Workshop on Internet Economics (WIE2018) Final Report

kc claffy
UCSD/CAIDA
kc@caida.org

David Clark
MIT/CSAIL
ddc@csail.mit.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 It has NOT been peer reviewed.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.

CCS Concepts

•Networks → Public Internet; •Social and professional topics → Broadband access;

Keywords

Economics, Internet, Interconnection, Network management

ABSTRACT

On 12-13 December 2018, CAIDA hosted the 9th interdisciplinary Workshop on Internet Economics (WIE) at the UC San Diego's Supercomputer Center. This workshop series provides a forum for researchers, Internet facilities and service providers, technologists, economists, theorists, policy makers, and other stakeholders to exchange views on current and emerging regulatory and policy debates.

To try to add clarity to a range of vigorous policy debates, and in pursuit of specific, actionable objectives, for this year's meeting the organizers used a slightly different approach to structuring the agenda. Each attendee chose a specific policy goal or harm, and structured their presentation to answer three questions:

1. What data is needed to measure progress toward/away from this goal/harm?
2. What methods do you propose to gather such data?
3. Who are the right entities to gather such data, and how should such data be managed and shared?

With a specific focus on measurement challenges, the topics we discussed included: analyzing the evolution of the Internet in a layered-platform context to gain new insights; measurement and analysis of economic impacts of new technologies using old tools; security and trustworthiness, reach (universal service) and reachability, sustainability of investment into Internet infrastructure, as well as infrastructure to measure the Internet.

1. INTERNET EVOLUTION IN A LAYERED-PLATFORM CONTEXT

An important connectivity trend over the last decade has been the growth of *private networks*. These private networks use the TCP/IP architecture as underlying technology, but they are not reachable, and thus cannot be measured, from the globally addressed and reachable TCP/IP-based Internet. Provisioning private networks has become easier with the emergence of network virtualization technology in routers, which allows for slicing of physical capacity into different shares. This capability has been available for decades on long-haul circuits and more recently over some access technologies, such as the hybrid fiber/coax infrastructure of the cable industry. This ability to slice capacity among different services will be a central capability of 5G technology, which will allow wireless networks to extend their reach all the way to the mobile user.

The benefit of such private networks is that they may exercise forms of traffic discrimination, including those that lead to a better quality of experience (QoE) for users. By some accounts, these private networks carry significantly more traffic than the public Internet, implying that problems of scale, and associated innovations to accommodate growth, will appear first on such private platforms. This matters in economic terms—a substantial part of the revenues for many service providers comes from provisioning these private networks. These private networks also matter to the regulator: to the extent that they enable delivery of consumer-facing retail services, they are part of what the consumer considers as "the Internet experience", but the scope of regulatory oversight of such networks is not clear. As these private networks extend all the way to the consumer's residence, they will become more important as a part of that experience, and thus more relevant to policymakers. But at the moment, there is no way to measure or characterize the variety and volumes of traffic carried over managed, hybrid vs. traditional globally accessible routed paths. Informing discussion of policy goals or fears (e.g., provision of universal service or concerns about anti-competitive practices) related to such networks would require obtaining data from the ISPs themselves, including possibly revenues and traffic associated with different virtual slices.

Another important connectivity trend over the last decade has been for large content providers to interconnect directly with access providers in pursuit of better quality of service for users, and lower impact on heavily aggregated public Internet links. These points of interconnection that cross platform layers (e.g., IP layer to content platform layer) between parties have attracted attention from policy-makers based on potential harms. Harms can include discriminatory interconnection practices, refusal to upgrade interconnection capacity to meet demand, and unreasonable access fees. Prohibiting these behaviors would require rules and oversight, but discouraging these practices might be effective with sufficient transparency. However, a range of data would be required to achieve such transparency, including

performance data such as utilization, latency, and packet loss, as well as cost data. Traffic matrices could reveal potential discriminatory behavior (or demonstrate lack thereof), including measures of traffic flowing from content providers across direct links into access providers, as well as flowing indirectly over transit interconnections. Scott Jordan (UCI) described his current research project to develop cost models of interconnection: what determines the cost of interconnection? what determines the value of interconnection? How should the price of interconnection be determined? Regulators could use such models to determine whether an offered interconnection arrangement is unreasonable, or unreasonably discriminatory.

Many measurements required as inputs to these interconnection models can only realistically be provided by the interconnecting ISPs themselves. Some measurements can be done from the edge (e.g., by the research community, or by the FCC through some future measurement program), but there are no current methods for a third party to measure key parameters such as link capacity, and the current FCC's Measure Broadband America (MBA) program does not consider interconnection measurement. Companies are generally resistant to any sort of public disclosure of wholesale cost and pricing data, but the appropriate agency could compel the disclosure of this data under a non-disclosure agreement. It is not clear if this form of transparency would be sufficient to discipline discriminatory behavior. Rulemaking and direct oversight might be necessary if interconnection issues surface, as the FCC noted in its 2015 Order.

We discussed another interconnection modeling effort that applied Nash equilibrium to interconnection pricing negotiations. This modeling effort assumes the the existence of a second-best solution: indirect connection through a transit provider. Assuming that the cost of transit is known to all parties, this information provides a basis for negotiation. If the cost of direct interconnection is lower than the sum of the costs to the two parties to connect via a transit path, then this option is preferred. To set the payment that one of the directly connecting parties would pay the other, Nash proposed that a fair outcome would equally split the total cost difference (in economic terms, the surplus) between direct and indirect interconnection between the two parties. Of course, the applicability of this approach depends on symmetry in bargaining that may not exist in the real world. Regardless of the degree of symmetry, modeling interconnection bargaining would require both data related to traffic flows and interconnection costs as inputs.

2. ECONOMIC IMPACTS

We considered measurement challenges specific to the layered-platform context, as well as another IT innovation that has resisted accurate macroeconomic accounting: accounting for the contributions of open source software to macroeconomic metrics of productivity.

2.1 Multi-sided market platform analysis

An important framework to reason about the economics of the Internet ecosystem is that of two-sided markets, and in particular two-sided platforms to serve them. The Internet service platform itself connects broadband subscribers, content distributors, content creators, and other ISPs; in platform terms, these customers of Internet connectivity may represent multiple "sides" of the market. If ISPs also sell advertisers access to distributing advertisements to subscribers, the advertisers would constitute an additional side of the platform's market.

The Internet ecosystem supports powerful platform operators who have successfully used multi-sided markets to capture large market share by exploiting scale economies, network externalities and high switching costs/barriers. One talk exposed the potential costs to consumers from such multi-sided markets: for example consumers may not fully appreciate the value they permit intermediaries to capture from privacy intrusions through mining consumer behavior, including web site visits, searches and emails or posts. One consequence is that consumers may have to pay more for goods and services when platform operators can more accurately assess their price sensitivity through data collection and analytics. On the other hand, platform operators can defray the cost of subsidies to end users with expansive data analytics that generate new revenue streams. At one point, AT&T provided a window on such value, when it offered reduced data mining for a monthly $29 payment from its wireline broadband subscribers.

The Supreme Court in a recent ruling (Ohio v. American Express) emphasized the need to assess both sides of a two-sided market to consider whether positive impacts on consumers can offset harms on the other side of the market: in this case a vertical restraint on trade. A single market analysis would detect harm to credit card competition with higher consumer costs, since vendors could not encourage consumers to use lower cost option. A two-sided market analysis identified offsetting consumer benefits, at least to AmEX card users. Regardless of one's view of the specifics of this case, it created a precedent that can help lower courts more effectively analyze trade-offs in consumer and competitive harms and benefits of multi-sided markets.

2.2 Economic impacts of open-source software

It is well-established that measurement of the Internet's impact on productivity has been challenging, and that much of the impact of the IT sector on human lives is not captured by macroeconomic statistics such as the Gross Domestic Product (GDP), an overall measure of final goods and services in a country. Shane Greenstein (HBS) led a discussion of one specific aspect of this daunting measurement problem: how to capture the impact of open source software on GDP. Since open source software is put into service at zero cost, it does not contribute to the current accounting of the GDP. A rough and early estimate of the shadow contribution of one open source system (the Apache web server) estimated that it contributes between $2B and $12B to the U.S. GDP. This is a significant amount, and should be taken into account when assessing returns on on the federal investment in IT R&D. The measurement method to estimate Apache server deployment used a port scan of the Internet; over 44% of Apache servers observed were in the U.S., which hints the wide range of impacts of such open source IT on GDP of different countries.

2.3 Costs and revenues across ecosystem

As a points of calibration and perspective on macroeconomic analysis, global annual revenues for the telecomm sector are over $2T/year, with some estimates closer to $2.5T/year. Advertising supports most applications that define the Internet experience (e.g., Facebook). But the total global ad-

vertising market is about $500B/year, with online or interactive advertising making up roughly half. This source of revenue will not keep growing forever. The implication is that advertising (and even content) are still dwarfed by connectivity in terms of revenue. However, notably, the actual cost of provisioning network connectivity – with the exception of residential access – is dropping rapidly, even despite the significant cost spent on marketing.

3. SECURITY

We use the term *trustworthy* to describe the societal aspiration that the Internet provide experiences that are sufficiently free of frustration, fears, failures, and financial losses that people are not deterred from using it. We discussed several aspects of this aspiration, all focused on how to get a better understanding of the resilience, robustness, and privacy of networked system components.

3.1 Resilience and risk asssessment

The Internet is not the first demonstration that we are capable of building something so complex that we do not understand it. We considered waya to devise stress tests for different actors, to allow an assessment of what would happen under adverse circumstances. For example, one might ask of an ISP, "What would the consequences be if 50% of your capacity was lost?"

Mike Lloyd (RedSeal) focused on resilience at the system level as a key to better security. Much is known about how to secure low-level system elements, from better code development practices to design checklists. At a high level, there are known corporate practices for governance, risk management, and compliance. In the middle layer, system elements are composed to make an overall system. This composition depends on complex networks to connect system elements, and the complexity of service composition leads to overall systems that are hard to understand and model. At this layer, resilience of the resulting system is often lost, and insecurities emerge as an overall property of the system.

Current commercial security measurements include efforts to understand behavior of the overall system by monitoring network activity. The research community needs to move to a higher level of integrated measurement, to gather evidence about which practices work, and where measures of cost can be balanced against measures of benefit. It is unclear that such measurements can be done from "outside" the system (just as there are limits to what can be learned about the operation of a network by probing it from the edge). But measurement from inside raises issues of what firms are willing to disclose.

Several speakers called for mandatory reporting of security incidents. Many states now have some sort of requirement for mandatory reporting of a data breach above a certain size. The resulting information has been valuable to both corporate actors and the research community. However, there are a number of questions that relate to reporting of security events. First, at what level should reporting be mandated? Second, at what level of detail and to what entity should reports go? Reports to consumers alerting them that some personal data may have been stolen need not reveal the details of how the breach was accomplished, but sharing that detailed information with suitable parties would be valuable in order to learn from our collective experience.

The National Transportation Safety Board has the authority to investigate accidents above a certain level, and (without assigning blame) provide an analysis of the circumstances that led to the event. One speaker argued that for security events above a certain threshold, reporting of the event to a similar investigative agency should be mandatory. Additionally, for events below that threshold, there should be a way for voluntary reporting to an agency that would analyze, anonymize and report on the event. The analogy was the "near miss" reporting now done by many airlines to the FAA Aviation Safety Reporting System. This reporting system provides a way for data to be gathered for study in ways that provides regulatory protection for the reporting actor from any enforcement action. The FAA, to increase the assurance that reports will not be used for any enforcement action, has outsourced the operation of the system to a separate government agency, NASA.

Ultimately, tracking and analysis of security events (such as data breaches) need to be translated into determinations about best practice—what defenses are actually valuable in the prevention of harms. Josephine Wolff presented an evaluation of a specific security hygiene practice—the use of two-factor authentication (in particular, Duo) looking at evidence from several universities. Is two-factor authentication of this sort actually justified, or a cascade of follow-the-leader assumption about best practice? The evidence from this study is ambiguous. When used in high-risk areas such as financial systems, there were no breaches in the universities being studied since the introduction of Duo. When Duo two-factor authentication was added to the student email system, the measured consequence was an increase in the forwarding of email to other systems such as Google to avoid the need for the two-factor login. At the same time, the complexity of the system required the universities to dedicate considerable time to operation and support, as well as pay for the cost of the Duo service itself. The evaluation ends up as an exploration of a multi-dimensional space of consequences. There is similar complexity in trying to evaluate the benefit of another commonly accepted best practice: the requirement that passwords be complex and frequently changed. One needs actual data on the resulting level of security, e.g., subsequent system compromises; simple simulations such as the relative length of time to break different sorts of passwords do not provide real insights.

We discussed CAIDA's ongoing project to measure another security hygiene practice – source address validation, i.e., ISP detection and blocking of packets from customers where the source address is invalid. So-called *address spoofing* allows end points (either malicious or corrupted with malware) to launch various spoofed distributed denial of service (DDoS) attacks. The measurement described in this talk did not attempt to validate the efficacy of spoof prevention. Rather, the project attempts to measure compliance with the best practice. This is a difficult measurement to undertake, since it requires sending packets with invalid source addresses from inside the net being tested, and seeing whether they are dropped. The approach is a crowd-sourced effort–volunteers download test software onto their computers, in particular portable computers. As these computers detect that they are attached to a new network, they then perform a spoofing test, the results of which are collected. This approach yields results with incomplete coverage, but with reasonable confidence about the results that are ob-

tained. Given the particulars of the test performed, the consequences to those who volunteer to perform the tests are probably minimal, but in general there is a concern with this kind of experiment, in which the experimentor is asked to carry out some sort of inappropriate action to see what happens.

Evidence from this study is that "naming and shaming" those who do not implement the blocking of invalid source addresses has some benefit, and seems to double the likelihood of remediation, but still does not lead to 100% remediation. The fundamental issue is the negative externality—the cost of implementing this practice falls on the ISPs, but the benefit accrues to edge providers in the form of fewer (or less intense and effective) DDoS attacks. Economics would argue that the only long-term remedy is to internalize this externality on the ISPs. "Name and shame" is a weak form of internalization. Regulation may be more effective, but may call for corroboration that enforcing this best practice is justified (i.e., that it leads to a reduced level of attacks in practice). Whenever there is a call to improve system security by mandating some sort of practice, there will be a counter-call to prove that the proposed practice is in fact effective in practice), which brings a focus back to the other talks in this session, which were concerned with gathering data about actual security incidents and the causes.

3.2 Routing Security

The stability of the interconnected Internet depends on valid routing information such that packets follow a correct path toward the actual destination. It has been known since 1983 that if an Autonomous System (AS, an independent network on the Internet) makes an invalid (forged) assertion that it owns a set of addresses, or a path to those addresses, traffic will blindly travel to the forging AS instead of to the proper destination. The persistent failure to deal with this vulnerability should be a concern to policy-makers as well as operators—why does this vulnerability persist? One element of the problem is that the highly distributed global routing protocol (the Border Gateway Protocol or BGP) does not have any ground truth by which to judge the validity of an assertion. In practice, it is not clear how any such ground truth could be derived. Addresses are allocated by Regional Internet Registries (RIRs) but RIRs do not know from where in the Internet these addresses should be announced. There is no trust framework from which to derive ground truth that is practical in the real world. The second issue is that most current proposals to improve routing security require that all actors modify their behavior, and add cost and overhead to the processing of routing messages, while the benefit of such action would accrue to others. So the system is plagued by coordination problems and negative externalities. While a regulator might wish to intervene, the problem is global, so there is no regulator with scope and authority to act. One medium term mitigating approach is a global observatory of reachability, to maximize the chance of crowd-sourced detection and mitigation of anomalies.

3.3 Outage Reporting During Disasters

Another presentation used as a case study the tracking of telecommunications infrastructure in Puerto Rico after the devastating hurricane of September 2017. Even well into the recovery period, both operators and government officials had incomplete visibility into the state of the communications infrastructure, in particular the cellular service. On the one hand, it is a burden on operators to require that they file extensive data while they are putting their resources into recovery, but on the other hand the government needs answers to critical questions such as which regions have no ability to call 911. This event, because of its magnitude (and time to recover) should be an excellent case study of what data should be required in what sort of time frame, and as well as a basis for a thought experiment as to how much data about actual outages could be gathered and reported by third-parties, including the citizens themselves and NGOs.

3.4 DNS Privacy

One speaker presented recent development and deployment of DNS over HTTPS (DoH) and the possible risks as well as benefits of this move. Major browser providers are moving to an alternative way of implementing the DNS system to resolve URLs: instead of invoking the DNS service provided by the underlying operating system of the host computer, they are making an encrypted connection to a DNS server selected by the browser. This change implies a major shift in power and control. Traditionally, the ISP providing Internet access also provided the initial DNS server to which queries would go. Now the provider of the browser, rather than the access ISP, has control over how the DNS query is processed. DOH is a much more centralized solution: would the DNS service operated on behalf of the browser be more more or less likely to manipulate results or block legitimate results than the DNS service of the ISP? More generally, the move to DoH demonstrates how a small number of powerful actors can materially change the character of the Internet infrastructure, which participants found intruiguing, potentially encouraging and alarming at the same time.

4. UNIVERSAL ACCESS

We extensively discussed a common and long-standing aspiration for telecommunications: universal service, getting the Internet to reach the unserved parts of the country, or world. There was clear consensus on the lack of accurate data on where broadband Internet is actually available. In the U.S., access providers are required to report their coverage, using FCC form 477. Several speakers stressed that this form is known to over-report areas of coverage (if there is one served user in a census block, the block is considered served), and there is interest in other measurements that could supplement these 477 filings. The goal should be to report accurately where broadband is *not*.

David Reed (CU Boulder) described a project to analyze broadband map data gathered by the California PUC, which supplements the sort of data gathered by the FCC with coverage maps of fixed wireless providers. These maps are accurate enough that in principle they can reveal not just the percentage area of a census block unserved, but the actual number of dwellings unserved. While more granular than FCC data, there is uncertainty about the accuracy of this data as well–it may over-estimate households served given uncertainty regarding the type of dwellings located within the wireless coverage areas. Early project results confirm that data is more likely to overestimate competition and service provider availability in rural areas due to the observed higher likelihood of partial coverage of fixed wireless providers over geographically-large census blocks.

Sascha Meinrath presented a study comparing Mlab speed test data to speeds reported on FCC form 477 for census blocks in Pennsylvania; preliminary results suggest that many users do not obtain the speed reported on Form 477, a discrepancy that appears to be worse in rural areas.

Another academic study looked for robust predictors of actual broadband deployment. Not surprisingly, population per road mile is a good predictor. As population density falls off, availability (and speed) of broadband drops, and probability that the region will receive supplemental funding goes up. Broadband availability is also correlated with median household income and median home value. However, little data is available on pricing, which makes it impractical to model the influence of cost on adoption. An ideal suite of measurements would include availability, usage, performance and pricing. Performance can be gathered from the edge, such as from the MBA platform. Pricing could be gathered from representative samples, and availability data should be fine-grained, ideally at the street level.

In the U.S., Federal subsidies have been used to stimulate broadband deployment in unserved areas. At the moment, separate programs subsidize residential service and mobile service. The U.S. universal service subsidy programs total $8.7 billion per year. Of that, $4.5 billion is directed to high-cost support, and of that $500 million is targeted for expanding access to 4G mobile broadband services (the Mobility Fund – Phase II (MF-II) program). In 2018, the FCC implemented MF-II using a complex reverse auction process to identify both the service areas that were eligible for funding and the reverse auction through which the funds were assigned to qualified providers. Since the 477 data on service availability is known to overstate the locations where broadband service is available, the FCC instituted a challenge process through which certain interested parties could submit measurement data demonstrating a lack of available mobile broadband service in areas that the 477 data indicates is already served by one or more un-subsidized providers. Identifying zones for additional subsidies is contentious because failure to qualify a zone may deny funding to providers and communities without adequate service coverage; whereas falsely identifying zones as unserved could result in subsidized competition that could threaten the economic viability of providers already providing mobile broadband to those communities. The challenge process thus represents a key mechanism for validating service provider submitted data.

The MF-II process offers an interesting opportunity to study the design of effective Universal Service policy programs. Most economists prefer the mechanism of a reverse auction for allocating universal service subsidies. The first such auction was used during the Mobility Fund – Phase I in 2012. But one serious concern is the lack of understanding of the cost structure of high-cost deployments, including rural areas and the developing world. While there is a perhaps understandable focus on the capital costs of initial deployment, the ongoing operational costs actually determine the long-term sustainability of high-cost deployments. Financial projections for these systems often underestimate the ongoing operational costs, due perhaps both to unjustified optimism but also to the difficulty of estimating ongoing operating costs. Case studies of actual deployments are critical here, and they must study deployed systems over time to learn about sustainability, not just initial deployment.

5. ACCESS—DEFINING AND MEASURING

There have been many attempts to define what constitutes broadband service. The most common measure of broadband performance is speed (downstream and upstream bandwidth). Increasing deployment of advanced cellular technology such as 5G may allow the same infrastructure to be used for mobile service and for residential service (fixed wireless). This possibility raises definitional challenges. For a cellular technology such as 5G to receive supplemental funding to serve as fixed, residential access, there must be an agreed specification of what broadband service it provides, in terms that allow comparison to other (e.g., wireline) options in a technology-neutral way.

Given that cellular service (at least at present) displays more speed (throughput) variation, both over time and location, than wireline offerings, a comprehensive measurement program would likely be required as part of qualifying a cellular service as a substitute for a wireline solution, and in particular for such a service to qualify for subsidy funding. Such measurement could be done via active probing or passive monitoring. Active probing from customer-provided devices may suffer distortions, e.g, from impairments in home WiFi. Such testing may also consume substantial network resources. Speed estimation based on passive monitoring is technically complex, likely to be imprecise, and could raise concerns about privacy. In consequence, one argument is that tools to measure access performance should be built into ISP-provided home interface devices (home routers or modems), as part of an integrated and ubiquitous quality measurement program.

A larger consideration is that for many, speed is no longer the most important measure of broadband service quality. Mark Johnson and Anita Nikolich (IIT) posed a challenge to develop a better, yet measurable definition of "good broadband" that includes speed, coverage, protection of privacy, lack of discriminatory treatment of traffic, cost, conformance to ISP best security practices (e.g., anti-spoofing filters), cost of service, clarity of advertising, corporate responsibility, and support for privacy-respecting scientific research. One speaker posed the challenge of defining a "privacy index" that could be part of such a definition. A resulting challenge would be to convert these metrics into a visual representation that aids consumers and governments in understanding and comparing different services.

6. OTHER POLICY CONCERNS

Harold Feld (Public Knowledge) described the consequences of the German law called NetzDG, which went into effect January 1, 2018. The law requires social media platforms with 2 million or more users to take down "clearly illegal" content within 24 hours of notification, or 7 days if the determination of illegality is more complicated. These platforms must publish a report twice a year containing: the number of the received complaints, the number and qualifications of employees who are handling the complaints, the network's association memberships, the number of times an external party has been used to decide the illegality of the content, the number of complaints that led to the content being deleted, the time it took to delete the content, and measures that were taken to inform the complainant and the member who posted the deleted content.

What can be learned from these transparency reports?

According to Reporters Without Borders, the law has led companies to delete large amounts of content that was in fact legal in an effort to ensure that they will not be punished under the Act. When deleting the content, Facebook and Google cite their community standards. In these standards they stipulate what kind of content users may share on their platforms and reserve the right to also remove content that is protected by communicative freedoms.

To consider a situation like this, Harold emphasized that it is important to start with fundamentals: what is the harm the law should actually be preventing? Is it to prevent the recruitment of terrorists (can a domestic law in one country be effective at this?), reduce crime, make users feel safer, or perhaps create a cleaner space to present ads? A complicated space like this should require that legislators or regulators articulate their actual objective, rather than skip that step on the grounds that it is "obvious".

Researchers need to socialize the complexity of the problem, while not suggesting that complexity should mean paralysis. All parties should acknowledge that problems may not become clear until after policies are implemented and scaled. To perform research in this area, we need to find ways to encourage or require platforms to give researchers access to raw data on how they do content moderation. If reporting obligations are going to be required by law, legislators and the research community need to think about what should be tracked long term. The requirement should be qualitative metrics, not just quantitative.

7. LOOKING TO THE FUTURE

Because so many Internet challenges are now rooted in economic or policy concerns, it is good to see the research community has expanded its use of measurements to attempt to inform these issues. But a recurring challenge for the network research community is to sustain a measurement infrastructure of sufficient scope and generality. ISPs might instrument their own customer premise equipment to measure specified parameters of service quality, but these devices will not generally be available for experiments proposed by third-party researchers. The research community needs to develop and put in place a scheme that provides an incentive for users across the Internet to allow their network connection to be used as a source for carefully managed measurement campaigns. Today researchers can attempt something similar by rewarding users who perform experiments through systems such as Mechanical Turk. However, a more effective scheme would be to motivate users to install a platform that allows researchers to instigate measurements in a more general (but controlled) way. One approach to providing an effective incentive might be to use some sort of crypto-currency reward when a user's device performs a measurement experiment.

8. PARTING THOUGHTS

Participants offered many interesting perspectives on what they learned at the workshop, in a closing session and an exit survey. We catalog some of these below.

- There is an expanding awareness that if policymakers hope to rely on academic or scientific research to inform policy, there will need to be increased accuracy and disclosure of data relevant to a given question. As the ecosystem evolves, required measurements/reporting could span from metrics such as security incidents; outages; broadband availability, cost, and pricing; cloud computing capacity and traffic; consumer usage patterns; how various parties in the ecosystem are using consumer data. Policymakers and academics must tie the need for these measurements to concrete harms that they would supporting monitoring or avoiding. There is also an increasing need to identify sustainable sources of funding for independent, open, trusted measurement of the Internet, and its communication to users and policy makers.
- One repeated "low fruit" suggestion was to require a programmatic API for accessing basic broadband service tier information, which would facilitate use of FCC MBA data, and also stimulate innovation of other measurement technology.
- The current theories and practices to deal with market concentration and antitrust are arguably failing to support the public policy needs of the IT space. The Internet ecosystem is distinctive with respect to speed of growth, mutation, amplification, prevalence of multi-sided markets, and network effects. Decisions about mergers and market concentration in multi-sided platform economics cannot rely on single-market metrics for evaluation.
- Increased data mining of consumers that allows perfect price discrimination may have the unintended effect of eroding the operation of capitalist markets, which depend on a degree of information symmetry.
- An important higher-level question is the character of the public space that is the Internet, how it is changing, and which actors have the power to influence that change. This is more important, but much more challenging to measure, than specific mechanisms such as routing or peering.
- The likelihood of federal regulation is increasing if only to mitigate the risk of dealing with a patchwork of state laws related to network management or piracy. The research community is in a position to inform regulation, and hopefully prevent poor regulations, as well as measure the impact of regulation (or at least what happens after regulation, since causation is difficult to establish). Measurement should be the foundation for a discourse about what would define good regulation. Actual data may be the best antidote to the current partisan divisiveness.

9. WORKSHOP PARTICIPANTS

- Co-Host: kc claffy (CAIDA/UC San Diego)
- Co-Host: Dave Clark (MIT/CSAIL)
- Co-Host: Amogh Dhamdhere (CAIDA/UC San Diego)
- Steven Bauer (MIT)
- Ken Calvert (NSF)
- Richard Clarke (AT&T)
- Constantine Dovrolis (Georgia Tech)
- Harold Feld (Public Knowledge)
- Rob Frieden (Penn State University)
- Dan Geer (IQT)
- Shane Greenstein (Harvard Business School)
- Geoff Huston (APNIC)
- Mark Johnson (U. of North Carolina)

- Scott Jordan (U. of California Irvine)
- William Lehr (MIT)
- Mike Lloyd (RedSeal)
- Sascha Meinrath (Penn State University)
- James Miller (FCC)
- Anita Nikolich (Illinois Institute of Technology)
- Andrew Odlyzko (U. of Minnesota)
- Jon Peha (CMU)
- Achilles Petras (BT Applied Research)
- David Reed (CU Boulder)
- Henning Schulzrinne (Columbia University)
- Marvin Sirbu (Carnegie Mellon University)
- Tony Tauber (Comcast)
- Josephine Wolff (Rochester Institute of Technology)
- Christopher Yoo (U. of Pennsylvania)
- CAIDA (8 total)
- Roderick Fanou (CAIDA/UC San Diego)
- Marina Fomenkov (CAIDA/UC San Diego)
- Alistair King (CAIDA/UC San Diego)
- Ricky Mok (CAIDA/UC San Diego)
- Joshua Polterock (CAIDA/UC San Diego)
- Mingwei Zhang (CAIDA/UC San Diego)

Acknowledgments

The workshop was supported by CAIDA members, and by the National Science Foundation's Computing and Networking Systems Division CNS-1513847, and the Office of Advanced Cyberinfrastructure OAC-1724853 grants. This report reflects the views of the authors and some workshop participants, but not necessarily the National Science Foundation.