Workshop on Internet Economics (WIE2015) 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 It has NOT been peer reviewed.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.

ABSTRACT

On December 16-17 2015, we hosted the 5th 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 inform current and emerging regulatory and policy debates.

The FCC's latest open Internet order ostensibly changes the landscape of regulation by using Title II as its basis. This year we discussed the implications of Title II (common-carrier-based) regulation for issues we have looked at in the past, or those shaping current policy conversations. Discussion topics included differentiated services on the public Internet, evolving approaches to interconnection across different segments of the ecosystem (e.g., content to access), QoE and QoS measurement techniques and their limitations, interconnection measurement and modeling challenges and opportunities, and transparency. The format was a series of focused sessions, where presenters prepared 10-minute talks on relevant issues, followed by in-depth discussions. This report highlights the discussions and presents relevant open research questions identified by participants.[1]

Categories and Subject Descriptors

C.2.5 [**Local and Wide-Area Networks**]: Internet; J.4 [**Social and Behavioral Sciences**]: Economics

Keywords

Economics, Internet, Network management

1. INTRODUCTION: QUALITY OF EXPERIENCE IN A TITLE II WORLD

In December 2015, UC San Diego's Center for Advanced Internet Data Analysis and MIT's Computer Science and AI Laboratory co-hosted the 5th interdisciplinary Workshop on Internet Economics (WIE) at the University of California, San Diego. In hosting this workshop series we recognize that the future of the Internet is shaped as much by economic factors as by technical innovations, and our goal is to provide a forum for researchers, commercial Internet facilities and service providers, technologists, economists, theorists, policy makers, and other stakeholders to empirically inform emerging regulatory and policy debates.

[1]Slides presented and this report are available at `http://www.caida.org/workshops/wie/1512/`.

The topic for the previous WIE workshop (2014) was "Aspirations for a Future Internet." This year (2015) we focused the workshop on implications of the FCC's February 2015 open Internet order, which changes the policy landscape by using Title II (common-carriage-based) Internet regulation.[2] We discussed the implications of Title II-based regulation for technology, business, and research developments that are shaping the future of the Internet. We held sessions on differentiated services on the public Internet vs. specialized services, evolving approaches to direct interconnection, scientific measurement techniques and their limitations, and data sharing. This report highlights the discussions and presents relevant open research questions identified by participants.

David Clark introduced the workshop with a framing talk that drew on insights from a recent workshop on measuring Internet quality of experience (QoE) that he co-chaired and the FCC and NSF jointly sponsored. "All measurement is political," so it is important to know what one is measuring and why it matters. A focus on QoE moves the measurement effort from the network to the user. The emerging field of QoE research [19] recognizes that measuring the user's QoE is harder than measuring the network, but even harder is meaningfully linking these two classes of measurements. Measurement of actual QoE requires assessments of complex subjective experiments, which means interviewing the user about how well a given application works relative to one's expectation. Many factors, including the user's mood, can affect such subjective measurements. This kind of experiential research is important but difficult and expensive.

Given this complexity, people instrument applications and measure network performance as a proxy for QoE, but sometimes experiential research reveals that the hypothesis linking performance optimization to QoE is flawed. As an example, much attention to adaptive coding can allow an application to send the best possible signal at any moment. However, QoE research reveals that rapid changes in coding lead to reduced reported QoE [19], and users notice when their experience degrades more than when it improves. Users also tend to remember the quality of experience at the end of a video more than at the beginning.

2. LINKING QOE TO QOS

While QoE is a complex, multi-dimensional concept, only certain aspects of QoE are likely strongly linked to network QoS. An open question is how deeply we need to under-

[2]The common carriage basis for U.S. regulation of telecommunications carriers is articulated under title II of the Communications Act of 1934 [1].

stand end user QoE in order to assess its relationship to network performance. Objective measurements may occur at various parts of the service delivery chain, but we lack methods to use these measures to estimate user QoE. There are two general approaches to doing so: objective QoE models, which map application layer performance indicators to quality metrics; or direct user feedback, which is less scalable but can inform objective QoE models. Some have suggested simple proxy measurements of QoE, such as a user aborting a download, but other factors may trigger such behavior, so one must approach such a proxy measure with skepticism. Patrick Zwickl (U. Vienna) emphasized the need for research to investigate, which requires measuring, the role of pricing in gauging a user's marginal utility of QoE. He noted that QoE has not really been tested in a real-life context where a user has to pay to improve quality in real time.

Another long-term research goal is to develop applications that can interact with the network to improve QoE, which gives rise to two questions: should there be explicit signals of network behavior from the network to the application, or signals of network requirements from the application? Especially in the resource-constrained wireless world, there is growing attention to how to use differentiated services (QoS) to improve application performance, and questions about whether we have the right architectural pieces to support it [9]. The wired world also exhibits noticeable performance impairment at times, which raises question of measurement, especially at interconnection points between network infrastructure providers where controversies have arisen.

A related challenge is merging multiple data sources, especially when incentives to share data are limited. Progress in integration of measurement data will require collaborations between application designers and network engineers, who do not necessarily see it in their interest to share data with each other.[3] Participants noted that this issue was part of a larger conversation underway in various forums potentially change how the Internet research community thinks about measurement, in particular to expand its scope to include application developers as well as network operators in pursuit of higher quality of experience from the Internet.

To balance the legitimate need to provide higher QoE to some users, with the risks of service discrimination that might harm consumers, the FCC's 2015 Open Internet Order described the boundaries of legitimate use of differentiated services by ISPs. Most concretely, if the user has control over service quality selection, it is probably acceptable, but provider control is likely to trigger regulatory attention. Many ISPs do use some differentiated service capabilities to improve QoS within their own networks, and there are emerging technologies for use within a network, including active queue management (a type of QoS) features in DOCSIS 3.1 to mitigate bufferbloat for gamers, without requiring additional complex network engineering. In contrast, the market for interdomain differentiated services (varying levels of QoS coordinated across multiple ISPs) remains unproven at the retail level, and requires intense engineering and operational and business complexity, so is likely to remain undeveloped for the foreseeable future.

For another example of the complexity of QOE-related measurement, the FCC's 2015 February Open Internet Order [12] also added a requirement that ISPs must publish loss rates, or participate as a safe harbor in the FCC's Measure Broadband America program in which the FCC plans to integrate loss rate measurement and reporting. This requirement begs the question: what should loss rate numbers be? A TCP/IP network with zero loss rate is not necessarily a well-functioning network – its primary transport protocol (TCP) functions by adapting (its sending rate) as a response to loss. Furthermore, what loss rate actually impairs the user experience? The questions of what is worth measuring, and to whom it should matter, are complex and political.

[3]Some discussion of how to incentivize cooperation was the focus of CAIDA's AIMS workshop in February 2016 [2].

3. MACROECONOMIC CONTEXT

Madura Wijewardena (Comcast) discussed macroeconomic factors influencing the broadband industry. He emphasized that limited economic growth and in particular limited household income growth constrains investment in broadband infrastructure. Some economists in the room argued that it was not merely a matter of access to capital but regulatory factors, or the reality of a real competitor (e.g., Google Fiber, Virgin Media in the U.K.) that observably incents or disincents network infrastructure investment in a region. Some participants noted that in many U.S. cities, facilities-based competition no longer induces as much investment in infrastructure as it did years ago; ISPs are instead growing revenue via managed services and vertical integration with content companies to enable bundling of services.

Andrew Odlyzko presented contrasting thoughts on economic incentives to provide differentiated services in light of history of other transport infrastructures. He noted the huge range of research literature but little commercial success in differentiated services in the telecom industry, unlike the transportation industry where it is the norm. He also noted that, historically, the push for differentiated services did not come from an objective need but from the desire of providers to squeeze more profits from customers. He used a recent extreme case of price discrimination, with Verizon charging up to $20/MB for international data, while T-mobile disclosed to the FCC that its actual cost of delivering mobile data was $0.002/MB, and the median wholesale charge to international carriers was $0.2, with a 95th percentile charge of 40 cents per MB!

Discussion turned to the Uber and AirBnB phenomena as examples of the search for choke points in a service delivery chain. Uber makes the transportation economy much more efficient, but at the cost of squeezing a plentiful resource: drivers. Uber can continue to drop rates, but inertia and network effects still yield high revenue. Scholarly journals are another example, in which publishers are squeezing inefficiencies out of academic libraries. More than two-thirds of the cost of running a university library system are internal, i.e., paying for librarians and libraries, both of which are gradually being eliminated.

There was not consensus on how much network infrastructure companies were themselves being squeezed out of profits by content companies. In terms of network infrastructure costs, the developing dominance of over-the-top video presumably reduces the revenue per bit while increasing network infrastructure costs. The bit transport industry may transform itself to become more efficient, but right now these businesses are typically not as profitable as the biggest content and/or hardware vendors. noted that the telecommunication companies would consider a 15% level of capital investment high, but it is low relative to electricity or rail-

road infrastructure providers, who invest closer to 50% of their capital. Historically, extraordinary profits, when they occur, tend to go into protecting choke points, i.e, inhibiting competition, and thus sometimes attract regulatory interest. Andrew noted the example of Mexico, one of the most ossified economies in the world, where the telecom provider makes huge profits charging exceptionally high rates, making Carlos Slim one of the richest men in the world.

Andrew noted another anomaly: corporate profits rebounded quickly after 2008 and have not dropped, while labor (wages) took the sustained brunt of that crash. Quantitative easing of monetary policies, in the U.S. and elsewhere, played a role in this historically unusual outcome. But the lack of wage recovery has also had an ossifying effect on the economy. Other measures of economic dynamics are also slowing down in the U.S., e.g., creation of new businesses. These are oversimplified observations of a complex economic situation, but they are all consistent with Madura's data.

4. MEASUREMENT OF DIRECT INTERCONNECTION

David Clark talked about what he considered an emerging phenomenon of interconnection in today's ecosystem, sometimes referred to as *direct interconnection*, but the phrase does not capture the most important aspect. It is not just that some settlement-free peering relationships are becoming paid peering relationships, but that fundamentally different types of companies are interconnecting, i.e., content providers with access providers (e.g., Netflix to Comcast). From a BGP routing policy perspective, these relationships look more similar to traditional peering than transit relationships, because the routing configuration implements connectivity to the access network, not the rest of the Internet. But historically, the business concept of peering was associated with networks that not only exchanged similar volumes of traffic, but also were in the same line of business. Another important distinction in today's direct interconnection is the ability of the content provider to originate flows of traffic across potentially many paths (including indirect connections) in order to distribute load and cost, and to improve performance. The research and policy communities have not yet understood the implications of how this mode of interconnection differs from traditional transit and peering, in technical and business terms.

kc claffy (UCSD/CAIDA) provided a distinct example with the recent ATT/DirecTV measurement project. In compliance with the merger agreement between AT&T and DirecTV, the FCC required that AT&T report on performance characteristics of traffic exchanged at interconnection points with its major peering parties and on-net only customers. The FCC further required AT&T to engage an Independent Measurement Expert (IME) to define the measurement and reporting methods.[4] This project had two objectives: to convince the FCC that AT&T was not using the points of interconnection as a source of leverage against competing interconnecting parties, but also to gain some insight into what is actually happening in the ecosystem. In this context, some of the most important data about an interconnection link is not directly available to AT&T, but only to the interconnecting party on the other end of the link. In particular, data about packet losses in the incoming direction can only be gathered from the router belonging to the interconnecting party. The IME also proposed a measurement method that required actively probing the interconnection link, and emphasized that cooperation of the interconnecting party was essential to achieve the most robust data about losses and jitter using such probes.[5] As a lever to incentivize cooperation, some economists noted that the FCC could deprioritize future complaints from companies unwilling to cooperate on measurement experiments.

Jason Weil (Time Warner Cable) gave a high-level overview of large scale measurement standardization in the IETF [6, 18, 7] partially motivated by the FCC's MBA program in 2012 [3], to support a standards-based broadband performance measurement architecture, metrics, data structures, and protocols to support communication between pieces of the architecture. Two concurrent standardization efforts – the IETF and the Broadband Forum – evolved into complementary specifications for broadband measurement architectures. The two primary use cases were measurements by the ISP itself and by the regulator.

5. INTERCONNECTION MEASUREMENT AND MODELING RESEARCH

Georgios Smaragdakis (MIT/TU Berlin) presented a recent measurement method – Constrained Facility Search (CFS) – that estimates the exact facility at which a given peering interconnection occurred. Annotating Internet interconnections with physical coordinates at the level of a building can help locate points of congestion or instability. The CFS method relies on PeeringDB and other public data about peering facilities, derives constraints through IP address alias resolution and geolocation, and uses targeted traceroute measurements to narrow the set of possible facilities hosting an interconnection, sometimes down to a single facility. A key insight of this method is that inference of the technical approach for an interconnection sufficiently constrains the number of candidate facilities such that it is often possible to identify the specific facility where a given interconnection occurs. Validation via private communication with operators confirmed the accuracy of this method. This study also revealed the different roles that interconnecting routers play – in many cases the same router implements both private interconnections and public peerings, in some cases via multiple Internet exchange points.

Vasileios Giotsas (UCSD/CAIDA) presented one of CAIDA's most exciting measurement research infrastructure achievements for this year: a unified web interface for to enable querying of all known open traceroute servers [15]. Looking glasses (LG) servers enhance visibility into Internet connectivity and performance by offering a set of distributed vantage points that allow data plane and control plane measurements. However, lack of input and output format standardization and limitations in querying frequency have hindered the development of automated measurement tools that would allow systematic use of LGs. Vasileios developed Periscope, a publicly-accessible overlay that unifies LGs into a single platform and automates the discovery and use of LG capabilities. The system architecture combines crowd-sourced and

[4]The FCC and AT&T selected CAIDA as the IME.

[5]The preliminary report of the IME was submitted at the end of December 2015 (after WIE) [16]. The IME is in the process of preparing a final version of the measurement and reporting methodology report for approval by the FCC.

cloud-hosted querying mechanisms to automate and scale available querying resources. Periscope can handle large bursts of requests, with an intelligent controller coordinating multiple concurrent user queries without violating the various LG querying rate limitations. As of October 2015, Periscope had automatically extracted 1,708 LG nodes in 262 Autonomous Systems. Periscope significantly extends the view of Internet topology obtained through RIPE Atlas and CAIDAs Ark, while the combination of traceroute and BGP measurements allows more sophisticated measurement studies. This infrastructure is now supporting several research projects [14, 13, 11, 4].

Scott Jordan (UCI) proposed an interconnection research challenge for the academic community: formalize a model of interconnection, which must answer basic questions to parameterize the model: Interconnection to do what? To deliver packets where? Should pricing be coupled to packet delivery distance? How do we know whether interconnection fees reflect service provided? A useful model must capture not only what service is provided, but the performance, and cost structure, including cost of colocation, cost of transit, incentives to localize traffic, and evidence of market failure. Some engineers noted that congestion is not necessarily a signal of market failure although persistent periods of unrelenting congestion likely merit regulatory attention. How does such a model handle the fact that a few customers may drive most of the operational cost of the network? Can researchers create a model that faithfully reflects the Rube Goldberg reality of the Internet interconnection ecosystem?

Later in the workshop Amogh Dhamdhere (UCSD/CAIDA) offered a partial response to Scott's challenge: a collaborative research project he started with another computer scientist and an economist at Georgia Tech to model the economics of contractual agreements for Internet interconnections. Motivating this inquiry is the basic public policy questions: Why do peering disputes keep happening? What should change to produce more stable peering relationships? He described four models of network formation to explore: (1) monopoly access network with subscribers on one side and content on the other; (2) duopoly or oligopoly; (3) access net and transit net in the middle with subscribers on one side and content on the other (no longer two-sided market); and (4) full mesh: multiple access and transit networks with subscribers on one side and content on the other. Some of these models are too complex to solve analytically, but are amenable to agent-based simulation. Among other ideas they will explore are more complex AS relationships, e.g., per prefix pricing, or pricing based on cost, demand, distance, traffic type, competition.

A fifth suggested model is a hybrid of (2) and (4), where content may be connected directly to a provider but also using transit elsewhere to reach that provider. Andrew noted that most studies model the Internet as a content delivery infrastructure, which is why the two-sided market model is so common. His view has always been that content is not as important as connectivity, which prevailing models fail to capture. Bill agreed that a two-sided market model only captures certain elements – it does not tell us anything about dynamics. These models break easily. An agent-based approach offers a promising alternative.

6. HIGH BANDWIDTH MEASUREMENT CHALLENGES

The FCC faces two primary technical challenges of measurement in the the wake of the Open Internet Order: its underspecified requirement for packet loss measurement; and scaling measurement techniques to gigabit speeds.

Several national regulators (including the FCC in the U.S., the European Commission, and Anatel in Brazil) have attempted to assess packet loss behavior, but they report on packet loss in different ways, and measurement of different loss-related metrics vary significantly. Packet loss measurement and interpretation of resulting data in a QoE context are both still research problems; there are still no plug-and-play solutions.

Walter Johnston (FCC) described the challenges of packet loss measurement, and scaling throughput measurements to gigabit bandwidths, in his update on the MBA measurement program [3]. Although the FCC did not launch the MBA program to be a permanent infrastructure, growing interesting in broadband performance has led the FCC to sustain and expand the program under the Open Internet Order, which establishes requirements for transparency of metrics related to bandwidth, packet loss, and latency. The FCC is also considering how to use the MBA program to improve their own understanding of interconnection performance. The MBA video performance tests are progressing slowly but they have managed to execute some tests of Youtube, Netflix, and Hulu. Unfortunately, the growth in access capacity has challenged the measurement capabilities of the MBA measurement infrastructure, which is starting to hit bottlenecks that are not always easily detectable. The MBA program is now trying to qualify a gigabit router to add to the measurement platform, which will hopefully enable them to include vantage points of Google Fiber customers. In parallel, Walter emphasized that the FCC is trying to leverage the existing MBA infrastructure to support the academic Internet research community where feasible.

As an example, in 2016, CAIDA will deploy measurements on MBA that attempt to detect the presence of CGN in U.S. broadband access networks. Strategic decisions to deploy CGN vs. IPv6 may also have observable implications for QoE. Comparison of IPv4 to IPv6 QoE is getting limited attention from MBA today, but the FCC would try to support researcher interest in IPv6 measurement.

The FCC has launched a mobile version of the MBA program, which has moved more slowly in order to carefully navigate the privacy issues associated with mobile data. Early mobile measurements also triggered technical questions about statistical validity that delayed public reporting.

Steve (MIT) spoke on measurement challenges specific to high bandwidth networks such as Google Fiber [8]. Many measurement systems and services do not perform accurately at high bandwidths, and potentially will shed some emerging services in an inappropriately negative light. Given regulatory interest in measurement results, this factor alone could deter investment in gigabit networks. Lack of current techniques and tools is one problem, but another issue is reasonable expectations for gigabit broadband service, given that the access provider does not control the end-to-end path. Complicating the picture is the fact most Internet access today involves local caching, CDNs, or other mechanisms to get content closer to the consumer to improve performance. Some municipal networks may choose to prefer local connectivity and prevalent use of caching over band-

width commitments. In the U.S., the FTC would argue for the need for truth in advertising claims and clear explanations to the user, which enable easy comparison with competing services.

A related issue is the growing disparity in access bandwidth across cable vs. DSL vs. satellite. The new consumer-focused broadband "nutrition" labels [5] are the FCC's attempt to facilitate consumer understanding and informed choice, and the FCC expects to evolve those reporting requirements as they understand more about what makes sense for consumers. If an ISP advertises gigabit bandwidth, some edge-based measurement technique should be able to verify the service is actually delivering that bandwidth. Scott noted that the disclosure requirements for these nutrition labels are loosely defined, e.g., the FCC did not specify the location of the server for bandwidth measurements.

Bill Lehr (MIT) propounded a provocative thought experiment: perhaps video-over-IP should be considered a specialized service in the current terminology, and moved off the public Internet? Internet video is the major driver of infrastructure investment, and the largest component of traffic on the Internet, much of which is being handled by CDNs. But does this traffic belong on the public Internet? Do we want an Internet optimized for TV? If so, is that still essential infrastructure? How might Internet differ if video-over-IP were not part of the public Internet? What are the technical, business, and policy implications of separating these networks from a regulatory perspective, and how could we get there from here, if we wanted to? This topic generated much lively discussion comparing different types of video on the network, much of which is amenable to quite dynamic broadcasting methods but is restricted to a linear streaming model for contractual (and copyright licensing) reasons. Perhaps not accepting the convergence of all video as inevitable, at least in the near to middle term, would lead to a better social policy outcome.

7. DISCLOSURE AND TRANSPARENCY UNDER THE OPEN INTERNET ORDER

Erin Kenneally talked about policy implications of the measurement and transparency requirements under the OIO [17], which had the strongest support from previous and likely future courts. She highlighted the range of disclosure and transparency (D&T) tools available to the FCC, and proposed a framework for D&T interventions (transparency requirements). One example is the MBA Transparency Report, which ISPs may use as a safe harbor for all the transparency requirements in the OIO. Another example is the measurement conditions on the AT&T/DirecTV merger. Each of these approaches has limitations, which the FCC acknowledges, but in combination they can go a long way toward addressing information asymmetries.

8. INSIGHTS AND LESSONS LEARNED

1. **QoE is a complex concept, only certain aspects of which are likely strongly linked to network QoS.** Although the workshop promoted new ways of thinking about measurement of QoE, one economist was pessimistic about the prospects for useful QoE measurement by academics, arguing that neither engineers nor psychologists would be able to measure how people feel about it or how much value they extract. The problem is highly non-linear with many unpredictable variables interacting simultaneously. The economics discipline has never figured out how to deal with it, and leaves it to marketing folks. Universities have institutional review boards to navigate when human subjects are involved. Only companies with millions of data samples from users who have already sacrificed their privacy in exchange for services could do such QoE measurement.
2. **Nonetheless, QoE research should and will continue.** Some researchers, and many companies, already do A/B testing by showing varying web page components to users and observing reactions. There has even been a study of using consumer grade EEG technology to measure the momentary state of the user [19]. Getting the most out of our information infrastructure will require cooperation between applications and networks, and synthesizing knowledge from a variety of measurement sources.
3. **Stakeholder perspectives on QoE measurement vary.** For some questions and stakeholders, a simple set of measurements may suffice. Much QoE research has high dimensionality but there may be value in mapping out a simpler abstraction.
4. **Transparency is the new opacity.** Complex product descriptions, bundled offers and other sources of artificial complexity are creating confusing user experiences, and challenges for measurement. Any time one comes up with better techniques to measure, companies may come up with better techniques to confuse regulators. The new consumer broadband nutrition labels were cited as an illustration of transparency that may do more harm than good by implying that consumers can really understand and compare services based on the labels. A revealing statement from the FCC about the complexity of financial reporting is that they have not been able to figure out how profitable telecom companies really are.
5. **The research community's ability to formally model interconnection dynamics is primitive.** A well-parameterized model requires understanding realistic ranges for inputs like cost of colocation and transit, incentives to localize traffic, and what constitutes evidence of market failure.
6. **Policy discourse requires expanded language.** Discussion of QoS tends to focus on fast and slow lanes, and resource scheduling, while the FCC must allow or constrain specific business transactions. For example, should the consumer be able to contract for a higher bandwidth virtual channel to Netflix? Do we see a future regulatory framework where Internet and specialized services coexist? In some sense, the Open Internet Order kicked the can down the road, and not only with respect to measurement and transparency, but these other policy questions. It was eye-opening to consider just how much streaming video is the driving force of telecom policy this century.
7. **Measurement methodologies and infrastructure must adapt and respond to research as well as technology advances.** Measurement of gigabit broadband networks will require managing expectations of regulators, users, and content providers. For

example, in a gigabit access network, even with uncongested interconnections a consumer may not be able to get their contracted bandwidth out of their broadband service. Additionally, the FCC needed a way to more nimbly integrate enhancements into MBA to give insight into path structure and dynamics, e.g., traceroutes in both directions. Such an enhancement would require changing the current testing methodology, at least with Netflix, since currently tests are rotated among servers like regular customer flows.

8. **Transparency and disclosure will require innovative approaches.** A recent comment to the FCC [10] imagined a scenario where an access provider might be asked to demonstrate to the FCC that their interconnects were not congested, or alternatively that any observable congestion could only be attributed to business decisions of the interconnecting party. Such demonstration might require disclosing business agreements to the FCC, similar to the conditions of the AT&T/DirecTV merger. It is fair to ask, for the health of the Internet, and to inform consumers, what measurements besides interconnection statistics should be required? A place to start is measurement tools to delineate home wireless network performance issues from those outside the home.

In closing, participants named their favorite emerging technology that might frame an agenda for a future WIE workshop, which included: DOCSIS 3.1, G.fast, 5G, information-centric networking architectures, universal encryption and its challenges for network management, and small cells.

9. WORKSHOP PARTICIPANTS

- Co-Host: kc claffy, (CAIDA/UC, San Diego)
- Co-Host: Dave Clark (MIT)
- Fred Baker (Cisco Systems)
- Steven Bauer (MIT)
- Roger Bohn (GPS/UCSD)
- David Clark (MIT/CSAIL)
- Lee Howard (Time Warner Cable)
- Walter Johnston (FCC)
- Scott Jordan (University of California Irvine)
- Erin Kenneally (DHS)
- William Lehr (MIT)
- Andrew Odlyzko (University of Minnesota)
- Michael Paradise (AT&T)
- Georgios Smaragdakis (MIT/TU Berlin)
- Tony Tauber (Comcast)
- Joel Watson (UC San Diego)
- Jason Weil (Time Warner Cable)
- Madura Wijewardena (Comcast)
- Patrick Zwickl (University of Vienna)
- kc claffy (CAIDA/UC San Diego)
- Amogh Dhamdhere (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-1414177. This report reflects the views of the authors and some workshop participants, but not necessarily the National Science Foundation.

10. REFERENCES

[1] https://en.wikipedia.org/wiki/Common_carrier.

[2] http://www.caida.org/outreach/workshops/aims/.

[3] https://www.measuringbroadbandamerica.com/.

[4] http://www.caida.org/data/routing/routeviews-prefix2as.xml.

[5] https://www.fcc.gov/document/fcc-unveils-consumer-broadband-labels.

[6] A. Morton and M. Bagnulo and T. Burbridge and P. Aitken and A. Akhter. RFC 7594: A Framework for Large-Scale Measurement of Broadband Performance (LMAP), 2015.

[7] M. Bagnulo, B. Claise, P. Eardley, A. Morton, and A. Akhter. Registry for Performance Metrics, March 2016. https://datatracker.ietf.org/doc/draft-ietf-ippm-metric-registry/.

[8] Bauer, Steven and Lehr, William. Improving the Measurement and Analysis of Gigabit Broadband Network Performance. In *Telecommunications Policy Research Conference*, September 2015.

[9] k. claffy and D. Clark. Adding Enhanced Services to the Internet: Lessons from History. In *Telecommunications Policy Research Conference (TPRC)*, Sep 2015.

[10] D. Clark, S. Bauer, and k. claffy. Approaches to transparency aimed at minimizing harm and maximizing investment, Sep 2014.

[11] Italo Cunha, Pietro Marchetta, Matt Calder, Yi-Ching Chiu, Bruno V. A. Machado, Antonio Pescapè, Vasileios Giotsas, Harsha V. Madhyastha, and Ethan Katz-Bassett. Sibyl: A practical internet route oracle. In *NSDI*. USENIX Association, March 2016.

[12] Federal Communications Commission. In the matter of protecting and promoting the open internet: Report and order on remand, declaratory ruling, and order, gn docket no. 14-28, 2015. https://apps.fcc.gov/edocs_public/attachmatch/FCC-15-24A1.pdf.

[13] V. Giotsas, M. Luckie, B. Huffaker, and k. claffy. Inferring Complex AS Relationships. In *IMC*, 2014.

[14] V. Giotsas, G. Smaragdakis, B. Huffaker, M. Luckie, and k. claffy. Mapping Peering Interconnections to a Facility. In *CoNEXT*, Dec 2015.

[15] Giotsas, V. and Dhamdhere, A. and claffy, k. Periscope: Unifying Looking Glass Querying. In *Passive and Active Network Measurement Workshop (PAM)*, Mar 2016.

[16] KC Claffy, Amogh Dhamdhere, David Clark, Steven Bauer. Report of AT&T Independent Measurement Expert Analysis of reporting requirements and proposals for measurement methods, December 2015.

[17] Lehr, William and Kenneally, Erin and Bauer, Steven. The Road to an Open Internet is Paved with Pragmatic Disclosure & Transparency Policies. In *Telecommunications Policy Research Conference (TPRC)*, 2015. http://ssrn.com/abstract=2587718.

[18] M. Bagnulo and T. Burbridge and S. Crawford and P. Eardley and A. Morton. RFC 7398: A Reference Path and Measurement Points for Large-Scale Measurement of Broadband Performance.

[19] T-Labs (Deutsche Telekom / TU Berlin), editor. *Quality of Experience: Advanced Concepts, Applications and Method*. Springer, 2015.