

Workshop on Internet Economics (WIE2012) Report

kc claffy
CAIDA/UC, San Diego
La Jolla, CA, USA
kc@caida.org

David Clark
MIT/CSAIL
Cambridge, MA, USA
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 It has NOT been peer reviewed.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 12-13 2012, CAIDA and the Massachusetts Institute of Technology (MIT) hosted the (invitation-only) 3rd interdisciplinary Workshop on Internet Economics (WIE) at the University of California's San Diego Supercomputer Center. The goal of this workshop series is to provide a forum for researchers, commercial Internet facilities and service providers, technologists, economists, theorists, policy makers, and other stakeholders to empirically inform current and emerging regulatory and policy debates. The theme for this year's workshop was "Definitions and Data". This report describes the discussions and presents relevant open research questions identified by participants. Slides presented at the workshop and a copy of this final report are available at [2].

Categories and Subject Descriptors

C.2.3 [Network Operations]: Network Management Public Networks; C.2.5 [Local and Wide-Area Networks]: Internet; J.4 [Social and Behavioral Sciences]: Economics

General Terms

Economics, Legal Aspects, Management

Keywords

Economics, Internet, Network management

1. INTRODUCTION

Building on the success of our first two workshops in this series [4, 5], we held the 3rd Workshop on Internet Economics (WIE). The theme for this year's workshop was "Definitions and Data", motivated by our sense that many of the debates today about effective regulation are clouded by lack of clarity about terms and concepts, and lack of real information about the current state of the communications infrastructure. Concepts that have resisted clean definition include network neutrality, reasonable network management, market power, and reliability. Stakeholders disagree on fundamental parameters of central concepts in the industry, such as interconnection, or the metrics for broadband quality itself.

Equally missing is good data on what is actually happening. Whether measurements are undertaken by the FCC, as with the current SamKnows effort, or by the research community or industry, good definition must precede good measurement, because collectively we must be consistent and clear what we are proposing to

measure and why. A guiding premise of this workshop was that attention to definitions can inform research in data gathering, which in turn can inform regulatory debate. Workshop discussions also focused on the impacts of the limitations of currently available data (such as undersampling) and how to gain more relevant data with minimal impact on personal privacy.

The workshop format focused discussion around six pre-selected topics: defining broadband (wired and wireless); Interconnection; definitions and metrics of market power; the emergence of private IP networks; regulatory distinctions in a converged world; and defining acceptable practice for data-gathering. We spent about two hours per topic, with at least two 10-minute talks followed by an hour for each discussion. Three promising future research directions emerged. First, we reached rough consensus on a proposed practical approach to measure a user's "quality of experience" (QoE), one that could frame not only a stable definition of broadband Internet service but also enable more rigorous description of "willingness to pay" for different applications. Second, most participants agreed that the rise of private IP networks as an alternative platform to the public Internet (and to the economically unsustainable PSTN) promise an even more opaque future at a time when it has become clear that much of current communications regulation lacks empirical basis. Third, there was recognition that both scientific research and sound public policy share the need to develop, maintain, and archive some classic data sets to develop some sense of history and to inform general models of network behavior. One possible goal for a future workshop is try to articulate an argument for data that might be valuable in the future, not only to support specific policy questions but also to begin to establish historical baselines and promote scientific inquiry. More formal and transparent ties between policymakers and researchers could frame ethical use of such data.

2. REGULATORY DISTINCTIONS AMID CONVERGENCE

We began with a fundamental and recurrent question: does the distinction between telecommunication and information service still make sense in a world where everything, including voice, is converging to use datagram networks for underlying transport. This regulatory artifact of public communications policy in the United States has yielded remarkable confusion when applied to the evolving Internet. Notably, this distinction does not exist in many countries, with important implications for regulation.

Barbara Cherry (Indiana University) introduced one problematic ramification of keeping this distinction at the core of policy debates, with her thought-provoking description of the concept of *shadow common carriage* [3]. In other industries the term "shadow" has referred to activities that have emerged or evolved in a form that

sits just outside a certain regulatory definition, so that they are not regulated in a way that more traditional variants were. Often such activities emerge as a direct consequence of deregulatory policies; as an example she cited the complex and highly leveraged instruments of the financial sector that arose after the deregulation of the banking industry. The problem with shadow activities is that they can impose a systemic risk on the industry. In the case of the financial sector, many rules intended to promote stability of the financial system were dismantled, leading to more unstable financial markets. Barbara described how the telecommunications industry has a similar type of shadow activity arising from deregulatory forces – specifically shadow common carriage, with VoIP its most obvious instantiation. As with shadow banking activities, these activities may carry substantial systemic risk, compounded in the telecommunications case by the legal gaps created between different bodies of law, which hinder policymakers’ ability to respond to problems.

A discussion followed about when the bright line regulatory distinction (between telecommunications and information service) began to fade, and how it could be clarified now. The U.S. Congress amended the Telecom Act (in 1996) to give the FCC explicit jurisdiction over the cable industry as a provider of information services, after which the FCC chose to classify cable broadband (and ultimately all broadband) as an information service, which meant it would be largely unregulated by the FCC. It is not clear what would happen if the FCC tried to reverse and reclassify broadband as a telecommunications service.

There was also debate as to whether the shadow activities were even an *unintended* consequence of the deregulation that led to them. With respect to VoIP, there was an explicit decision to let commercial VoIP services develop unhindered by regulation; one could argue the consequences were intended. But the increasing prevalence and variety of voice (VoIP) services triggers the obvious question: how do we regulate different forms of the same service, e.g. telephony, in a converged world? These questions are playing out in different ways in different places around the world, unfortunately without common definitions of key terms among people in different disciplines (engineers, economists, lawyers).

Shane Greenstein tried extending the financial sector metaphor, pointing out two critical issues that increased the system risk: reduced capital requirements, and mispricing of risk because of the auditor’s inability (or willful neglect) to examine a given instrument enough to analyze it. Does the absence of outage reporting requirements, for example, of the national network infrastructure put the system at risk?¹ Barbara’s conclusion from her comparative research is that just as markets need underlying regulation (e.g., property rights, contract law) to function effectively, there appears to be a similar requirement for a specific set of regulations for networks to emerge with properties we associate with effective function, namely that they be widely available, affordable, and reliable. Such regulations are the kind of rules that we generally associate with common carriage: universal service on non-discriminatory terms at reasonable prices. This conclusion did not have consensus among participants; another opinion was that a sector-specific Internet regulator may not even be needed – that consumer protection could be handled by the Federal Trade Commission, reliability by the Department of Homeland Security, etc.

Rob Frieden (Penn State University) summarized recent work analyzing costs and benefits associated with discontinuation of com-

mon carrier, wireline voice telephone service [14]. While common carriage imposes burdens on the carriers, it also brings benefits that would not be automatically extended with the deregulation of the public-switched telephone network (PSTN): rights-of-way, spectrum and pole access, tax benefits, interconnection rights, universal service funding. At one extreme, legacy carriers might be reclassified as VoIP carriers, which would include several costly regulatory obligations (911, CALEA, number portability, outage reporting), but perhaps without some of the benefits they enjoyed as PSTN common carriers. At the other extreme, they could be reclassified as information service providers, which would allow them to leverage access to their networks as well as charge for higher service qualities, although at risk of triggering regulatory attention to interconnection terms and consumer protection. Rob presented three case studies of disputes (Comcast/Level3, Fox-Hulu/Cablevision, and GoogleVoice/ATT) that might have triggered substantial regulatory attention had parties not resolved them on their own.

Scott Jordan (UC Irvine, with recent experience on sabbatical as a Congressional staffer trying to help policymakers understand Internet technology) gave us a primer on how regulatory models map to layers of the Internet architecture and implications for policy [15]. The networking community sees service offerings in terms of layers (physical infrastructure (1), IP (2), TCP (3), Applications (4), Content (5)).² In the Internet, barriers to entry for new services tend to be low at layers above IP and high at the IP layers and below (physical infrastructure). So, for example, in order for Skype (layers 3-5) to compete with Cox (layers 1-5) at delivering voice with service quality guarantees, Skype needs access to features of layers 1-2, such as QoS, which Cox does not currently allow. Scott’s conclusion from his own research is that the important demarcation point is between (what is at and below) layer 2 (IP) and (what is at and above) layer 3 (TCP). This “bright” line delineates where the business activity is, regardless of technology (copper, wireless, fiber), and clarifies regulatory debate. For example, universal service should be about subsidizing only the lower layers (1-2). He believed that creating policy that respects this delineation would promote innovation at higher layers in the industry by preventing vertical foreclosure, i.e., when networks close some services off to competitors. He also suggested that the distinction between telecommunication and information service would remain useful if updated to define telecommunication as layers 1-2 and information service as layer 3-5.

Different regulatory regimes have tried to place this demarcation at different layers. Many countries (including the U.S. in the past) have put the cut point just above the physical layer, and required owners of facilities (typically copper pair telephone circuits) to make their wires available to competitors. This regime made sense when the layering model was telephony on top of copper, but does not cleanly apply to different technology (e.g. the HFC facilities of the cable industry) and higher level service models. Putting the demarcation at a higher logical level, e.g., above the IP layer, provides some insulation from variations in technology.

Scott’s talk led to a discussion of innovation and how capital flow supports it. Innovation has been greatest at the application layer, although layer 1 has experienced plenty of innovation too, e.g., the DOCSIS3.0 protocol, cellular performance. Innovation at the lower layers, especially in access networks, has been driven by direct payment from users purchasing Internet access; funds to support innovation in lower layers flows only indirectly from higher layers of the stack, since providers of applications and information do not normally pass substantial payments to broadband access providers.

²Scott noted the difference from the OSI layers, where IP is layer 3, TCP layer 4, applications layers 5-7 and content is layer 7.

¹In February 2012 the FCC extended some outage reporting requirements to interconnected VoIP service providers, but not to broadband Internet services [9]. It required reporting only complete outages, and deferred action on measurement or reporting of any performance degradation for either VoIP or broadband services.

3. DEFINING BROADBAND

This session began with three talks to spur discussion on how to develop useful, flexible (across media) and stable (across time) definitions of broadband. Bill Lehr (MIT) focused on the challenges in defining and measuring mobile broadband, especially since the nature of mobility environments is rapidly changing and increasingly opaque. Since the distinction between wired and wireless is only at Layer 1, the bigger question is how to define the Internet (or broadband) for statutory purposes when telephone, cell, and cable infrastructure have all merged into one network? This discussion related to Scott's point that the correct interface at which to focus regulatory concern is the one above the IP layer.

Gabor Molnar (CU Boulder) presented his two-step approach to analyzing empirical data on the U.S. broadband market structure and service quality. His first step was to examine economic factors that determine market entry, including number of existing providers and economic factors in the region. His second step was to assess whether the number of wireless and wireline providers affect service quality. He found wireline and wireless providers have similar entry patterns, and imperfect substitutes for each other. He found a positive correlation between competition (number of providers) and quality of wireline access, with the biggest jump in quality occurring at the jump from 1 to 2 providers.³ A more multi-dimensional approach to measuring quality, would include pricing information and the ability to cross-reference existing databases of performance data, e.g., M-Lab, Ookla, Netflix.

Andrea Soppera (BT Labs) presented a survey-based approach to measuring broadband customer satisfaction, recognizing that customers value quality of experience (of application and performance) as much as access speed and low latency. Through his measurements and surveys of 139 individual customers, he confirmed that metrics for quality of experience vary with the application. For specific applications like Skype, he can map technical parameters of QoS to a subjective perception of quality of experience (QoE). Different parts of the network will have different QoS, but for voice applications a user typically has a few other endpoints he communicates with, so a provider could combine the mapping from QoS to QoE with a customer's call graph to tailor a more realistic estimate of expected performance than a general service guarantee.

There was agreement on the need for standard methods and policies that define what dimensions of the "broadband experience" we should be measuring, and how to aggregate those measurements into simple meaningful indices that reflect real customer experiences and inform customer decisions. We are reaching a point where 100MB/s download speed does not matter as much as monthly download (or upload) capacity measured in terabytes (which makes the distinction between wired and wireless obvious). Regulators could force ISPs to be more transparent about not only data rate but also about latency and jitter which matter to QoE perceptions. But even data rate is challenging to measure, and since some progress has recently been made on measuring this metric, e.g., SamKnows, it will probably be used as a proxy for a while. Engineers and network researchers could certainly contribute to a developing framework for measuring quality of experience (QoE) that could influence how economists, lawyers and regulators think about the issues in the future.

4. DEFINING MARKET POWER

We next moved to methods for describing and assessing market power in today's Internet. Elisabeth Maida (Akamai) presented her

³Shane noted that the national broadband plan claimed that the biggest jump was from 2 to 3 providers.

MIT masters thesis [8] research on assessing market power of networks. She first undertook a historical analysis focused on limited network access competition, including consumer pricing and likely harmful discrimination against traffic or users. More recently there has been increasing focus on the role of interconnection agreements (and control over them) in market power. She built on an idea examined by D'Ignazio and Giovannetti [7] called betweenness (centrality) of an Autonomous System (AS): the number of BGP paths between two other ASes that traverse that AS. The metric inherently does not apply for access networks, since they are at the edge of the topology, so she developed a new potential metric that she called access variance, to capture the variance in how content is delivered to an access network's end users. One key challenge of this metric is quantifying the viability of paths, since performance varies considerably across them. She listed several network characteristics that could help quantify this variability, most of which require network provider cooperation (network type, link capacities, traffic profile).

Shane Greenstein (Northwestern) offered some lessons from his recent role at the National Bureau of Economic Research (NBER) studying "digital dark matter" – economic activity in the digital economy that standard GDP measurement fails to capture. Through the U.S. census and Bureau of Labor Statistics survey activities, we have reasonable data on about two-thirds of the U.S. economy, but the Census Bureau has not expanded its scope to cover the contribution of the Internet to the economy. GDP measures reflect the revenue affiliated with Internet access, but not the benefits that arise from the existence of multi-sided platforms in the digital economy except for their advertising revenue. Shane suggests the prevalence of digital dark matter likely renders traditional approaches to measuring GDP and market power inadequate.

Shane also discussed some of the issues that relate to assessing market power in an industry, which requires not only identifying appropriate data to use to measure market share, but (first) defining the correct market and its scope. For merger reviews, a recognized metric is called SSNIP (small significant non-transitory increase in price). The SSNIP metric is used to gauge whether two products are in the same market by asking whether a price increase in one product (of perhaps 5%) would cause consumers to switch to the other product. If consumers would switch, the two products are considered to be in the same market. Thus one could ask whether a SSNIP in the price of wireline broadband would cause people to switch to wireless, and conversely whether a SSNIP in wireless would cause people to switch to wireline access. Based on this sort of analysis, AT&T lost the T-Mobile (horizontal) merger based on the presumption that mobile was the only market relevant to the merger review. On the other hand, Comcast and NBC was a vertical merger with potential for concentrated distribution; in this case the market for distribution was defined as not including wireless.

Andrew Odlyzko (U. of Minnesota) offered some estimates of the value of different slices of what Shane referred to as the digital economy, and their implications for policymakers. The economics of the network have changed dramatically in the last two decades, with long distance and switching costs much lower than they used to be, and only access costs remaining high. As baselines, world telecommunications service revenues are almost \$2 trillion annually; world advertising revenues \$500B; Google's revenue (in 2011) \$38B; and Hollywood revenues about \$100B. The total income of the cable networks is about \$97.6B; \$30B from advertising and \$70B from consumers, confirming Andrew's belief that "content is not king; communication is still king." Over half of the world's telecom revenues are now from wireless (mostly from voice and texting), although various estimates indicate that only 4%

of global traffic goes over radio. U.S. wireless network revenue is closer to \$150-170B, reflecting higher value for much lower data volume than the cable industry. In terms of traffic (and the architecture required to support it), content providers are primary, but most of the actual value (what generates revenue) in the network is much lower volume content (advertising, wireless service).

John Chuang (UC Berkeley) re-introduced us (with his talk from WIE2011 [1], upon request) to the concept of a two-sided bilateral oligopoly, i.e., where a set of access networks, each representing a terminating access monopoly with respect to their customers, are connected to a set of powerful content providers, each with monopoly control over their respective sets of content [12]. The problem with the recently popular two-sided market model in describing the current Internet is that it assumes vertical integration between access and CDN providers to form the two-sided platform between the two ends. Such vertical integration is not common today. The two-sided bilateral oligopoly model can better accommodate actual industry structure today. Important characteristics of a bilateral oligopoly include: (1) negotiation and bargaining leading to long-term contracts that create entry barriers for both fields of business; (2) an appropriate balance of power across the two sides. These characteristics can result in lower consumer prices and increased consumer surplus. He offered two other example industries where the presence of a bilateral oligopoly helped counteract the market power of terminating access monopolies: the hospital market and the health insurer markets, and the music distribution and record label markets. His main conclusion was that market power is not only due to concentration in a given market, but also influenced by the level of competition in adjoining markets.

Milo Medin pointed out that today most people still use multiple CDNs as well as transit providers, so CDNs do not have significant leverage. John acknowledged that CDNs do not have a lot of market power vis-a-viz access networks, but can still exercise market power between gateways on each side. Many content providers may choose to multihome to CDNs, but the industry could still be quite concentrated. One needs to clarify for what purpose market power is being discussed: e.g. abuse or merger review. The relationship between CDNs and networks is also shifting. Networks must analyze the cost and benefits of deploying additional equipment vs. subscribing to a CDN. Lately more networks are trying to charge the CDNs (e.g. Akamai) for access to end users.

We spent time identifying specific questions whose answers would inform the discussion of market power as well as the larger discourse on broadband. There were repeated calls for a clearer articulation of specific questions that needed answers, since some performance-related data can be (and already is being) gathered with limited (or no) cooperation from industry, e.g., MLab, Speedtest, Netflix's performance measurements. Policy researchers noted that M-Lab is sharing large chunks of their performance data although it was not curated in a way that they could make use of it. But traffic data, fundamental to most questions of market power and interconnection, remain tightly guarded by players with powerful economic incentives to not disclose it. Changing laws could make it easier for providers to provide relevant information, but it may take a crisis to inspire change since telecom policy has been on a deregulatory trajectory for decades, funded by heavy investment into lobbying by incumbent providers.

5. INTERCONNECTION

The patterns and characteristics of ISP interconnection are poorly understood. We discussed what kind of information about interconnection would inform future policy debates, whether (and how) we should track the capacity and utilization of Internet intercon-

nections, and how traffic data might inform debates about payment patterns for interconnection.

Tony Tauber (Comcast) gave his perspective on network engineering, where network growth planning is based mainly on internal traffic data, and connectivity is mainly focused on connecting users to content. While mobility (user movement) represents microscopic changes to the network, content movement can now represent macroscopic change. Specifically, the traffic and interconnection landscape is now dominated by large content providers and CDNs, who source large amounts of traffic (hyper-giants), using BGP, DNS, and geolocation technology as inputs to dynamically adjust traffic load and path (overlay routing). CDNs' ability to dynamically adjust traffic load across different paths in a matter of minutes complicates the task of network management for other network service providers, and renders previous empirical projection methods inapplicable.

Matthew Luckie (CAIDA/UCSD) described the notion of customer cone as a metric of influence ASes have on the global Internet routing system, and the algorithm they use to compute it. An AS's *customer cone* is the set of ASes that can be reached from this AS using only provider-to-customer links. His talk generated discussion on how closely BGP-topological rankings might reflect traffic or transit revenue rankings.

Jonathan Liebenau and Silvia Elaluf-Calderwood (London School of Economics), presented some work trying to understand the characteristics of quantitative data on aggregate Internet traffic. There are differences between the U.S. and Europe Internet industries, both in business models (for peering and transit) and how regulators and consumers view issues such as pricing and network neutrality. However, there are even more important trends shared across both regions. For example, although it is nearly impossible to get traffic data, aggregated statistics posted by exchange point operators suggest that while traffic and bandwidth usage is growing globally, annual rates of traffic growth have been dropping. But perhaps more importantly, traffic is still growing much faster than revenue [16], for a number of reasons. The cost of transit or associated capital expense to deliver traffic is being driven down for a number of reasons. In addition to the raw effects of Moore's law driving performance up for the same prices, the proliferation of Internet exchanges (IXes), or more precisely the alternative of private peering at an IX, has driven transit prices down. However, the growth in traffic means that while operators are getting much more capacity than they used to for the same price, they are still paying about the same amount overall. There was discussion of the implications for regulators if network operators move toward an investment model of maintenance rather than growth, and related interest in understanding the relationship between EU investments and how things operate in Europe.

6. THE EMERGENCE OF PRIVATE IP NETWORKS

In the face of continued decline of PSTN subscription and revenues, private IP networks that provide basic VoIP telephone service (and other services) have emerged that may warrant measurement and definition. David Clark (MIT) led a discussion of how to approach this challenge. The FCC has introduced the concept of "specialized services": IP-based services offered by a facilities owner using the same physical infrastructure as the public Internet, but distinct from the Internet in some way. VoIP is an example of a specialized service, and when provided by the facilities owner as a specialized service may offer a better quality of experience than an over-the-top voice service like Skype because the operator of

the infrastructure can enhance the QoS in ways that are not offered on the public Internet. The FCC's Open Internet Report and Order [10] attempts to define the distinction between a specialized service and a service covered by the Order, but the definition is not well-crafted, and the FCC Open Internet Advisory Committee (which David chairs) has struggled to agree on a working definition of specialized services. The lack of clarity is illustrated by debate over a recent Comcast offering: an on-demand IPTV service that uses a customer's Xbox as an alternative to the traditional cable set-top box to deliver video. Comcast asserts that this should be viewed as a specialized service, but some critics have claimed that in fact it is a service running on top of the public Internet and should be subject to the rules of the Open Internet Order.

In the meantime, specialized services fall under Barbara's earlier definition of *shadow service*.⁴ Such private services raise concerns regarding services available to consumers over the public internet, in particular whether the public Internet can sustain its role as a platform to encourage third-party innovation. If the public Internet enables sufficient innovation, there may be no reason to worry about the proliferation of private IP-based networks. But should private IP-based services that leak onto the public internet be covered by the FCC's network neutrality guidelines [10]?

Referring back to Scott's discussion of layering, the EU takes the approach of lower-level unbundling at the physical layer, which has different implications for an operator that might like to sell consumer-facing services that run over IP but not over the public Internet. Since the point of demarcation is under the IP layer, rather than on top of it, the facilities owner cannot implement private IP services at the same time that a competitor uses the copper pair to offer Internet access. Barbara tied this distinction to historical notions of common versus private carriage, which established whether tort obligations applied to a given entity/individual. But the obligations only applied in retail relationships: carrier to customer. Statutory law (industry-specific or antitrust) had to further develop what to do in wholesale relationships among providers. Issues that will arise in the context of private IP networks and specialized services have yet to be well-articulated, and represent important research challenges going forward.

7. ACCEPTABLE PRACTICES FOR DATA-GATHERING

Alissa Cooper (Center for Democracy & Technology) gave a thought-provoking example of all the data already being collected by (in her example cellular) operators for operational and marketing purposes, the collection of which most consumers cannot opt-out of. Examples from a popular operator's terms of service (ToS) include: website addresses visited, location of device, use of applications and features, data and calling features, device type, and demographic categories. They then package up this data and sell it to interested parties as a product. She sees a window of opportunity where researchers could make the case that access to some of this data already being collected could serve some social good in researchers hands (like Google's Flu Trends), as opposed to only serving the carriers and their advertising partners.

Erin Kenneally (CAIDA/UCSD) then led a discussion on how to formalize some of the data access/indexing methods. She agreed that the research community could do better job of articulating the benefits of engaging researchers as stakeholders in addressing issues central to Internet regulation. She also saw a gap between

⁴Although the FCC does require some limited reporting for interconnected VoIP services outages [9], they retracted an initial plan to require reporting for other broadband Internet services.

the "data-rich" and "data-poor", amplified by the current ad hoc, opaque, and unscalable approach to data sharing: providers tend to give data only to researchers they already know and trust, since the risks of sharing are perceived as greatly outweighing potential benefits. Ironically, those same providers are eager to sell data to marketing or advertising entities. She advocated a more standardized yet flexible data exchange framework that addresses data sensitivity risks in parallel with intended utility needs [13, 6]. Although this framework is designed for those already incented to share data, such interested but risk-averse providers can use it to identify, communicate and defensibly demonstrate their sharing decisions. Further, the framework generalizes across various data sharing scenarios. The approach includes assessment of risks and benefits of the data exchange, and a set of tunable operational disclosure controls that can accommodate a provider's risk profile. Following these presentations, there was a group discussion about data gathering in support of specific research goals, and how the empirical research community could contribute to the ongoing debates. Our most productive example was the possibility of redefining our current understanding (and practice) concerning broadband service quality using (QoE) measures (per service, and given a composition of common services) that could map directly to more technical performance parameters, using a parameter space in which technical factors can be traded off to optimize QoE at least cost.

8. FUTURE RESEARCH DIRECTIONS

Thematically, three high-level directions emerged as particularly important: QoE measurement, private IP services, and baseline data collection to support research.

We achieved some consensus on a method to measure "quality of experience", a consumer-facing metric that could be useful to both consumers and network engineers, and even frame a standard parameterization of broadband Internet service. Instead of starting with technical parameters such as jitter or average bandwidth [11], the proposed approach would select a basket of perhaps ten popular applications for residential Internet users, and then determine what performance characteristics would be required to adequately serve that basket of applications without noticeable impairment to them. For any given Internet service, one could compute a "QoE score" for each application, capturing the degree of impairment, and then compute an overall figure of merit by combining the QoE scores of the given Internet service for a blend of these different applications.

This "QoE score" could also be used to frame "willingness to pay" for different applications. For web-browsing, the willingness to pay might be a function of download time for a typical web page, but for a videoconferencing application, willingness to pay would more likely be a function of percent of packets received within the time frame for a group-of-pictures. Assuming that willingness to pay is a function of the QoE score, one could work backwards to compute a willingness to pay as a function of technical performance statistics, e.g., mean, variance, tail probability of parameters such as e2e delay, loss, throughput. Like the overall figure of merit, these estimates of willingness to pay could be aggregated across the basket of applications, as a function of time devoted (say, per month) to each application. The research challenge is how to derive an evolving standard definition (performance parameters to support a given basket-of-popular-apps) that constitute a baseline ("universal") Internet service. The operational challenge is how to allow consumers to pick among services based on their own anticipated usage patterns, and derive an overall figure of broadband performance based on the blend of application QoE they want.

The second recurrent theme of the workshop was the rise of private IP services as an alternative platform to the public Inter-

net, which is likely to be a growing issue for regulators, service providers and application providers. It will require new theories of regulation (perhaps causing regulation to return to a more "layered" approach). It will also raise many issues for the empirical research community, because these networks are not open platforms on which third-party measurements can easily be performed, but restricted platforms on which only specific higher-level services can be run. This reality will lead to an even more opaque future at a time when sound regulation (especially market-based self-regulation) demands more data and more transparency.

Finally, and related to recent calls for ex post (measurement-based) rather than ex ante (rules-in-advance) oversight, there was recognition of the need to develop, maintain, and archive some classic data sets (perhaps "day in the life" snapshots), to develop some sense of history, and to inform general models of microscopic and macroscopic behavior. David Clark noted the benefit of identifying classes of data that have enough utility that we will want them in 5 years, even though we do not know what questions to ask of such data now. He reported a conversation with a scientist from another discipline who said to him that: "Good data outlasts bad theory." His point was that data gathered for one purpose, or based on an informed intuition of its utility, often proves useful later, including in unanticipated ways. Therefore, baseline data gathering should not be gated by a specific question. On the other hand, for data that might be personally identifiable, some privacy policies (specifically in Europe) require that data be used only for the declared purpose, a well-motivated restriction that limits research use of the data. However, there is no reason for network operators to be motivated to gather and archive data without an immediate benefit to them, a reality that suggests that either public funding of third-party data gathering or regulation of certain data reporting will be required to develop this sort of historical baseline. The research community needs to develop and articulate a strong argument for the sorts of data that might be valuable in the future.

A future workshop could try to distill policy research questions that people are attempting to answer, data sources they have looked at, gaps in the available data, and their ideal data. Such a workshop should be structured to elicit discussion of data sources that researchers may not know are already available, as well as proposals for responsible access to data, both to inform specific policy discussions as well as enable more general access, in many cases for not yet defined purposes. More formal and transparent ties between policymaking agencies and the research community could support a framework to safeguard ethical use of such data.

9. WORKSHOP PARTICIPANTS

- Co-Host: kc claffy, (CAIDA/UC, San Diego)
- Co-host: Dave Clark (MIT)
- Rebecca Arbogast (Comcast)
- Silvia Elaluf-Calderwood (London School of Economics)
- John Chuang (UC Berkeley)
- Barbara Cherry (Indiana University)
- Alissa Cooper (Center for Democracy & Technology)
- Rob Frieden (Penn State University)
- Shane Greenstein (Northwestern University)
- Scott Jordan (UC Irvine)
- William Lehr (MIT)
- Jonathan Liebenau (London School of Economics)
- Elisabeth Maida (Akamai Technologies)
- Milo Medin (Google)
- Gabor Molnar (UC, Boulder)
- Andrew Odlyzko (U. Minnesota)

- Andrea Soppera (BT Labs)
- Jesse Sowell (MIT)
- Tony Tauber (Comcast)
- Nicholas Weaver (ICSI)
- Erin Kenneally (CAIDA/UCSD)
- Matthew Luckie (CAIDA/UCSD)
- Brad Huffaker (CAIDA)

ACKNOWLEDGMENTS. Thanks to Barbara Cherry for contributing her notes for this workshop report, and thanks to all participants for their insights. The workshop was supported by CAIDA members, and by the National Science Foundation's Computing and Networking Systems Division CNS-1017064 "The economics of transit and peering interconnections in the Internet." This report reflects the views of the authors and most workshop participants, but not necessarily the National Science Foundation.

10. REFERENCES

- [1] Workshop on Internet Economics (WIE), 2011. <http://www.caida.org/workshops/wie/1112/>.
- [2] Workshop on Internet Economics (WIE), 2012. <http://www.caida.org/workshops/wie/1212/>.
- [3] Barbara Cherry. Radical experimentation under deregulatory broadband policies: The rise of shadow common carriers. June 2011.
- [4] k. claffy. Workshop on Internet Economics (WIE2009) Report. *ACM SIGCOMM Computer Communication Review (CCR)*, 40(2), Apr 2010.
- [5] k. claffy. Workshop on Internet Economics (WIE2011) Report. *ACM SIGCOMM Computer Communication Review (CCR)*, 42(2), Apr 2012.
- [6] Coull, Scott E. and Kenneally, Erin E. A Qualitative Risk Assessment Framework for Sharing Computer Network Data. In *Telecommunications Policy Research Conference*, Sept 2012.
- [7] Alessio D'Ignazio and Emanuele Giovannetti. "Unfair" Discrimination in Two-sided Peering? Evidence from LINX. *Cambridge Working Papers in Economics*, 0621, 2006.
- [8] Elisabeth Maida. Assessing Access Network Market Power, Masters Thesis, MIT, 2011.
- [9] Federal Communications Commission. FCC Adopts Outage Reporting Rule for Interconnected VoIP Services, Feb 2012.
- [10] Federal Communications Commission. Preserving the Open Internet, Broadband Industry Practices, Report and Order, 2012.
- [11] J. Klensin. RFC 4084. Terminology for Describing Internet Connectivity, 2005.
- [12] John Chuang. Loci of Competition for Future Internet Architectures. *IEEE Communications Magazine*, July 2011.
- [13] Erin Kenneally and Kimberly Claffy. Dialing privacy and utility: a proposed data-sharing framework to advance Internet research. *IEEE Security and Privacy*, July 2010.
- [14] Rob Frieden. The mixed blessing of a deregulatory endpoint for the public switched telephone network. *Telecommunications Policy*, July 2012.
- [15] Scott Jordan. A Layered Network Approach to Network Neutrality. *International Journal of Communication*, 2007.
- [16] Scott Marcus. Network Operators and Content Providers: Who Bears the Cost? 2011.