# Workshop on Internet Economics (WIE2017) Final Report

kc claffy
UCSD/CAIDA
kc@caida.org

Geoff Huston
APNIC
gih@apnic.net

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 It has NOT been peer reviewed.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 December 13-14 2017, CAIDA hosted the 8th 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. The FCC's expected decision (released during the workshop, on 14 December 2017) – to repeal the 2015 classification of broadband Internet access service as a telecommunications (common carrier) service – set the stage for vigorous discussion on what type of data can inform debate, development, and empirical evaluation of public policies we will need for Internet services in the future.

`http://www.caida.org/workshops/wie/1712/.`

## 1. WHITHER THE PUBLIC INTERNET

We knew going into this workshop that it would involve vigorous exchange of diverse opinions. We also knew the expected U.S. FCC announcement regarding the reclassification of broadband Internet access service as an information (non-common carrier) service offered an appropriate context for our theme of how to think about regulating telecommunications from a "clean slate" perspective, i.e., not having to fit segments of the ecosystem into decades old classifications created before the Internet existed. The primary goal of the workshop was to bring together different perspectives on how we build, operate, use, and think about the Internet as a communications fabric, and examine these perspectives through an economic lens, and where possible, an empirical lens. There was substantial discussion of a post-network-neutrality network and the risks and benefits of such a network in the evolving political economy.

We started with talks about what forces are shaping the Internet today, and how evolution of those forces has yielded a world in which a few economic actors capture a large share of economic rents. The revenue picture in the Internet carriage provider network has been relatively steady for over a decade now, but peak traffic demands have grown tremendously, leaving some stark economics for broadband deployment. Some ISPs are seeing a relatively consistent growth of peak traffic demands within their network, the bulk of which is per-user content-based growth, generally in households where access capacity allows content streaming.

Enabling much of this traffic growth has been the rise and concentration of massive content distribution networks (CDNs) and e-commerce platform operators. These players capture large market share by exploiting economies of scale, network externalities and high switching costs (network effects) that inherently leads to winner-take-all industry actors, e.g., Airbnb, Amazon, Facebook, Google, LinkedIn, Twitter, or Netflix. In the case of advertising-based platforms, massive subscriber populations generate massive data that helps these companies capture large shares of advertising revenue in a market, making it possible to fine-tune their data analysis internally and fund acquisition of existing, or potential competitors.

More interesting from a regulatory perspective, CDN operators are evolving away from using public Internet transit services to carry content to users, instead establishing dedicated communications capacity to handle the enormous traffic levels and growth. These private interconnections open up the possibility of reconceptualizing, and re-architecting, them as something other than public Internet services – what the FCC Open Internet Order classified as "specialized services" that would remain outside the oversight of the then-proposed Title II provisions for general Internet access services. (Under the 2018 repeal of the Title II for broadband access services, there is no FCC oversight of any interconnections, dedicated or shared.)

An underlying motivator for these large CDNs is their interest in avoiding persistent choke points (near-monopoly control over access to a facility) – in this case the access ISPs. Choke points are a potential signal of market failure, typically accompanied by falling levels of real investment, and often associated with a masquerade of confusology taking the place of data and facts. There was consensus among participants that today's telecommunications market is dominated by choke points, not only for network access itself, but search engines, auction sites, and app stores. The proliferation of these choke points at different layers of the ecosystem is again a reflection of natural network effects in an increasingly connected economy. The open question – and one resistant to study given the lack of available empirical data – is whether exertion of market power via such choke points is harming other areas of economic activity.

Another tension lies between application/service behavior vs. network operators themselves. A growing number, and already some of the largest, content providers deliberately obfuscate both content and signaling information from network operators providing transit for the traffic. Consistent with supporting user privacy but also consistent with preserving exclusive access to economic opportunities in monetization of user traffic, this edge provider strategy inhibits (or prohibits entirely) the use of network middleware technologies to improve traffic engineering, police (or secure) usage, and improve their own services. Increasingly, the access and transit network operators have less insight as to the nature of the traffic, and fewer effective traffic management tools.

The wireless realm is experiencing similar ecosystem dynamics but with more confounding (perhaps hype-induced) finances. U.S. mobile service providers are increasingly investing outside of their core competence to generate market growth. For example, T-Mobile has recently purchased Layer 3 TV (https://bloom.bg/2zyGbFS), Verizon has purchased the combination of Yahoo and AOL (called, curiously, "OATH"), and AT&T's efforts to acquire Time Warner continue through the courts. The once lucrative Average Revenue Per User (ARPU) levels are falling, and the only way mobile providers can raise revenue levels is through extensive bundling of the product with content.

The fundamental result of this convergence between content and carriage is that what we used to conceptualize as a public communications network is in the process of being thoroughly privatized, with little visibility into the fiscal and social costs of resulting infrastructure and service models. The disruptive changes that this rapid deployment of information technologies is bringing to our society are similar in scope of impact to those that occurred during the industrial revolution. There was consensus at the meeting that we are experiencing a time of uncertainty, where it is appropriate to ask some fundamental questions about the purpose of regulation, how to apply it, and the ever-present question in public policy and science: how can we tell the difference between facts and suppositions?

## 2. DATA: MINING THE NEW OIL

If the Internet is heading down the same path as newspapers, where advertising is a substantial source of funding for content and services, what happens when advertising revenue reaches its growth limits? For example, today's global GDP is about $80T. Annual telecommunications revenue is $\approx$ $2T. Total expenditure on advertising is $\approx$ $500B. In 2016 Google reported a revenue level of $90B, the same order of Hollywood annual revenues worldwide ($100B). These numbers suggest that in revenue terms, perhaps content is not as significant as telecommunications access service itself. On the other hand, most of the world's largest market capitalization companies are in the content business: Google, Microsoft, Amazon, Tencent, Facebook, Alibaba. Geoff Huston referred to the current epoch as a new gilded age, similar to the U.S. in the 1890s, where a small number of players held most of the economic returns: Carnegie, J.P. Morgan. The relative level of absolute market dominance by these players in an emerging global market and the economics of that time also kept the U.S. in a role of economic world leadership for the next 100 years. Whether intending to or not, today's platform companies are building an economic architecture for the 21st century. Some noted also that the size of these dominant companies is much smaller today, 17K employees at Google vs. 60K at Standard Oil in 1909 (or 340K at U.S. Steel in 1943).

The daunting and insidious difference between traditional media advertising and Internet advertising is the effect on individual privacy, and consumers' lack of any realistic valuation on the personal information they exchange for access to online goods and services. The current asymmetry of information on privacy and its monetization poses a challenge for regulators that the U.S. has not yet attempted to tackle. The Europeans are leading the way with the impending General Data Protection Directive, the implementation and impacts of which remain to be seen.

## 3. REGULATION OF WHAT?

While no one believed that prospects were good for major reform, everyone agreed that the current legislative regime was ill-matched to current industry needs as well as consumer protection needs. Some argued that current models of economic regulation (competition policy) could manage market power concerns of the new platform companies. But given the continuing up-ending of the ecosystem, even our models of how we regulate merit re-examination. More specifically, the U.S. telecommunications regulator normally looks at the service, not the network facilities used to deliver the service. For example, the U.S. FCC defines a "Broadband Internet Access Service" as a service that sends and receives Internet packets within certain service parameters. The current US Telecommunications Act[1] uses different titles to classify services, but these days almost every service runs as an application over IP. The convergence of TV/telephony/data over IP begs the question of what we would consider a reasonable universal service obligation.

## 4. REGULATION OF TRANSPARENCY

Building on the FCC's recent commitment (some say weak) to transparency, we entertained the metaphor of a *Bureau of Internet statistics*, a federal open data management agency where the only regulations would involve incentivized transparency. Many in the room emphasized that transparency was not sufficient to discipline the market, especially in today's world, in the words of one participant, where "Transparency is the new opacity." Another historical comparison to the 19th century gilded age: the first serious engagement of the U.S. federal government to regulate price discrimination and other common carriage aspects of the railroads, took decades of agitation from small businessmen (largely farmers) before leading to the establishment of the Interstate Commerce Commission (ICC) in 1887. Even after that the ICC had trouble enforcing anything due to weakness of the Congressional legislation that created it.

In today's hyper-connected world, one could argue that every agency has an open data role – the fundamental problem in the broadband ecosystem is that transparency is not effective, and the FCC's commitment to transparency is diminishing: the 2017 Order reduced transparency with respect to pricing, what consumers care about most. Research on mandated transparency disclosures after the 2015 order found that they were buried too deep on web sites for people

[1]https://www.fcc.gov/general/telecommunications-act-1996

to find them, among other barriers to utility. We also know that privacy disclosures on web sites do not leave consumers generally informed. If a data disclosure model were likely to work at all for network management or data mining (serious doubts in the room), it would require public interest firms processing them and making lots of noise.

However the FCC chooses to respond to third-party measurements that demonstrate harm (e.g., Comcast BitTorrent resets in 2007) some form of transparency is required to enable ex post audit enforcement mechanisms. That is, most enforcement problems are first discovery problems. And the discovery problem is not trivial for most policy-relevant issues. Measurement and data analytics require infrastructure (hardware and software), computational tools and skills that are rare in regulatory agencies. The coupling of (e.g., two-sided) markets requires data analysis cognizant of both upstream relationships and the larger economy and downstream consumer impacts. Further complicating the issue is the "innovation" in automation of public filings which effectively enables denial-of-service attacks on the FCC filing system by opponents of FCC actions. This system received over 22 million public comment filings in response to its latest proposed reconsideration of Title II provisions. This episode demonstrated the infeasibility of any manual processing of these filings to capture and represent a thorough understanding of the range of responses on often subtle issues, and more generally points to the challenge of managing and distilling crowd-sourced regulatory debates today. As WIE participants have avowed before, the FCC (and related regulators) need more data analytic and information/communications technology experts.

## 5. SPECTRUM MANAGEMENT

The radio spectrum usable for wireless communication spans frequencies from 3kHz to 300Ghz. To limit interference, governments have generally developed and used models to allocate an exclusive right to use certain frequencies to certain classes of uses and specific operators. Public policy has historically attempted to match propagation and performance properties of spectrum bands to specific capabilities and service requirements. Commercial AM radio operators received band assignments of the medium wavelength band from approximately 530Khz to 1700Khz, which yield low bandwidth but wide dispersion. Low frequencies have superior penetrative properties but limited bandwidth. Higher frequencies offer higher bandwidth but lower propagation distance.

Advances in mobile technology and service demands have strained this allocation model; many governments have shifted to spectrum auctions, intended to improve the efficiency of spectrum use and promote innovation of radio technology and infrastructure deployment. In the developed world, the spectrum is now fully allocated; any sector that wants to expand its use into other assigned bands can only do so if another sector releases its allocation.

In response to the perceived scarcity, Qualcomm originally developed a proposal for enabling LTE (4G commercial radio communications technology) to leak into unlicensed spectrum during peak periods, with large industry players engaged in controversy about the degree of harmful interference to unlicensed users from such spectrum sharing.[2]

[2]https://en.wikipedia.org/wiki/LTE_in_unlicensed_spectrum

Another stakeholder growing in significance and volume is the revitalized satellite industry. A new round of proposed MEOs, LEOs and HAPS are all trying to co-exist within the same spectrum bands. Smaller spacecraft, and increasing launch capabilities, are fueling interest in the use of thousands of satellites for broadband services. Finally, Software Defined Radio (SDR), which allows highly agile transmitters and receivers, has led to a rise in the intensity of use from individual users and applications.

To support mobile services, some countries are opening higher frequency bands, e.g., 3.5GHz-7GHz, previously considered too expensive to exploit at scale. The U.S. FCC was first to allocate the 28Ghz band, and are considering proposals for use in 95Ghz and higher bands. Given the limited propagation and penetrative qualities of these higher frequency bands, their feasibly for mobile data services is unproven. To this point, one participant presented an analysis of street view data to identify potential 5G base stations, considering height, location, trees, and vegetation to generate a projected coverage map. They used a simulation of part of a major city with a 3.5Ghz spectrum service to deliver a 1Gbps 5G service, and found that strategic deployment of 92 poles on could provide service to 93% of residences. Because 3.5Ghz spectrum (and lower) bands are so heavily allocated, they explored the potential viability of the deployment of the recently hyped option of millimeter wave services for a retail 5G service offering in the 28Ghz bands. The same deployment simulation in these higher frequency bands found a requirement for 1,400 poles, and yielded an acceptable level of service to only 82% of residences. That's a lot of poles, requiring a lot of power, and hundreds of billions of dollars of investment for each individual service provider with exclusive use spectrum allocations. The obvious question is whether the mobile data services market is capable of generating revenue to match this investment. Current carrier mergers with content providers (§1) suggests that access revenue itself could not cover this sort of investment. Furthermore, since U.S., household income in all but the top 10% has been relatively flat, there is limited household capacity to spend more money on Internet access services, 5G or otherwise.

## 6. CONCLUSIONS

We did not achieve, or seek, consensus on what the "public Internet" means, much less how to measure it. One view was that the public Internet is that portion of IP networking with strong public interest, that is, it uses public resources such as spectrum, rights-of-way, or tax subsidies, or deeply affects social factors such as economic performance and civic engagement. A more narrow technical model – the union of routable IP addresses and traffic between them – is commonly used in measurement research, but we need to connect this model to services that operate over the network, and then classify those services such that the classifications can help drive the policy goals specified by Congress. For example, at least in the U.S. today, the FCC regulates services, or the behavior of those offering them, and does so based on a set of social goals, e.g., emergency communication and public safety, universal service, consumer protection from cartel or monopoly behavior.

The rise of the large content providers brings up obvi-

ous regulatory concerns, but there are even less clearly delineated boundaries with respect to what harms one would regulate against, how to implement and target such policies, and even what disclosure/transparency requirements would best foster development of new policies. With respect to investment, the balance between incentivizing private sector investment in the infrastructure of a public communications platform, and protecting consumers from exploitation in the form of unduly high economic (or data mining) costs, is a never-ending challenge. A key question to answer is what options for self-provisioning (municipal and community wifi networks, emerging ad hoc networks) exist at the margin and if enough of those exist or may be ensured to exist, is it enough to protect key interests. In economic terms, this is a question about contestability. In network terms, it may be a question about multi-homing, community clouds, and other network architectural substitutes to relying on mega cloud providers.

# 7. FUTURE WORKSHOPS: DATA-DRIVEN POLICY

We identified areas where data could, and should, drive policy development, implementation, and evaluation: penetration (uptake), discriminatory behavior at choke points, pricing, and security and stability vulnerabilities and compromises. We plan to pursue these topics in a measurement/transparency (data disclosure) context at WIE 2018.

### *Penetration.*

This topic did not get much attention at this year's workshop. The only substantive conclusion we came to on mapping is that the $350 million dollars that NTIA spent on broadband mapping in 2009 mostly went to industry rather than independent third parties, and the resulting maps were questionable in quality and accuracy.

### *Discriminatory behavior at choke points.*

There are opportunities for potentially harmful forms of discrimination at choke points, whether they are network facilities, mobile device operating systems, search engines, app stores, etc. A example that recently motivated a regulatory response in the U.S. is interconnection links connecting access providers to their peers, transit providers and major content providers. The FCC's 2015 Open Internet Order asserted regulatory authority over those links, although the FCC acknowledged they lacked sufficient expertise to develop appropriate regulations thus far. With knowledge of interconnection link capacity, traffic link statistics are the most directly revealing form of information about whether a given link is experiencing congestion. But such aggregated traffic statistics are not currently shared, even with regulators (at least in the U.S.), for reasons that mostly amount to providers historically regarding that data as private. A variety of other measurements present complementary – and each incomplete on its own – views of interconnection performance. But without a basis of knowledge that relates measurement to justified inferences about actual impairment, different actors present opportunistic interpretations of data to support their points of view. Network operators themselves noted they would like more data on application-level performance and performance of components of an end-to-end path. For example, many issues with quality of the data streaming experience may be attributable to WiFi performance issues in the customer network, and it would be helpful to have data on this to assist customers in understanding why they are experiencing performance issues.

### *Pricing.*

Another obvious means of discrimination is via pricing, and pricing as well as terms of interconnection are also considered trade secrets. Such price occlusion admits the potential for discrimination and price gouging when market information is asymmetric.

### *Security and stability of critical infrastructure.*

We spent no time on this topic, but recognized it as an obvious area for future deeper discussion. We are likely to see mandated disclosures for security-relevant data sooner than for other types of data.

# 8. WORKSHOP PARTICIPANTS

- Co-Host: kc claffy, (CAIDA/UC San Diego)
- Co-Host: Dave Clark (MIT/CSAIL)
- Johannes Bauer (Michigan State University)
- Richard Clarke (AT&T)
- Amogh Dhamdhere (CAIDA/UC San Diego)
- Bradley Huffake (CAIDA/UC San Diego)
- Robert Frieden (Penn State)
- Emanuele Giovannetti (Anglia Ruskin University)
- Geoff Huston (APNIC)
- Arnaud Jacquet (BT)
- Mark Johnson (MCNC)
- Erin Kenneally (DHS)
- William Lehr (MIT)
- Jason Livingood (Comcast)
- Milo Medin (Google)
- James Miller (FCC / OET)
- Vishal Misra (Columbia University/Google) (remote)
- Scott Jordan (University of California Irvine)
- Walter Johnston (FCC)
- Andrew Odlyzko (University of Minnesota)
- Marvin Sirbu (Carnegie Mellon University)
- Tony Tauber (Comcast)
- Gautam Akiwate (CAIDA/UC San Diego)
- Esteban Carisimo (CAIDA/Universidad de Buenos Aires)
- Marina Fomenkov (CAIDA/UC San Diego)
- Bradley Huffaker (CAIDA/UC San Diego)
- Ricky Mok (CAIDA/UC San Diego)
- Josh Polterock (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 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.*