

Workshop on Overcoming Measurement Barriers to Internet Research (WOMBIR 2021) Final Report

kc claffy
UCSD/CAIDA
kc@caida.org

David Clark
MIT/CSAIL
ddc@csail.mit.edu

Fabián E. Bustamante
Northwestern
fabian@cs.northwestern.edu

John Heidemann
USC/ISI
johnh@isi.edu

Mattijs Jonker
University of Twente
m.jonker@utwente.nl

Aaron Schulman
UC San Diego
schulman@cs.ucsd.edu

Ellen Zegura
Georgia Tech
ewz@cc.gatech.edu

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 It has NOT been peer reviewed.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

In January and April 2021 we held the Workshop on Overcoming Measurement Barriers to Internet Research (WOMBIR) with the goal of understanding challenges in network and security data set collection and sharing. Most workshop attendees provided white papers describing their perspectives, and many participated in short-talks and discussion in two virtual workshops over five days. That discussion produced consensus around several points. First, many aspects of the Internet are characterized by decreasing visibility of important network properties, which is in tension with the Internet's role as critical infrastructure. We discussed three specific research areas that illustrate this tension: security, Internet access; and mobile networking. We discussed visibility challenges at all layers of the networking stack, and the challenge of gathering data and validating inferences. Important data sets require longitudinal (long-term, ongoing) data collection and sharing, support for which is more challenging for Internet research than other fields. We discussed why a combination of technical and policy methods are necessary to safeguard privacy when using or sharing measurement data. Workshop participant proposed several opportunities to accelerate progress, some of which require coordination across government, industry, and academia.

CCS CONCEPTS

• **Networks** → **Public Internet; Public Internet; • Social and professional topics** → **Broadband access, Economic Impact; Economic impact; Governmental regulations;**

KEYWORDS

Economics, Internet, Measurement, Security, Policy

1 INTRODUCTION

Although the Internet originated as a U.S.-government funded research project, it has established itself as among the most critical infrastructures in society. It has now been a quarter of a century since the U.S. government decommissioned the National Science Foundation's research backbone, part of a carefully implemented policy of transitioning the Internet infrastructure to the private sector. The Internet has transformed the world, but has also itself transformed over this period, expanding in scope, scale, complexity, and functionality.

Today's growing interest in applying machine learning and artificial intelligence methods to understanding and developing network systems amplifies the need for data to support research and development. Application of data-driven ML techniques to Internet infrastructure research brings many challenges: each network is unique, dynamic, typically not instrumented for scientific measurement much less manual labeling of data, characterized by anomalies and misbehaviours that complicate the creation of training data sets, and usually proprietary.

In the field of Internet research, academic researchers can independently probe the Internet from the edge, draw their own conclusions, subject these to comparison and peer review, and publish results. But often edge measurements allow for *inference* of properties or behavior, but not direct assessments. Much of the data that would help inform better

research is gathered internally by network operators. Indeed, operators collect substantial data on their own networks, but typically with a narrow focus and almost always with limited availability and corporate interest in the messaging, since the data may reveal aspects of their business practices that they hold close. The research community is dealing with one troubling consequence of these proprietary data sets. Sometimes a group of researchers manages to negotiate a one-time data-sharing agreement with a commercial firm, and get access to such data in order to perform research. Some papers resulting from this sort of collaboration appear to report important findings. But often an employee of the associated company is an author on the paper, which triggers concerns regarding scientific objectivity. Compounding the problem, because the data is not available to other researchers, there is no way to validate or replicate the analysis. The importance of the Internet requires that the research community move beyond this mode to a more sustained, scientific engagement.

The National Science Foundation (NSF) has recognized these challenges, most recently illustrated by their request for information on research community data needs [1]. In January and April 2021, NSF sponsored a two-part virtual Workshop on Overcoming Barriers to Internet Measurement (WOMBIR) [2]. The goal was to gather feedback from researchers on barriers related to collection, curation, management, and privacy-preserving sharing of Internet infrastructure measurements. This paper summarizes that workshop.

We describe several focus areas with research goals that measurements help answer (§2). Participants discussed common measurement challenges: visibility into today’s network architecture (§3), the need for stable, long-term observations (§4), and opportunities to apply new privacy-preserving technologies and policies (§5). We considered how to balance incentives for different stakeholders to make measurements sustainable (§6). The outcome of the workshop is a set of recommendations for enabling new network measurement and data sharing, and supporting activities to meet these challenges. (§7).

2 FOCUS AREAS FOR MEASUREMENT

The white papers and discussions covered many areas. We select three examples (not a complete list) that illustrate measurement barriers and opportunities.

Infrastructure Security. The recent Cyber Solarium Commission report [3] set out a strategic plan to improve the security of cyberspace. Among its many recommendations is that the government establish a Bureau of Cyber Statistics, to provide the government with the information that it needs for informed planning and action. A recent report from the Aspen Institute echoed this call [4]. This proposal suggests an opportunity to consider the relationship academics

could or should have with such a government function [5]. Many questions about the security, stability, and resilience of critical infrastructure will require cooperation between the private sector and academia, with the encouragement and support of governments. Topics of interest include: studying significant network outages; hijacking of routing and naming layers of the infrastructure; botnet origin, scope, and proliferation; persistent or recurring congestion and performance impairments; or simulation of “what if” questions about how the Internet would respond to disruption due to error, natural disaster or malicious attack.

Properties of Internet Access. Understanding Internet access is critical, because access shapes the vantage point of real users, and it often creates performance bottlenecks and affordability challenges. Understanding access is a grand challenge because it inherently comes with grand scale, and deep societal importance. Technologies such as 5G and low-earth-orbit satellite expand the range of options for access, as well as range of performance and affordability of these options. Access properties of interest include deployment coverage, availability, adoption, throughput, latency, reliability, and usage. Some of this data is notoriously hard to acquire, and public debate on how to document and quantify differences in these properties across the country (the digital divide) has continued for decades.

The challenges in studying access fall into two main categories: making effective use of existing data, and creating new data sets. The scale and longitudinal challenges of understanding access require creative and technically sound methods to use all forms of data collection, even those that contain inaccuracies. The FCC has a central role in requiring providers to report access data, in documenting progress in fixed and mobile broadband deployment [6], and in deploying spectrum to reduce the digital divide. The federal government’s Data Catalog indexes over 600 broadband data sets [7]. This data catalog includes data from the FCC’s Measuring Broadband America (MBA) program’s measurements on US fixed broadband access services since 2011. The MBA data provides a sample-based view on metrics such as throughput, latency, jitter, DNS performance, and several more, that constitutes a longitudinal data set that has been used by several researchers in their study of U.S. broadband [8–10].

Federating data to maximize utility calls for standardization of reporting, methods to characterize and overcome measurement bias (e.g., from crowdsourced measurements), multi-level spatial analysis and representation, and support for local contributions to national data sets that preserve privacy.

Workshop participants also discussed the limits of existing data sets. Effectively mapping broadband access over time requires technical measurements that go beyond basic access,

to quality of service and quality of experience, reliability, and usability, combined with assessment of affordability. The interaction of pricing with affordability is not well understood, and the tails of access occur in under-examined communities such as Native American reservation lands, rural communities, and low-income urban neighborhoods.

Mobile and Wireless Networking. Measurement challenges in access differ by technology. For wired access, providers know where they have deployed, which customers have which price plans, and customers have a path to complain if performance expectations are not met. Mobile broadband access, primarily achieved via cellular service but also through WiFi, is a different story. Providers deploy cell towers in known locations but have limited models for propagation and service quality as user endpoints move away from the tower. Mobile broadband QoE can be significantly affected by time-varying congestion in the cell, device type and operating system, roaming, carrier aggregation, mmWave base stations, backhaul links, and environmental factors such as topology, land cover, and buildings. Rural geographic regions lack the latest technology, dense deployments, and high-bandwidth backhaul needed to provide the same QoE as in metropolitan areas.

Beyond issues of access, another barrier to end-to-end studies of wireless networks is the data rate of raw wireless capture: a single 20 MHz downlink channel of an LTE base station produces 7 TBytes of data per day, infeasible to store and analyze at the same timescales as typical network traces, e.g., days. Another challenge is location-dependence. Scaling mobile data collection to cover many providers and geographic areas is costly, requiring huge manual effort.

Data sharing is also an issue. The community does not have a taxonomy of wireless data sets, tools, standardized sharing/metadata format etc. The CRAWDDAD data repository played an important role in the past, but held data sets mainly for upper layers and had no capabilities for storing large raw wireless captures. Its total size was about 500GB. Workshop discussions also covered the privacy issues that inhibit collection and sharing of mobile network data (§5).

An important intellectual barrier is the gap between researchers that work in spectrum/5G, and Internet measurement researchers, in part due to different theory and practice at the different layers. As a result, many questions receive less attention that they deserve, e.g., how do mobile carriers interconnect with the Internet, and other carriers? What is the role of in-network third parties, like performance optimizing middleboxes, or CDNs? What is the role of edge-components of cloud providers in mobile carrier network architectures?

3 VISIBILITY CHALLENGES

We discussed trends that are reducing visibility, from the physical layer to the application layer, impeding the ability to perform independent research.

Internet Service Providers (ISP). ISP and cloud interconnectivity trends that challenge measurement include virtualization, and sophisticated traffic engineering methods including remote peering and anycast.

Content Delivery Networks (CDNs). CDNs, including those run by large ISPs and clouds, operate rich networks of servers, often using anycast or DNS-based network traffic redirection. Such servers are often hosted in third-party networks, partially masking the CDN's presence from observation.

Home networks. Today's home networks often include wireless links, multiple wifi access points, and repeaters. The resulting heterogeneity in physical layer and media access control protocols complicate inferences of home network properties. Amplifying the challenge is the increase in number and diversity of wirelessly connected devices, including e-readers and smart home appliances

Application layer. Changes at the application layer make cloud and ISP activity less visible. Moreover, some applications that previously connected to nearby service endpoints now talk to the cloud instead, e.g., DNS resolution moving away from ISP-provided recursive resolvers. Most cloud services, such as on-line data storage and web-based documents, are complex applications whose underlying architecture and dependencies are opaque to external users.

Protocol encryption. Compounding opacity at the application-layer opacity are privacy-motivated protocol trends toward greater use of encryption, e.g., QUIC at the transport layer, TLS wrappers around plain text application-layer protocols. The tension between protecting privacy and enabling legitimate measurement and inspection is not new, but the voice representing scientific research is largely lost in the debate. One opportunity to consider is the ability to enable observations on user endpoints (with consent), so that communication can remain private in general, but researchers can work with users to get some useful information.

Cellular infrastructure. Many elements of cellular infrastructure (i.e., radio access and core networks) are not visible in end-to-end Internet measurements. Preliminary tools such as MobileInsight [11] have provided visibility into cellular network behavior on a rooted smartphone, but scaling these measurements to many phones is an open challenge. The rise of open and interoperable components in mobile network design (i.e., OpenRAN) offers exciting leverage to pursue more visibility into mobile carrier networks.

4 ADVANCING LONGITUDINAL INTERNET RESEARCH

Many properties of Internet infrastructure are important to track over time, longitudinally. If existing results are not publicly refreshed, one cannot know which remain safe to use. Yet, as challenging as it is to establish and maintain measurement infrastructure for the duration of a funded project, it is far harder to sustain such infrastructure once the project funding runs out (or the student graduates!). Incentives for publication, funding, and graduation/promotion also favor one-off snapshots that may quickly become stale. Specifically, program committees favor novelty over additional analysis of previous results, and publishing replication studies can be quite challenging, especially if those results have not changed. Evaluation of scientific promotion does not always value artifacts such as data sets or infrastructure.

These practices are mutually reinforcing, with structural limitations of funding agencies. Most funding sources fund short (three-year) research projects, and most programs are structured with budgets at the granularity of a “grad student year”. These cycles are not synchronized with longer periods of time needed to maintain measurement operations after initial development, even if a relatively small amount of on-going funding would sustain measurement. A related barrier is that funding agencies do not yet have a way to evaluate longitudinal Internet measurement research, nor an explicit program to review and renew longitudinal activities. It is important to learn how other fields of science and critical infrastructure research have addressed their data challenges.

With regard to community incentives, existing attempts to encourage public data and revisiting of results have included community awards, reproducibility badges, and reproducibility tracks at conferences. These have had only partial success due to their low professional impact relative to promotion, publications, and degrees. What the community can do on its own is simply not worth enough.

5 PRIVACY

Participants discussed advances in technology and policy tools to safeguard privacy in the context of data use. We consider the risks in sharing possibly sensitive data, evolving technologies to support such sharing, and the role of policy to augment technical methods for disclosure control.

Measurement researchers and their industrial research partners perceive privacy laws and regulations as a barrier to collecting, using, and sharing measurement data. Institutional Review Boards (IRBs) are tasked with ethical and regulatory oversight of measurement research that involves the collection, use, or sharing of personally identifiable information, consistent with ethical principles, and more recently with privacy laws and regulations [12, 13]. In this rapidly

evolving research ecosystem, IRB decisions are surprisingly variable across institutions and the community would benefit from more uniformity.

Another challenge is that researchers often do not understand the application of privacy laws and regulations to university-based research. The most pertinent regulations are the European General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA). Although companies may be subject to the GDPR and/or the CCPA, it is unclear that university researchers are.¹ Both the GDPR and the CCPA encourage forms of data minimization such as pseudonymization and de-identification; these developments are increasing interest in disclosure control technologies that can perform such data minimization.

5.1 Data Disclosure Technologies

Measurement data spans a spectrum of identifiability, from personally identifiable information (PII) that includes, e.g., an email address, to information aggregated such that it cannot be related to an identifiable person. Most Internet measurement data lies between these extremes. Common examples are data sets that include source and destination IP addresses, location, and/or portions of packet payloads. There are several technological frameworks to support work with PII [14–19]. For example, with differential privacy, a researcher does not obtain direct access to a data set but may submit queries; the amount of distortion is calibrated to ensure that a specified metric of privacy leakage remains below a specified threshold. A critical gap remains *identifying how these privacy preserving technologies can be applied to networking problems*, and where networking questions do not fit (for example, when a few queries would consume the entire privacy budget). This gap persists in part due to the steep learning curve that measurement researchers face with the advanced privacy-preserving frameworks. A thorough investigation of this area will require creating a taxonomy of data and understanding in more depth the concerns that arise about sharing, to inform design of repeatable practices to enable legitimate research access to various data types.

5.2 Disclosure control policies

There will inevitably be some questions that cannot be investigated with technical privacy-preserving tools. Fortunately, there are well-understood practices, used in this and other sectors, to responsibly share data with qualified independent scholars (Table 1) to allow replication or original research that builds on previous work.

¹The GDPR applies to entities in the European Union, to data processing related to the offering of goods or services to European subjects, and to the monitoring of the behaviour of European subjects; see GDPR Recitals 22-24. The CCPA applies to for-profit businesses; see CCPA Section 1798.140(d).

If the data minimization process results in information stored solely in a form in which it cannot reasonably be linked to a particular person, then a code of conduct can obligate the researchers to maintain the information in this de-identified form, and to not attempt to re-identify the people to whom the information relates. For research questions that require access to raw data, key-coded data may be warranted. [20]. In this case, data minimization results in two data formats: for researchers with demonstrated need, a detailed format in which the data is linkable to people; for most researchers, a less detailed format that cannot be reasonably linked to particular people. An appropriate code of conduct obligates researchers to: (a) use either form solely for research purposes, (b) use the most privacy-preserving algorithms that enables the research questions to be answered, and (c) limit access to the detailed format as much as possible. For example, a justification for access to data that includes PII that is not de-identified is to allow the joining of different data sets where the common field is the PII, e.g., an IP address (in some cases).

GDPR and the CCPA's encouragement of pseudonymization and de-identification can inform development of such codes of conduct. Such codes of conduct can serve as a model for innovative partnerships with industry to provide access to measurement data collected by companies.

IRBs would benefit from better understanding best practices associated with Internet research and approaches articulated in privacy laws and regulations, even if university research is not subject to those requirements. A limitation of IRB processes is that they require a well-defined question and experimental protocol. Exploratory research does not always start with a research question.

6 CHALLENGES IN INCENTIVES

There was recognition that even with disclosure control technologies and policies, one cannot overcome all risks to private sector stakeholders from sharing data. Release of certain data could conceivably lead to adverse commentary on some stakeholder, or policies adverse to the stakeholder's interests. While these are legitimate concerns, governments will need to encourage and participate in a solution to the tremendous counter-incentives to share data to support Internet science. In computer security the Menlo Report [12] has proposed approaches to navigate these challenges in the case of specific incident and threat data.

There are compensating benefits to the private sector in a program of data sharing. Each actor in the Internet ecosystem may have an accurate view of their part of the system, but not about the state of their competitors, or the larger ecosystem. Allowing neutral third parties to obtain data from

multiple actors can give the private sector, as well as governments and society, a global view of the state of the Internet. But the government will have to find ways to limit liability as a result of responsible sharing of data for documented scientific research.

Another benefit of sharing is the academic training of STEM professionals to work with large data sets. While synthetic network data can be used for classroom exercises, serious research of the sort that leads to professional development requires real data, with the genuine potential for new discovery.

The U.S. government could send a strong signal to the private sector that builds and operates the Internet: data sharing is a necessary aspect of sustaining critical infrastructure, the Internet has now reached this level of maturation, and (as is true in other aspects of society) responsible data sharing needs to be part of normal practice. Developing this model now is a worthwhile activity before some future Internet catastrophe forces an ad-hoc approach to Internet data sharing that would be less beneficial to operators, policymakers, and citizens. We may learn from consortia such as UIDP that are dedicated to promoting such collaborations [23].

A case study—data needs of the FCC. The FCC is a great source of research questions for the research community to tackle and could be a good consumer of results, but the FCC is not a grant-making organization, and cannot (as currently structured) directly support the research community. The NSF is good at peer review and is structured to make funding awards. Cooperation between NSF and the FCC to help the research community identify interesting and important research challenges that are of practical relevance to the FCC would benefit all parties. The Spectrum Innovation Initiative is a potential model for NSF/FCC cooperation where NSF serves as the glue between the FCC and researchers [24]. A Center would allow an MOU with the FCC for data sharing.

There are many ways that researchers can engage with FCC, from submitting research results as public comments through the Electronic Comment Filing System [25] and FCC Daily Digest [26], contacting FCC staff to set up meetings to present research, inviting FCC staff to attend workshops, hosting workshops that explicitly bring together FCC staff and researchers [27], and presenting research at a policy conference that government staff attend, such as the Telecommunications Policy Research Conference.

For this process to be effective, the research community will have to make results known and digestible for policy makers with limited technical background. This will require a different form of presenting research results than a typical paper at a technical conference. The correct incentives must exist for this to happen, which is not currently the case.

- Data is made available in curated repositories, or otherwise provided in ways that allows adequate access for legitimate scientific research
- Access requires registration with data source and legitimate research need
- Standard anonymization methods are used where needed
- Recipients agree to not repost corpus
- Recipients agree that they will not deanonymize data
- Recipients can publish analysis and data examples necessary to review research
- Recipients agree to use accepted protocols when revealing sensitive data, such as security vulnerabilities or data on human subjects
- Recipients agree to cite the repository and provide publications back to repository
- Repository can curate enriched products developed by researchers

Table 1: *Codes of conduct have been developed that enable responsible sharing of data in ways that protect stakeholders while allowing research [21, 22].*

7 RECOMMENDATIONS

The discussions yielded rough consensus on a number of recommendations.

7.1 Structuring Programs

Several recommendations are organized around different ways to structure research programs.

There was support for **building on models for interdisciplinary and cross-sector collaboration** that NSF has pioneered, including the Smart and Connected Communities and the Convergence Accelerator Programs. Such an approach could support networking researchers collaborating with social scientists and local community stakeholders to facilitate, e.g., highly granular measurement studies of broadband deployment and uptake. Creative approaches could connect to K-12 and lifelong STEM education initiatives through a citizen science model for broadband measurement.

There was enthusiasm for **center-scale efforts** focused on grand challenge problems, similar to the Spectrum Innovation Initiative to foster FCC and academic research community collaboration. Center-scale efforts could serve as clearinghouses and repositories of measurement tools and data sets, vehicles for collaboration, cross-fertilization, and actionable knowledge transfer to policy makers. This is one avenue to **stabilizing existing efforts in the research community to collect and curate critical data**.

Facilitate small-scale efforts. Workshop participants agreed it would be ideal if funding agencies could foster small-scale efforts to support longitudinal measurements, e.g., small levels of multi-year funding to continue measurements that led to successful peer-reviewed research, with an easier submission process and consideration for the impact of the previous project. One vehicle could be REUs or fellowships to contribute to productionizing existing measurements, as a supplement to existing grants.

Fund significant measurement projects, leveraging scientific research networks where possible. There was agreement on the value of a new large infrastructure project targeting visibility of security, reachability, performance, and resilience properties observable from volunteer vantage points. NSF may be able to encourage networks in the scientific community (national labs, research institutes) to share measurement data and ground truth data under appropriate research use agreements. NSF's successful programs to support infrastructure in HPC (CICI, C* programs) and connectivity (going back to the NSFNET) may also provide role models for data infrastructure.

Measurement coverage will always have gaps, so there was support for development of new research **methodologies and tools for Internet measurement research that can overcome skewness (bias) in data collection** and other limitations due to sporadic/spotty data collection.

In the area of **mobile and wireless measurement**, NSF could provide leverage by (1) supporting emerging open source initiatives, e.g., open RANs, mobile edge computing, and cellular core networks (2) promoting collaborations between wireless spectrum and wireless Internet communities, and (3) exploring new ways to incentivize development and deployment of privacy-respecting apps that gather data from mobile devices.

7.2 Challenge Goals

The workshop recognized the importance of targeting research on specific goals. In addition to the focus areas of §2, three more specific targets were identified.

Annual state-of-the-Internet report/conference. The community should consider a new conference that generates an annual community-led state-of-the-Internet report, highlighting what academic researchers know and what they

would like to know. The emphasis would be **advancing longitudinal data collection and sharing** to expand empirical coverage of network properties over space or time, as well as contributions of data artifacts. Such a conference might expedite publication of work that revisits previous empirical measurements, and include supplemental appendices with code/data in addition to a live talk at the conference. Such a track or conference might include curated guides to the data available for that year.

Test-of-time awards for data sets Conferences could consider providing test-of-time community awards to recognize efforts, including longitudinal data sets, maintenance of which has helped the community over the last decade.

Skills in data science and ethical data use There is an acute need for new coursework and training to give students skills to create, curate, and ethically use Internet infrastructure data sets, including mitigating bias in crowdsourced data, and understanding participation incentives.

7.3 Promising Mechanisms

Finally, the workshop identified new technologies or mechanisms than can assist network measurement (or for IRBs, existing mechanisms we can improve).

The federal government can play an important role in lowering the barriers to applying **privacy-preserving techniques** to Internet infrastructure data, by promoting cross-fertilization among the fields of Internet measurement, privacy-preserving algorithms, and privacy laws and regulations. Researchers need to know what privacy-preserving algorithms are available, and the benefits that they offer, without having to become experts in this field.

University Institutional Review Boards (IRBs) have an important role, to understand privacy preserving techniques and privacy laws sufficiently to evaluate privacy risks. NSF could also promote the creation and operation of an **oversight committee** (a kind of meta-IRB) to oversee some community measurement platforms, develop **best practices** around data anonymization, and **support matchmaking** between researchers and data providers.

Government can contribute by **navigating misalignment of incentives that impede data sharing**. This includes shepherding data use agreements with providers to facilitate industry contribution of large, shareable data sets for research and STEM workforce training.

Facilitate standardization of data practices To navigate the data management lifecycle, funding agencies could consider promoting (funding) the creation of working groups to standardize rules of data set generation and sharing of data artifacts, and to create common application platforms and tooling for maintaining and sharing best-practice pipelines for issuing, processing, and publishing measurements.

8 WORKSHOP PARTICIPANTS

Co-Hosts: kc claffy (CAIDA/UCSD), Dave Clark (MIT/CSAIL), and John Heidemann (USC/ISI); with Fabián Bustamante (Northwestern U.) and Mattijs Jonker (U. Twente).

Participants: Mark Allman (ICSI), Lamy Alowain (UIUC), Malte Appel (IJ), Tarun Banka (Juniper Networks), Marinho Barcellos (U. Waikato), Elizabeth Belding (UC Santa Barbara), Randy Bush (IJ & Arrcus Inc), Fabián Bustamante (Northwestern), Matt Calder (Microsoft / Columbia), Robert Cannon (FCC), Esteban Carisimo (Northwestern), Richard Carlson (DOE), Michael Chen (UIUC), David Choffnes (Northeastern), Kaushik Chowdhury (Northeastern), kc claffy (CAIDA/UCSD), Richard Clayton (U. Cambridge), Alberto Dainotti (CAIDA / UC San Diego), Bernhard Degen (KTH / NII), Ram Durairajan (U. Oregon), Flavio Esposito (Saint Louis U.), Dubem Ezeh (Drexel U.), Nick Feamster (U. Chicago), Simone Ferlin (Ericsson AB), Alessandro Finamore (Huawei), Darleen Fisher (NSF), Romain Fontugne (IJ), Avi Freedman (Kentik), Simson Garfinkel (ACM), Dan Geer (In-Q-Tel), Monisha Ghosh (U. Chicago), James Griffioen (U. Kentucky), Arpit Gupta (UC Santa Barbara), Stephen Hayne (Colorado State U.), John Heidemann (USC/ISI), Kurtis Heimerl (U. Washington), Nguyen Phong Hoang (Stony Brook U.), Mattijs Jonker (U. Twente), Scott Jordan (UC Irvine), Ethan Katz-Bassett (Columbia), Mariam Kiran (LBL), Maciej Korczynski (Grenoble Alpes U), Padma Krishnaswamy (FCC), Georgios Lazarou (TAMU - Kingsville), Ann Von Lehmen (NSF), Simon Leinen (SWITCH), Zizheng Liu (Purdue), Jason Livingood (Comcast), Qasim Lone (TU Delft), Aniss Maghsoudlou (Max Planck Institut für Informatik), Tarun Mangla (U. Chicago), Alex Marder (CAIDA), Deep Medhi (NSF), Jelena Mirkovic (USC ISI), Leandro Mondin (RNP - UFRGS), Leandro Mondin (UFRGS), Rodrigo Moreira (Federal U. of Viãgosa), Andrew Morris (GreyNoise), Alex Moura (RNP), Arvind Narayanan (U. Minnesota), Marcin Nawrocki (Freie U.), Anita Nikolich (UIUC), Jaudelice de Oliveira (Drexel U.), Christoph Paasch (Apple), Cristel Pelsser (U. Strasbourg / ICube), Chunyi Peng (Purdue), Shahrooz Pouryoucef (U. Mass.-Amherst), Lars Prehn (Max Planck Institute), Aanjhan Ranganathan (Northeastern), Leobino Sampaio (Federal U. Bahia (UFBA / RNP), Patrick Sattler (TUM), Jennifer Schopf (Indiana U.), Aaron Schulman (UC San Diego), Henning Schulzrinne (Columbia), Vyas Sekar (CMU), Ivan Seskar (Rutgers), Esther Showalter (UCSB), Alan Teixeira da Silva (Unicamp), Joel Sommers (Colgate), Dave Taht (Teklibre), Tony Tauber (Comcast), Ross Teixeira (Princeton), Cecilia Testart (MIT), Kevin Thompson (NSF), Paul Vixie (Farsight, Inc.), Sara Wedeman (Behavioral Economics Consulting), Walter Willinger (NIKSUN, Inc.), Tzu-Bin Yan (UIUC), Feng Ye (U. Dayton), Sean Yun (FCC), Ellen Zegura (Georgia Tech), Zhi-Li Zhang (U. of Minnesota), Johannes Zirngibl (TUM).

ACKNOWLEDGMENTS

We thank participants for contributing their insights, and for feedback on this report. This workshop was supported by NSF OAC-1724853, NSF CNS-2111828. Opinions expressed do not necessarily reflect views of the NSF. Any errors are the responsibility of the authors.

REFERENCES

- [1] M. Martonosi, “Request for Information on the specific needs for datasets to conduct research on computer and network systems.” NSF Dear Colleague Letter 21-056, at <https://www.nsf.gov/pubs/2021/nsf21056/nsf21056.jsp>.
- [2] “NSF-sponsored Workshop on Overcoming Measurement Barriers to Internet Research (WOMBIR 2021) – Parts I and II,” 2021. <https://www.caida.org/workshops/wombir/2101/>.
- [3] “Cyberspace Solarium Commission report,” 2020. <https://www.solarium.gov/report>.
- [4] “A National Cybersecurity Agenda for Resilient Digital Infrastructure,” *The Aspen Institute*, 2020. <https://www.aspeninstitute.org/longform/a-national-cybersecurity-agenda-for-resilient-digital-infrastructure/>.
- [5] National Academy of Sciences, “Principles and Practices for a Federal Statistical Agency, Edition 7.” <https://www.nap.edu/resource/25885/P&P%207%20Highlights.pdf>.
- [6] “Telecommunications Act of 1996, Section 706(a): Advanced Telecommunications Incentives.” (codified as 47 U.S.C. 1302).
- [7] U.S. General Services Administration (GSA), “Broadband Data Resources in data.gov,” 2021. <https://catalog.data.gov/dataset?q=broadband>.
- [8] Zachary S. Bischof and Fabián E. Bustamante and Rade Stanojevic, “Need, Want, Can Afford —Broadband Markets and the Behavior of Users,” in *Proceedings of the ACM Internet Measurement Conference*, 2014. <https://doi.org/10.1145/2663716.2663753>.
- [9] J. P. Rula, Z. S. Bischof, and F. E. Bustamante, “Second chance: Understanding diversity in broadband access network performance,” in *Proceedings of the 2015 ACM SIGCOMM Workshop on Crowdsourcing and Crowdsharing of Big (Internet) Data*, C2B(1)D ’15, (New York, NY, USA), p. 9–14, Association for Computing Machinery, 2015.
- [10] Z. Bischof, F. Bustamante, and R. Stanojevic, “The utility argument —making a case for broadband slas,” in *Passive and Active Measurement - 18th International Conference, PAM 2017, Proceedings* (S. Uhlig, J. Amann, and M. Kaafar, eds.), Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), pp. 156–169, Springer Verlag, 2017.
- [11] Y. Li, C. Peng, Z. Yuan, J. Li, H. Deng, and T. Wang, “MobileInsight: Extracting and Analyzing Cellular Network Information on Smartphones.” <http://www.mobileinsight.net>, 2016.
- [12] Kenneally, Erin and Dittrich, David, “The Menlo Report: Ethical Principles Guiding Information and Communication Technology Research,” 2012. <http://ssrn.com/abstract=2445102>.
- [13] Dittrich, David and Kenneally, Erin and Bailey, Michael, “Applying Ethical Principles to Information and Communication Technology Research: A Companion to the Menlo Report,” 2013. <http://ssrn.com/abstract=2342036>.
- [14] C. Dwork, F. McSherry, K. Nissim, and A. Smith, “Calibrating noise to sensitivity in private data analysis,” in *Theory of Cryptography* (S. Halevi and T. Rabin, eds.), 2006.
- [15] C. Staff, “Differential privacy: The pursuit of protections by default,” *Commun. ACM*, vol. 64, Jan. 2021.
- [16] D. Evans, V. Kolesnikov, and M. Rosulek, “A pragmatic introduction to secure multi-party computation,” *Foundations and Trends® in Privacy and Security*, vol. 2, 2018. free version at <http://securecomputation.org/>.
- [17] Y. Lindell, “Secure multiparty computation,” *Comm. ACM*, Dec. 2020.
- [18] P. Francis, S. Eide, and R. Munz, “Diffix: High-utility database anonymization,” in *Annual Privacy Forum*, pp. 141–158, 06 2017.
- [19] H. Corrigan-Gibbs and D. Boneh, “Prio: Private, robust, and scalable computation of aggregate statistics,” in *USENIX Conference on Networked Systems Design and Implementation*, NSDI’17, 2017.
- [20] J. Polonetsky, O. Tene, and K. Finch, “Shades of gray: Seeing the full spectrum of practical data de-identification,” *Santa Clara Law Review*, p. 593–628, 2016. <http://digitalcommons.law.scu.edu/lawreview/vol56/iss3/3>.
- [21] “CAIDA Data Stewardship Agreements,” 2020. <https://www.caida.org/about/legal/>.
- [22] kc claffy and David Clark, “Comments on the American Research Environment,” 2021. in response to Request for Information from the National Science and Technology Council, https://www.caida.org/catalog/papers/2020_comments_rfi_american_research/.
- [23] University Industry Demonstration Partnership, “Uidp website.” <https://uidp.org/>, 2021.
- [24] NSF, NTIA, and FCC, “Memorandum of agreement between the NSF, the NTIA and the FCC regarding the Spectrum Innovation Initiative.” https://www.ntia.doc.gov/files/ntia/blogimages/sii_moa_fcc_nsf_ntia.pdf, Jan. 2021.
- [25] FCC, “Electronic comment filing system.” <https://www.fcc.gov/ecfs>, 2021.
- [26] FCC. <https://www.fcc.gov/proceedings-actions/daily-digest>.
- [27] NFS and FCC, “NSF-FCC workshop on tracking quality of experience in the Internet.” Workshop website <http://aqualab.cs.northwestern.edu/conference/276>, Oct. 2015.