can motivate more organizations to contribute measurement platforms and maintain them at a high level of availability. This advantage is particularly pronounced when users are required to contribute a measurement platform before they are allowed to utilize the shared infrastructure.

Some challenges are unique to shared infrastructures. Shared infrastructure often contains more diversity in hardware, operating systems, and software than infrastructure deployed by a single organization. Moreover, timely communication about changes in the platform configuration, platform availability, or experiments that are damaging the infrastructure can be difficult. For this reason, single-organization infrastructures have been historically the most *reliable, persistent* sources of data across time.

3.2.2 Contention

When a single organization has deployed and controls infrastructure, communication with measurement platform hosts about what experiments will be run and coordination between various uses of the infrastructure is relatively simple. Shared infrastructure requires coordination and enforcement mechanisms to ensure that measurements can be run and that they are scientifically valid—the process of one measurement is not substantially changing the results of a simultaneous measurement. Moreover, the process of resolving problems with use of shared infrastructure, whether the problem

is a hardware failure or difficulty running an experiment, becomes much more complicated because it can involve a lengthy chain of inter-organization contacts.

3.2.3 Acceptable Use

All widely-deployed measurement infrastructures face the challenge of providing a uniform interface and set of capabilities while complying with a wide variety of site-specific acceptable use policies. For infrastructure dedicated to a single experiment, this constitutes making sure that the practice of performing the measurement is acceptable to all of the host sites. For infrastructure running many experiments, this requires identification of the subset of activities that are allowed across all sites. Host sites for all measurement infrastructures have local access to measurement platforms and upstream network devices and thus retain the ability to disable machines they determine are violating their AUP policies.

Shared infrastructures pose additional challenges for acceptable use policy creation and enforcement. Permission to perform experiments in a given location is often based on painstakingly established trust relationships between the hosting site and the researcher(s) who are running an experiment. As benevolent intent is not easy to correlate with a given action, many hosting sites place greater restrictions on measurement infrastructure that is used by many researchers for many purposes.

Because shared infrastructure is often used for many different experiments, often simultaneously, enforcement of acceptable use policies can be quite difficult. Even if a policy violation is detected, tracing that back to the researcher responsible can be difficult. Researcher compliance with acceptable use policies is important, as a bad experience for a host site can result in the loss of a measurement platform.

The difficulty in setting and enforcing acceptable use policies in a shared infrastructure environment has a large payoff in terms of experiments allowed and researchers aided. The necessary work required to develop platforms that allow a variety of network measurements is worth the investment of time and resources.

4. CONCLUSION

At this workshop we discussed motivation, obstacles, and platforms for community-oriented network measurement. The motivations are community frustration with limited, one-shot experiments, the need for vastly more data than is currently available from existing infrastructures, and the financial limitations of the Internet research community in sustaining or building new measurement infrastructure. For both passive and active measurement, collaboration offers increased rigor by subjecting results to academic scrutiny by repetition and cross-validation. Collaboration also offers the ability to assemble prior results to support new measurements and inferences. For example, measures of capacity can help determine available bandwidth as in Spruce [27], and measures of geography can bootstrap inference of link latency and link metrics [15]. Further, collaboration can extend deployment to sites in disparate geographic and network locations.

The workshop raised more questions than it provided answers regarding how to best use measurement infrastructure funding to support the needs of the network research community, and we expect this discussion to continue. For active measurement infrastructure, the primary considerations are how to coordinate measurement requests from a large community of researchers, how to ensure the integrity of the data if gathered by an unknown party, and how to limit perceived or actual network damage (e.g., DDOS attacks). For passive measurement infrastructure, the primary considerations are the cost of hardware for high speed trace collection and preserv-

ing privacy while supporting access to trace data. In both cases, a community-oriented program is likely to be necessary to accommodate the diverse needs of the community as cost-effectively as possible. Given the limited funding available to invest in measurement infrastructure, an objective cost-benefit analysis of the payoffs of a proposed infrastructure is essential.

5. WORKSHOP ATTENDEES

Although this report is based on minutes from the workshop, the authors are solely responsible for the text, and meeting attendees may not agree with everything written here. The CONMI workshop attendees were: Mark Allman (ICSI), David Andersen (CMU), Rob Beverly (MIT), Nevil Brownlee (U. Auckland/CAIDA), Kc Claffy (CAIDA/UCSD), Mark Crovella (Boston U), Timur Friedman (Univ. P. & M. Curie), Gianluca Iannaccone (Intel Labs), Jim Kurose (U. Mass Amherst), Tony McGregor (U. Waikato), Joerg Micheel (U. Waikato), David Moore (CAIDA/UCSD), George Riley (Ga.Tech), Colleen Shannon (CAIDA/UCSD), Neil Spring (UMD), Rick Summerhill (Internet2), Kevin Thompson (NSF), Mike Witt (U. Oregon), Matt Zekauskas (Internet2).

6. ACKNOWLEDGMENTS

This work was partially supported by National Science Foundation grants OCI-0532233 and ANI-0137121.

7. REFERENCES

- [1] Dimitris Achlioptas, Aaron Clauset, David Kempe, and Cristopher Moore. On the bias of traceroute sampling, or: Power-law degree distributions in regular graphs. In *ACM Symposium on Theory of Computing (STOC)*, Baltimore, MD, May 2005.
- [2] Active Measurement Project. <http://amp.nlanr.net/>.
- [3] José Ignacio Alvarez-Hamelin, Alain Barrat, Mark Crovella, Benoit Donnet, Timur Friedman, Matthieu Latapy, Philippe Raoult, and Alessandro Vespignani. [traceroute@home](http://tracerouteathome.net). <http://tracerouteathome.net>.
- [4] David G. Andersen, Hari Balakrishnan, M. Frans Kaashoek, and Robert Morris. Resilient overlay networks. In *Proceedings of the ACM Symposium on Operating Systems Principles (SOSP)*, pages 131–145, Banff, Alberta, Canada, October 2001.
- [5] Ran Atkinson and Sally Floyd, editors. IAB concerns and recommendations regarding Internet research and evolution. Internet Engineering Task Force Request for Comments RFC-3869, August 2004.
- [6] Suman Banerjee, Timothy G. Griffin, and Marcelo Pias. The interdomain connectivity of PlanetLab nodes. In *Proceedings of Passive & Active Measurement (PAM)*, pages 73–82, Antibes Juan-les-Pins, France, April 2004.
- [7] Paul Barford, Azer Bestavros, John Byers, and Mark Crovella. On the marginal utility of network topology measurements. In *Proceedings of the ACM SIGCOMM Internet Measurement Workshop (IMW)*, pages 5–18, San Francisco, CA, November 2001.
- [8] kc claffy, Tracie E. Monk, and Daniel McRobb. Internet tomography. *Nature, Web Matters*, January 1999.
- [9] Computer Science and Telecommunications Board, National Research Council. *Looking Over the Fence at Networks: A Neighbor's View of Networking Research*. The National Academies Press, 2001.
- [10] Benoit Donnet, Philippe Raoult, Timur Friedman, and Mark Crovella. Efficient algorithms for large-scale topology discovery. In *Proceedings of the ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, Banff, Canada, June 2005.
- [11] Gianluca Iannaccone, Christophe Diot, Derek McAuley, Andrew Moore, Ian Pratt, and Luigi Rizzo. Como: An open infrastructure for network monitoring – research agenda. <http://como.intel-research.net/pubs/como.agenda.pdf>, September 2004.
- [12] Jinliang Fan and Jun Xu and Mostafa H. Ammar. Crypto-PAN: Cryptography-based Prefix-preserving Anonymization. *Computer Networks*, 46(2), October 2004.
- [13] Sunil Kalidindi and Matthew J. Zekauskas. Surveyor: An infrastructure for Internet performance measurements. In *INET'99*, June 1999.
- [14] Anukool Lakhina, John Byers, Mark Crovella, and Peng Xie. Sampling biases in IP topology measurements. In *Proceedings of the IEEE Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, pages 332–341, San Francisco, CA, April 2003.
- [15] Ratul Mahajan, Neil Spring, David Wetherall, and Thomas Anderson. Inferring link weights using end-to-end measurements. In *Proceedings of the ACM SIGCOMM Internet Measurement Workshop (IMW)*, pages 231–236, Marseille, France, November 2002.
- [16] David Meyer. University of Oregon Route Views project. <http://www.routeviews.org/>.
- [17] Ruoming Pang and Vern Paxson. A high-level programming environment for packet trace anonymization and transformation. In *Proceedings of the ACM SIGCOMM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication*, pages 339–351, Karlsruhe, Germany, August 2003.
- [18] Vern Paxson. Strategies for sound Internet measurement. In *Proceedings of the ACM SIGCOMM Internet Measurement Conference (IMC)*, pages 263–271, Taormina, Sicily, Italy, October 2004.
- [19] Vern Paxson, Andrew Adams, and Matt Mathis. Experiences with NIMI. In *Proceedings of Passive & Active Measurement (PAM)*, April 2000.
- [20] Larry Peterson, Thomas Anderson, David Culler, and Timothy Roscoe. A blueprint for introducing disruptive technology into the Internet. In *Proceedings of the ACM Workshop on Hot Topics in Networks (HotNets)*, pages 59–64, Princeton, NJ, October 2002.
- [21] PlanetLab Consortium. Planetlab acceptable use policy. <https://www.planetlab.org/php/aup/PlanetLab.AUP.pdf>, February 2004.
- [22] Y. Shavitt and E. Shir. DIMES: Let the internet measure itself. *SIGCOMM Computer Communication Review*, 35(5):71–74, 2005.
- [23] Charles Robert Simpson, Jr. and George F. Riley. NETI@home: A distributed approach to collecting end-to-end network performance measurements. In *Proceedings of Passive & Active Measurement (PAM)*, Antibes Juan-les-Pins, France, April 2004.
- [24] Neil Spring, Larry Peterson, Andy Bavier, and Vivek Pai. Using PlanetLab for network research: myths, realities, and best practices. In *Proceedings of the Second USENIX Workshop on Real, Large Distributed Systems (WoRLDS)*, San Francisco, CA, December 2006.

- [25] Neil Spring, David Wetherall, and Thomas Anderson. Scriptroute: A public Internet measurement facility. In *Proceedings of the USENIX Symposium on Internet Technologies and Systems (USITS)*, pages 225–238, Seattle, WA, March 2003.
- [26] Stuart Staniford, Vern Paxson, and Nicholas Weaver. How to Own the Internet in your spare time. In *Proceedings of the USENIX Security Symposium, 2002*.
- [27] Jacob Strauss, Dina Katabi, and Frans Kaashoek. A measurement study of available bandwidth estimation tools. In *Proceedings of the ACM SIGCOMM Internet Measurement Conference (IMC)*, pages 39–44, Miami, FL, October 2003.