

Workshop on Internet Economics (WIE2013) 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 It has NOT been peer reviewed.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 2013, CAIDA and the Massachusetts Institute of Technology (MIT) hosted the (invitation-only) 4th interdisciplinary Workshop on Internet Economics (WIE) at the University of California's San Diego Supercomputer Center. This workshop series provides a forum for researchers, commercial Internet facilities and service providers, technologists, economists, theorists, policy makers, and other stakeholders to inform current and emerging regulatory and policy debates. The theme for this year's workshop was the economic health of the Internet ecosystem, including emphasis on the cost of and revenue sources to support content delivery, the quality of user experience, economic and policy influences on and effects of emerging specialized services, and the role of data in evaluating ecosystem health. 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 <http://www.caida.org/workshops/wie/1312/>.

Categories and Subject Descriptors

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

Keywords

Economics, Internet, Network management

1. INTRODUCTION

In December 2013, CAIDA and the Massachusetts Institute of Technology (MIT) hosted the 4th 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.

The topic for the previous WIE workshop (2012) was "Definitions and Data", and three themes emerged from those discussions: QoE measurement, private IP services, and baseline data collection to support research. This year we explored some of these themes in depth in pursuit of a better understanding of the economic

health of the Internet ecosystem. The Internet ecosystem – or more specifically, the consumer-centered part of the ecosystem – derives income from consumer spend on access and applications/content, and from advertising. We were interested in how these moneys flow into the various parts of the ecosystem, and how they sustain investments and innovations. Within that general framework, we held specific sessions on:

1. The cost of and incentives for efficient content delivery
2. The sustainability of the advertising revenue stream supporting the "free" Internet experience
3. The quality of the user experience (QoE, not QoS)
4. The role of specialized services as a driver of investment and/or threat to OTT (over-the-top) innovation
5. New models for regulation as a driver of a healthy ecosystem
6. The role of data: theory vs. empirical framings

The workshop format focused discussion around these topics: We spent about two hours per topic, with at least two 10-minute talks followed by an hour for each discussion. Three main conclusions emerged, each of which suggests new (or not yet overcome) research challenges. First, we refined last year's emphasis on the need for a standard (but evolving) definition of what constitute a baseline Internet service. Second, the economics of the Internet (advertising, content, access) ecosystem are remarkably opaque, but even generous estimates of the global advertising market suggest the need for new sources of revenues to cover the growing cost of many services that advertising now supports. A related issue is how to equitably cover the cost of infrastructure needed to support widespread sustained high-bandwidth flows (i.e., not Netflix). Third, specialized services will bring both benefits and challenges to the political economy of the Internet, not least of which is related to finding consensus on which services count as specialized and how the regulatory stance toward them should differ. In all cases there were lively discussions on how existing measurement infrastructure or data sources could shed light on the debates, or what new data would be needed to do so.

2. COST OF AND INCENTIVES FOR EFFICIENT CONTENT DELIVERY

The first day of the workshop focused on the interplay of cost and quality of content delivery, and its impact on the economic health of the Internet ecosystem. The emphasis on content delivery, as opposed to generic Internet data transport, was motivated by the observation that well over half of consumer download traffic is video and audio from content providers. Such high-bandwidth content has become a dominant component of traffic that access ISPs carry. An open question is which source(s) of capital will fund infrastructure improvements to accommodate this growth.

The discussion of the cost of content delivery did not provide any data on actual costs, but illustrated the complex engineering and business tradeoffs content providers make in choosing how to deliver content. Content providers normally employ caches to position content near the consumer, but these caches can be positioned at several places along the path to the consumer: within the content provider (e.g., Google's global network attaches to the Internet at many interconnection points), in transit networks, in access ISPs, or in the consumer's home network (e.g., a DVR such as Tivo). Different cache locations imply different costs and quality of delivery.

Technically, cache positioning is driven by at least two considerations: quality of delivery, and reduction of backbone utilization achieved as a result of repeated downloads from the cache ("cache hits"). The closer the cache is to the consumer, the lower the round trip, and thus the more responsive and less variable the data transfer. However, the closer the cache is to the consumer, the fewer consumers can be satisfied by the same cache and thus the less the backbone bandwidth savings. In the limit, if the average number of downloads for any content object approaches 1, the backbone bandwidth savings is zero.

For some content delivery today, the quality of delivery may be more important than the reduction of backbone traffic, which leads to pressure to position caches as close to the consumer as possible. These caches thus may end up in the access ISP's network, rather than the content provider's, which in turn implies a business negotiation around the terms under which the access ISP hosts the cache. A content provider might argue that there is benefit to the access ISP from having the cache inside its network, but if the reduction in backbone traffic is not material, the access ISP may view the cost of installing, powering and attending the cache as outweighing any reduction in network load. Negotiations about cache placement thus have somewhat the same flavor as peering negotiations.

The content provider has many options to control the quality of the download and playback, including cache positioning, selecting from among the available cache locations, format and encoding of the data, and size of playback buffer. The access ISP has fewer options. To first order, provisioning of adequate capacity will resolve most of the quality issues associated with the access ISP: with adequate capacity, there is no queueing at points of congestion, which minimizes jitter (variation in latency) and packet loss.

One view is that in a wireline access network, capacity is too cheap to provision to justify the use of techniques such as explicit QoS; with adequate capacity, there are no queues, and thus no need to use QoS to manage those queues. A competing view is that for any flow there will always be some point of minimum capacity, which may be at the access link itself. Congestion can arise there and cause jitter, so the use of QoS at this point (or capacity isolation mechanisms such as a per-flow queuing) may help to eliminate jitter for sensitive applications.

3. SUSTAINABILITY OF INTERNET SERVICES SUPPORTED BY ADVERTISING

Only a few sources of capital feed the Internet infrastructure: advertising, users, or government.¹ Andrew Odlyzko provided some data on the financial structure of the Internet industry. In support of his argument that "content is not king", he observed that Google's revenue was around \$50B in 2012, or about 2.5% of global revenues for the telecom industry that year. Global video revenues for Hollywood were around \$100B. Globally, telecom revenues

¹Another workshop could focus on another aspect of infrastructure capital needs, namely how to reduce the cost of network infrastructure dramatically, e.g., via more efficient use of spectrum.

(at around \$2T) were around 3% of GDP. This number may have peaked in 2001 at around 3.6%, and slowly declined to a perhaps stable 3%. World advertising (in all forms) was around \$500B, a quarter of the revenues of the telecom industry. Telecom industry expenditure on capital investments tended to be around 14% of revenues, far lower than in industries such as power or highways. Even though telecommunications is not a particularly capital-intensive industry, his numbers suggest that neither content nor advertising has the revenues to sustain the investment required for the infrastructure. The implication is that the primary benefit of broadband is not access to content, but simple connectivity.

David Clark elaborated on some of the advertising numbers provided by Andrew Odlyzko. Data from the Interactive Advertising Bureau (IAB) indicates that the spending per U.S. household on interactive (Internet) advertising is around \$35/month, 40% of which appears to be captured by Google. (Additional revenues received by Google from third-party web sites are passed on to the web site owners.) All of the "free" advertising-supported Internet services and applications are being sustained on this (estimated) \$35/month per broadband household. There is some evidence that as more services chase this advertising money, the value of an ad placement (and related activities such as behavioral tracking) is dropping, making it hard for some providers of services and content to survive on ad revenues alone.

The level of per-household advertising varies widely from country to country. Spending in the U.K. and in parts of Scandinavia are somewhat similar to the U.S., but in parts of Europe the numbers are far lower. Per household interactive spending in Italy was \$11.60 in 2012; in Spain it was \$8.90 (IAB estimates). A possible explanation for this is that ad spending in a region will be responsive to online commerce, and countries with weaker economies engage in less e-commerce. Plotting interactive ad spending vs. e-commerce for different countries shows high correlation. A possible conclusion from Clark's analysis is that e-commerce drives advertising, which in turn drives the "free" advertising-supported Internet experience, and thus the success of online commerce in any region is key to a healthy Internet ecosystem in that region.

Christopher Yoo provided a further perspective on the role of advertising in sustaining the creation of applications, services and content. If a consumer directly pays for access to some content, the price point will necessarily exclude consumers who find the content of low value. But the signal of value from the user is the willingness to pay. If the content is ad-supported, the value is reflected in the consumer response to the advertising, which has little to do with the value of the content itself. The only signal that the consumer generates about the value of the content is the decision to view. Content that is ad-supported will tend to be biased toward demographics most responsive to advertising, and will be subject to advertisers' preferences and concerns. Yoo reported that CBS derives 1/8th of the revenue per viewer as HBO. Programming with small audiences cannot survive under an ad-supported regime. Applying this model to the Internet, we can expect to see the emergence of both ad-supported and payment-supported applications, as well as tuning by providers of free vs. paid content.

Yoo raised the question of the appropriateness of payments from content providers to access (last-mile) providers. The answer appears to be different for payment-supported and ad-supported content. However, any price point charged to content providers for delivery of content will exclude from the market any application where the provider cannot appropriate the value of that application to a sufficient degree to cover costs. A user might attach a high value to a piece of ad-supported content, but the content provider may not be able to appropriate that value. Delivery schemes such

as “zero rating”, where the user is subject to usage tiers, but some content providers may choose to make a payment to exempt their content from the tier cap, allow both the provider to signal the value of delivery (by payment to avoid the cap), and the user to signal the value (by choosing to consume content subject to the cap).

4. QUALITY OF USER EXPERIENCE (QOE)

Another discussion centered on the topic of quality of experience (QoE). The term QoE is used to describe the subjective perception of quality by a user of an application, as opposed to the technical and objective measures of quality of service (QoS) such as bandwidth, delay and jitter, and packet loss. Since the ultimate goal of a network is to deliver a good QoE, QoS matters only to the extent that it has an influence over QoE. However, measurement of QoE involves engagement with users, and must be evaluated for different applications. There has been considerable research on specific applications, including streaming video, audio, use of the web, and gaming, that relates QoE for specific applications to the underlying QoS parameters of the delivery. If successful, this coupling can allow an operator (who cannot directly measure QoE) to estimate the quality of the user experience based on the underlying QoS parameters that can be more easily measured by the operators.

A regulator could also use such information to estimate what set of service parameters (e.g., access speed) are adequate for a typical basket of applications. Access service today is described in terms of technical parameters, most commonly peak speed. But how much speed is enough? One way to estimate required peak speed is to consider the speeds required to achieve a high QoE for a specified basket of applications, and typical degree of simultaneous usage from multiple devices in the home. This notion of a “basket of apps” gained consensus at last year’s WIE workshop and was reconfirmed this year, as one approach toward establishing a standard for 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.² The research challenge remains 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 select 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.

Complicating matters, many factors that can degrade QoE are not under the control of access providers. Many computers (especially those running older versions of operating systems) are shipped with configuration parameters that limit the peak speed of data transfers to a rate well below what access providers can deliver today. Wireless networks often introduce impairments. Customers frustrated that they cannot obtain the advertised speeds may be limited by their own home environment, not the service of the ISP.

This session ended with a discussion of different potential objectives for telecommunications regulation. In one view, the objective is to encourage competition. However, competition may not always protect the provider of higher-level services and content. An anecdote concerning France was discussed, in which the bargaining

²As discussed last year, this parameterized QoE score could also be used to frame willingness to pay for different applications, i.e., framing might be in terms of throughput for web browsing to loss or delay for interactive voip applications. Operators could further incorporate information such as a call graph (typical endpoints of communication for a given user and application) to enable more precise estimates of expected performance for a specific consumer.

over payment for interconnection seems to have blocked the entry of Netflix into the French market, even though there are competitive access ISPs in France.³ This discussion provided context for the final topic of the day: the relationship between the technical design of the Internet and potential regulatory goals.

5. SPECIALIZED SERVICES: DRIVER OF INFRASTRUCTURE INVESTMENT, AND THREAT TO INNOVATION

David Clark presented findings from a paper jointly written with kc claffy, which proposed a framework to guide regulators in understanding Internet technology. The framework distinguished two layers of infrastructure: an IP layer operated by a telecomms provider (a private platform that runs the Internet Protocol) and the global Internet that is formed by connecting these private platforms together. Not everything that runs over the Internet Protocols (such as Voice over IP or IP television) is running over the Internet, a distinction that is not always well-understood by non-technologists. If regulators determine that some regulatory intervention is justified (e.g., open access or reasonable network management obligation), the regulator should understand the consequences of imposing regulation on one or the other layer. In particular, regulation should not create a perverse incentive for operators to move away from a converged IP infrastructure. That is, a service should not be able to escape or acquire a regulatory burden by moving to IP.

The position proposed by Clark and claffy is that if the service provided to the global Internet by a provider is “good enough” for third party designers of applications and services, the regulator need not be concerned with what happens on the private IP layer. It remains an open question whether one can address the issue of “good enough” is by using metrics of QoE discussed above; as of now there are no standard metrics used by industry for this purpose.

A second view challenged the ability to practically define many terms, especially “specialized service”, even with Clark and claffy’s proposed framework, and that if the regulatory goal is to protect third-party developers, the regulator should consider imposing open access obligations on the private IP layer of an access provider.

A third view was that there is no arguable justification to protect third-party developers, and the goal of regulation should be to encourage competition, whether at the facilities level (competing private IP platforms) or retail competition to provide the global Internet service on top of the private IP platform of the access provider.

6. REGULATION AS A DRIVER OF A HEALTHY ECOSYSTEM

Christopher Yoo began the discussion of regulatory models by comparing the impacts of U.S. and European broadband policies in terms of investment and resulting broadband coverage (not subscriptions). Using EU and FCC data for comparison (although there was considerable controversy over the quality of these data), he found that the U.S. had higher coverage for most higher bandwidth options: FTTP, cable, and LTE. He also found that investment of capital into Internet infrastructure was almost twice as high in the United States, per capita or per household. However, download speeds and latency tended to be better in Europe.

William Lehr (MIT) talked about the regulatory challenges of the mobile industry, with its more complicated cost structure due to increased granularity of service and variety of rapidly evolving technologies, devices, and software in the ecosystem. There are more

³In March 2014, French ISP Orange’s CEO announced on-going negotiations with Netflix for service to Orange customers in France.

“common pool” resources (spectrum, towers, standards) associated with mobile services than with wired access. Higher shared costs complicate cost recovery, and make price discrimination a likely approach to recover costs, since incremental cost pricing for all goods will not suffice. Yet mobility increases consumer choice, and decreases the ability to price discriminate, since cross-price elasticity increases.

Sharon Gillett (Microsoft) gave a thought-provoking talk exploring the scenario of *innovation* as the primary goal of communications policy, explicitly recognizing the importance of a dynamic communications ecosystem to economic growth. She acknowledged the traditional regulatory view of innovation as a positive side-effect of achieving the regulatory goal of competition, i.e., we like competition because it fosters innovation. But she suggested that innovation may be the more primary goal – it is the foundation of economic growth and societal development, and it can also foster competition, if policies let it (e.g. mobile). Policies to mitigate the risks of network concentration need to foster both innovative networks innovative customers, devices, applications and services.

7. THE ROLE OF DATA: THEORY VS. EMPIRICAL FRAMINGS

kc claffy (CAIDA/UCSD) gave a brief overview of the types of active and passive Internet measurement data that CAIDA collects, curates, and shares. CAIDA’s Macroscopic Topology Project measures connectivity and latency using active probing to a stratified cross-section of the commodity IPv4 and IPv6 Internet. CAIDA also collaborates with organizations that operate network infrastructure to passively capture IP packet header traffic on selected links, anonymize IP addresses to allow trace sharing, and in some cases publish near real-time statistics of traffic captured from these links. CAIDA has developed and uses a privacy-sensitive data sharing framework that employs technical and policy means to balance individual privacy, security, and legal concerns against the needs of researchers and scientists for access to data. CAIDA maintains a list of publications by CAIDA researchers and collaborators, as well as publications by external researchers who are required to report back published use of CAIDA data (and publications we find via searching bi-annually). kc acknowledged that many policy questions require specific experimental design and cannot be addressed by existing CAIDA data, while at the same time observing that many existing data sources are underutilized for the purposes of understanding the Internet ecosystem.

Jonathan Liebenau and Silvia Elaluf-Calderwood (LSE) talked about which metrics are needed for Internet governance and who has relevant data. A lively discussion around exactly what data might be needed to answer what concrete questions regarding traffic flow, performance, and the economics of routing, peering, and traffic engineering. One point of view was that most technical data in CS networking papers is not actually relevant to policy, and in fact almost all data relevant to policy was proprietary. There was some consensus that performance problems (that negatively impact the consumer experience) due to congestion resulting from unresolved interconnection disputes was a question with definite public policy implications, and an open research challenge is how to measure such phenomenon from the edge.

8. FUTURE DIRECTIONS

A final thread of discussion focused on a wishlist of topics participants would like to see covered in future workshops, which included: more comparisons of regulatory policies and metrics of performance across countries; methods to diagnose path perfor-

mance and locate problems (e.g., at specific interconnection points); efficiency of overall services, and how to equitably determine who should fund infrastructure upgrades; how wireline technology can support mobility as fixed and mobile networks converge; studies of congestion of specific spectrum bands. We hope to pursue some of these and related questions at WIE2014.

This workshop series is motivated by our interest in how data can inform our understanding of the future, as well the deliberations of regulators. Both technical and economic data share the problem that many of the interesting facts are hidden from outside observers. For example, our discussion of efficient delivery of high-volume content illustrated that the barriers to technical efficiency may be unresolved business conflicts among parties, all mostly proprietary. Essential to progress will be models that guide us to understand what data may be the most important and revealing to pursue, and concerted cooperation between the public and private sector in pursuing and safely using it.

9. WORKSHOP PARTICIPANTS

- Co-Host: kc claffy, (CAIDA/UC, San Diego)
- Co-Host: Dave Clark (MIT)
- Silvia Elaluf-Calderwood (London School of Economics)
- Alissa Cooper (Center for Democracy & Technology)
- Rob Frieden (Penn State University)
- Sharon Gillett (Microsoft)
- Shane Greenstein (Northwestern University)
- Jennifer Holt (UC Santa Barbara)
- Scott Jordan (UC Irvine)
- William Lehr (MIT)
- Jonathan Liebenau (London School of Economics)
- Elisabeth Maida (Akamai Technologies)
- Denver Maddux (Microsoft)
- Kevin McElearney (Comcast)
- Milo Medin (Google)
- Gabor Molnar (UC, Boulder)
- Andrew Odlyzko (U. Minnesota)
- Achilles Petras (BT)
- Ivan Philips (Microsoft)
- Jesse Sowell (MIT)
- Tony Tauber (Comcast)
- Sanjay Udani (Verizon)
- Nicholas Weaver (ICSI)
- Madura Wijewardena (Comcast)
- Christopher Yoo (U Penn)
- Amogh Dhamdhare (CAIDA/UCSD)
- Bradley Huffaker (CAIDA/UCSD)
- Erin Kenneally (CAIDA/UCSD)
- Matthew Luckie (CAIDA/UCSD)

ACKNOWLEDGMENTS. 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.