

Approaches to transparency aimed at minimizing harm and maximizing investment

David D. Clark, MIT, ddc@csail.mit.edu
Steve Bauer, MIT, bauer@mit.edu
KC Claffy, UCSD, kc@caida.org

Executive Summary

Embedded in a challenging legal and historical context, the FCC must act in the short term to address concerns about harmful discriminatory behavior. But its actions should be consistent with an effective, long-term approach that might ultimately reflect a change in legal framing and authority. In this comment we do not express a preference among short-term options, e.g., section 706 vs. Title II. Instead we suggest steps that would support any short-term option chosen by the FCC, but also inform debate about longer term policy options. Our suggestions are informed by recent research on Internet connectivity structure and performance, from technical as well as business perspectives, and our motivation is enabling fact-based policy. Our line of reasoning is as follows.

- 1) Recent discourse about Internet regulation has focused on whether or how to regulate discrimination rather than on its possible harms and benefits. For four reasons, **we advocate explicit attention to possible harms, their causes, and means to prevent them.** First, the court has stated that while the FCC cannot ban traffic discrimination unless it reclassifies Internet access providers under Title II, the FCC does have the authority to remedy harms. Second, a focus on harms provides a possible way to govern specialized services, which are currently not subject to traffic management constraints. Third, if the FCC chooses Title II, it will open up many questions about which parts to enforce, which will require a discussion of the harms vs. benefits of selective forbearance. Fourth, any new regulatory framework would be well-served by a thorough understanding of potential harms and benefits that result from behavior of various actors.
- 2) Impaired quality of experience (QoE) is a meaningful indicator of harm, essential to effective regulatory analysis focused on harms, but not well-understood. Unfortunately, tools to measure and analyze QoE today are primitive, which is a key obstacle to using it as part of a sound basis for regulation. **The FCC should promote research, tools and capabilities to measure, quantify, and characterize QoE,** and explore metrics of service quality that better reflect our understanding of QoS and QoE for a range of applications.
- 3) **Interconnection is not separable from discrimination when evaluating harmful behavior,** since interconnection practices can induce harms that do not materially differ from those induced by discrimination. Different modes of connection raise different potential harms, and an overall analysis of such harms can suggest reasonable constraints on ISP and edge provider behavior.
- 4) One such constraint is that **to the extent that terminating monopoly ISPs are under no obligation to peer, they should have to provide uncongested interconnection (peering and transit) links.** However, this approach requires cognizance of the distinction between small edge providers (“congestion-takers”) and large edge providers (“congestion-makers”). That is, **large edge providers who can send enough traffic volume to induce congestion, must responsibly manage their traffic sources** and negotiate in good faith for direct interconnection or other solutions where appropriate, so that ISPs can reasonably fulfill their obligation.
- 5) **A longer-term challenge is to create a regulatory framework that promotes improvements in infrastructure performance, robustness, and security while also promoting freedom and innovation in a highly dynamic ecosystem.** An examination of the Internet ecosystem from an *industrial platform* perspective can help scope consistent policy discourse, in particular around *specialized services*, and anchor a framework that balances these aspirations.
- 6) **A regulatory requirement for transparent consideration of harms and benefits will support any direction the FCC chooses** This requirement brings four challenges: ensuring such analyses are sufficiently detailed to allow independent, third-party evaluation; obtaining independent evaluations from objective parties; adapting to changing reasonable expectations about QoE over time; and capturing specific as well as more general societal harms, such as the effects of under-investment in capacity.

I. WE SHOULD EXPLORE HARMS AND BENEFITS, THEIR CAUSES, AND HOW TO PREVENT HARMS

Network discrimination is a means by which an operator can cause harm or benefit to various classes of actors. Discourse about regulation today has centered on this means, and the potential to regulate it, rather than on the harms and benefits themselves. Looking at different classes of entities, and the specific harms that might fall to each of them, provides a basis for reasoning about possible interventions, now or in the future.

We taxonomize harms by asking on whom they are imposed, e.g., edge providers vs. end users, and what types of providers. Not all harms arise from discrimination: for example harms include extraction of monopoly rents and other unjustified pricing to customers, peers, or edge providers. The FCC should consider whether the focus on regulation of traffic discrimination addresses the most important of the realistic harms.

ISPs may be able to hinder *edge providers* (i.e., of application and content), one consequence of which may be a slowing of innovation and associated economic benefits. Harms to *end users* are less well-defined (and the NPRM left them unspecified) since it is impossible, for example, to measure harm to an end user from a hindrance to edge-provider innovation that reduces the number of market offerings. Other possible harms to end users include unreasonable pricing, data caps, or performance. In the past, the FCC has used protection of the end user from harm as a justification to limit discrimination on edge providers: *“To encourage broadband deployment and preserve and promote the open and interconnected nature of the public Internet, consumers are entitled to access the lawful Internet content of their choice.”* (The first of the four original (2005) neutrality principles [4]). It has become clear since that time that harm to the end user cannot serve as a proxy for all harms that may affect edge providers.

One way to assign the “harm” label to a specific discriminatory treatment of packets is to find evidence that it impairs the user’s Quality of Experience (see section II). However, many traffic discrimination scenarios may benefit users, the most obvious being protecting latency-sensitive traffic from the consequences of co-mingling with other traffic. The need for such discrimination arises especially across broadband access links, where a single end user (or single home full of devices behind one access modem) may operate several applications at once – a multi-player game, a voice call, a streaming video and a bulk data transfer – and the traffic flows may interfere with each other, impairing QoE. With enough knowledge one may be able to configure a home router to mitigate such interference, but users may want to simply purchase a broadband service that has reduced jitter to support voice and gaming. In other words, users have varying demands for QoE, and varying interest in paying for it; today’s public Internet cannot support such user choice. Thus, *we advocate a careful, structured exploration of the space of possible harms and benefits, seeking a framework that minimizes the harms and maximizes the benefits.* This approach offers four advantages in the near and long term. First, the court has stated that while the FCC cannot ban traffic discrimination unless it reclassifies Internet access providers under Title II, the FCC does have the authority to remedy harms. Second, the FCC allows Internet access providers to offer their own *specialized services* (IP-based services offered over the same physical infrastructure as the public Internet), which constrains the FCC’s ability to minimize harms or protect the Internet through regulation of Internet traffic discrimination. A focus on harms provides an alternative way to govern specialized services. Third, if the FCC chooses Title II, it will open up many questions about which parts to enforce vs. forbear, which will undoubtedly require a discussion of the relative harms vs. benefits of selective forbearance. Fourth, any significant new regulatory framework for Internet infrastructure providers would be well-served by a thorough understanding of potential harms and benefits that result from behaviors of various actors.

II. IMPAIRED QUALITY OF EXPERIENCE IS A MEANINGFUL INDICATOR OF HARM

Quality of Experience (QoE) is a measure of whether a use of the Internet meets the needs and (“reasonable”) expectations of the user and provider in the context of any specific application. As such, it is one basis to judge if harms have been imposed on the end-user. QoE is a subjective measure of quality. Was voice quality degraded? Is the video stream experiencing re-buffering delays? While *quality of service (QoS)* is used to describe the technical parameters of any service (peak speed, latency, jitter, packet loss), assessment of QoE is application-specific. Impairments to QoE derive in part from underlying QoS parameters (e.g., jitter impairs voice QoE but not email), and can arise anywhere in the network.

One long-term measure of success for the Internet would be a wide variety of innovative edge-provider services available with good QoE. Achieving this goal would require all parties to manage (and invest) to eliminate barriers

to high QoE, and refrain from discriminatory treatment that impairs QoE differently among similar applications.¹ For example, persistent congestion on paths degrades QoE on certain applications. We are developing and have demonstrated a research method that can detect such congestion at points of interconnection [8], [12]. In a recent collaboration our research groups (at MIT and UCSD/CAIDA) have found that business contention (usually well-reported in the press) correlates with evidence of substantial persistent congestion at points of interconnection – some episodes that have persisted for days or even months – but we have not found evidence that such persistent congestion is widespread. As another example, jitter is normally a consequence of queues forming and dissipating as instantaneous episodes of congestion occur.² Our current research method focuses on persistent rather than instantaneous congestion, but with careful instrumentation and refinement of such methods, the research community could also develop tools that they, the consumer or the FCC could use to gather evidence on whether jitter is arising at points of interconnection or elsewhere in the network.

We acknowledge that the community’s general understanding of QoE is relatively primitive, as is the way we describe broadband service quality. Today, retail broadband access services are classified by peak speed. Users try to purchase enough capacity to run their preferred applications with their preferred QoE, but peak speed does not determine whether ISPs are always delivering that QoE, since transient or persistent congestion or other impairment on certain links may reduce (sometimes to zero) a user’s ability to run a desired application.

Since our understanding of congestion and its impact on QoE is still rudimentary, this approach will require advances in measurement, although not dramatic ones. To illustrate, standard speed tests to a server across an interconnection link (to an “off-net” server) may detect congestion at that point, if a user knows how to run the appropriate test and there is congestion at most at one point in the path. But for networks such as Google Fiber that offer high speed local access, the user might discover that an off-net speed test could not run at the full access speed, while at the same time the typical apps could achieve adequate QoE off-net. Direct measurement of QoE (e.g., instrumentation of applications), combined with active measurement focused on QoE rather than raw speed, will better reveal problems than a simple speed test to distant parts of the Internet. *The FCC should promote research, tools and capabilities to measure, quantify, and characterize QoE, and explore metrics of service quality that better reflect our understanding of QoS and QoE for a range of applications.*

III. INTERCONNECTION IS NOT SEPARABLE FROM DISCRIMINATION WHEN EVALUATING HARMFUL BEHAVIOR

Traditionally, interconnection has been an issue among ISPs, who purchase transit and arrange peering to route traffic among themselves. But since interconnection practices can induce harms that do not materially differ from those induced by discrimination, *we believe the FCC should expand its scope of attention to include interconnection.* Edge providers use different modes of access to reach end users, and to understand potential harms to edge providers, we distinguish among three different modes of access: direct, indirect and third-party. The distinction is important because the different modes may raise different regulatory issues.

Direct interconnection is used by large edge providers such as Netflix, Youtube or the major CDN providers to gain direct access, and thus faster delivery of content, to the customers of a broadband access provider.³

Indirect access is how smaller edge providers reach customers of access providers, by purchasing transit service from one or a few providers, and then relying on those providers to access the rest of the Internet.

Third-party connection uses a service such as a content delivery network (CDN) to reach end consumers; the CDN has likely negotiated direct interconnections with major access ISPs in order to provide customers the benefits of close proximity without the customers having to engineer these paths for themselves.

Because large and small edge providers face different options in reaching customers, they are subject to different potential harms. Small providers do not typically purchase direct interconnection, so they are not directly subject to harms that might arise as a consequence of discriminatory treatment of those interconnection agreements. On the other hand, small providers can be collaterally harmed by issues that arise from disputes involving large providers, if they send traffic over the same links – and note this harm does not derive from traffic discrimination. Smaller

¹Such a long-term principle must periodically be translated into rules applicable to current technology and requirements. Adjusting rules as needed to re-align behavior with the same long-term goal is a formula for *adaptive regulation*, a recent movement to increase the pace of regulation, which we critically analyzed in [7].

²[2] provides background on different definitions of Internet congestion that are used within the community.

³[3] describes constraints on bilateral negotiation between networks contemplating interconnection. In most cases (but less so for large providers such as Netflix), commodity transit prices serve as a cap on the price of paid peering interconnection.

content providers may be more resistant to harm from explicit, directed traffic discrimination, since their flows may be harder to isolate, especially if they are serving content via a CDN that aggregates across many content providers. Nonetheless, such discriminatory treatment represents an obvious harm and should be prohibited.

Smaller edge providers today benefit from a rich set of intermediate cloud, storage, and CDN providers prepared to serve small providers. These intermediate service providers implement *third-party access*, i.e., small firms use a cloud provider for storage and computation, and a CDN or other overlay service for delivery and customer access. These service providers in turn negotiate interconnection agreements with large providers. The health of these intermediate providers, and whether they struggle to negotiate for direct interconnection, is a signal of the presence or absence of barriers to innovation for small providers. In other words, while small edge providers do not directly negotiate for direct interconnection, they are dependent on the ability of these intermediate providers to negotiate suitable interconnection. *Monitoring the competitiveness of these third-party industries (cloud, storage, CDN) could yield potential warning signals of future harm.* That ecosystem appears robustly competitive today (we have not done a careful analysis), but poor health of that support ecosystem will hinder smaller edge providers trying to bring new applications to market. The challenge for regulators is that pricing and other terms of these negotiations are entirely opaque; the agreements are usually covered by NDAs. Thus there is no way today, except waiting for a complaint of material deterioration in the quality of this intermediate platform, to tell if discriminatory (or other) treatment by access ISPs is harming this part of the ecosystem. *Approaches to increasing transparency of the interconnection system include mandatory disclosures of operating practices that might induce harms, and measurement and data sharing capabilities deployed by operators and others.*

IV. IF TERMINATING MONOPOLY ISPs ARE UNDER NO OBLIGATION TO PEER OR AGREE TO DIRECT INTERCONNECTION, THEY MUST PROVIDE UNCONGESTED PATHS TO THEIR NETWORKS

Access ISPs presumptively have market power as a terminating monopoly, and can impose both technical and economic harms as part of a business negotiation, or favor their own higher-level services. In particular ISPs today have no obligation to negotiate a direct interconnection agreement with any party. One approach to limiting the potential exercise of this power is to *require that ISPs provide adequate means for edge providers and off-net users to reach their customers over interconnection and transit links.* This standard provides some protection to the end user from impaired QoE, and also gives a basic level of protection to the edge provider. Failure to agree to terms for direct connection with an edge provider is not an unreasonable harm to that edge provider, so long as there is adequate overall interconnection capacity (including transit capacity) so that the edge provider can get its traffic to customers of the ISPs. Failure to agree to direct interconnection terms means that the edge provider will pay the cost of transit to reach the customers of the access ISP, but provided that transit continues to be priced as it is today in the U.S., and the transit path into the access ISP is uncongested, this outcome is not unreasonable. (Edge providers who purchase transit service from an ISP may under-provision these links at their own risk.)

However, this approach requires cognizance of the distinction between small edge providers (“congestion-takers”) and large edge providers (“congestion-makers”). Small edge providers, who generally use indirect access, have few options to avoid congestion on links carrying their flows; large edge providers may send enough traffic volume that links must be engineered or upgraded to carry it. So if ISPs are expected to provide uncongested paths into their network, *such large volume edge providers must responsibly manage their traffic sources, and negotiate in good faith for direct interconnection with the ISP or other mutually agreeable solution where appropriate, so that the ISP can reasonably fulfill its obligation.*

In asking whether an ISP is fulfilling its minimum commitment to its users and to the ecosystem of edge providers using indirect connections, the FCC should focus on the transit connections of an ISP, which are the “routes of last resort” for that ISP. Persistently congested links are an indicator that users with traffic flowing over those links will not achieve their desired service quality, and should attract regulatory attention, but is another example where there is no transparency today. Negotiations and agreements are covered by NDAs, there are no requirements for disclosure of any operational data, and thus no way for outside observers and regulators to judge whether parties are fulfilling their obligations and expectations. *We recommend the FCC adopt a policy that triggers when persistent congestion is detected on major paths in the Internet, to compel access to information from ISPs relevant to the cause of the congestion. If the FCC establishes a framework where impaired QoE is considered a harm, other causes of persistent impairment would also be a basis for the FCC to require disclosure of relevant information.*

V. A LONGER-TERM POLICY CHALLENGE IS TO CREATE A REGULATORY FRAMEWORK THAT PROMOTES IMPROVEMENTS IN INFRASTRUCTURE PERFORMANCE, ROBUSTNESS, AND SECURITY WHILE ALSO PROMOTING FREEDOM AND INNOVATION IN A HIGHLY DYNAMIC ECOSYSTEM

The FCC has used the term *specialized services* in several orders; the European Parliament borrowed the term in their recent ruling on network neutrality [14]. The definition is murky in these documents; describing such services in terms of *industrial platforms* helps clarify the regulatory intentions [11]. Broadband providers are now engineering their network infrastructures as *single-firm IP-based platforms*, on which they provide a range of services, only one of which is the global Internet. Other services include IPTV and other special services with QoS support that is not offered on the global Internet. Building on this, several large broadband providers have interconnected their internal IP platforms with QoS agreements to create a global network for high-quality voice (and perhaps other) services. Such alternative platforms allow service differentiation not available on the public Internet, i.e., improved QoE, with associated revenue opportunities to charge for it.⁴

The emergence of these alternative platforms implies that there are material impairments to QoE in the global Internet, since otherwise investment in these enhancements would be hard to justify. This implication triggers four questions, the most immediate of which is: what is driving these QoE impairments today, technical issues or business/engineering decisions? For example, there is no technical reason for congestion at interconnection points – the interconnecting parties can remedy the problem using current technology, e.g., additional capacity or traffic engineering. In contrast, the current generation of last-mile (broadband access) technology can trigger QoE issues that are harder to eliminate. Even if the consumer has purchased enough capacity to meet his nominal needs, a phenomenon called *buffer bloat* [10] can induce jitter on access links. Such impairments cannot be eliminated without upgrading to new technology that includes traffic management schemes, e.g., discriminatory treatment of traffic. Specialized services today achieve improved QoE using such features to isolate specialized services from Internet traffic. With regulatory approval, ISPs could use these techniques to implement beneficial traffic discrimination for Internet applications, such as protecting latency sensitive traffic from impairments such as buffer bloat.

A second question is what these QoE impairments on the public Internet imply for edge-provider innovation: do they inhibit development of classes of applications? Real harms might arise from an access provider's ability to negotiate direct interconnection on its private IP platform with arbitrary levels of individual specialization and discrimination. One such harmful trajectory is that innovative edge providers who need special QoE, and can afford it, will migrate away from the global Internet onto alternative specialized service platforms, eroding the centrality of the global Internet as the universal platform for innovation. The future might be an "all IP" world, but with parallel IP platforms competing both for end user and edge-provider access. Inherent harms of such a fragmented world include inconsistent access to end users, and constant concerns about discriminatory treatment.

The third question that arises from the emergence of specialized services is whether they fill a substantial market demand, or are primarily a way for infrastructure owners to maximize return on their capital investment? Payment for QoE may have benefits that exceed harms to consumers, and it is not obvious that the potential loss of benefits implied by neutrality are small compared to potential harms (although the NPRM presumes it [6, para. 27]).

Finally, designing a sustainable regulation requires acknowledging the degrees of freedom of all parties to evade it. Banning discriminatory traffic treatment, while it may prevent certain harms, may also prevent the public Internet from competing with alternative private IP platforms with superior QoE, an undesirable outcome. A structured focus on potential harms and benefits of specific discrimination and pricing behavior of ISPs can help frame a debate about how to maintain the Internet as a vigorous platform that can compete with alternative private platforms. *Specifically, allowing ISPs to sell QoS enhancements that lead to improved QoE on the public Internet may reduce the drive to offer specialized services, preserving the Internet as the unified service platform.*

QoS enhancement on the public Internet may conflict with ISP preferences to bring proprietary schemes to the market. But the benefit of disclosure and standardization outweighs the unproven benefit of proprietary solutions. Proprietary solutions would allow ISPs to compete with each other by offering differentiated enhancements, but at the expense of the edge providers, who would have to adapt their services to different proprietary schemes. Further,

⁴In recent work [11] used multi-sided platform theory to model key technical and business aspects of today's industry. Among other examples, we described the cellular industry's (GSMA) use of the Internetwork Packet Exchange (IPX) framework to create such private interconnected IP platforms, which allow for isolation of voice from other traffic to optimize end-to-end QoE as well as support specific payment terms. <http://www.gsma.com/technicalprojects/technical-programme/ip-exchange>.

there is no evidence that such differentiation would influence consumer selection of provider. QoE enhancement is better seen as a way for the industry as a whole to bring new services to the Internet, provide a potential basis for additional revenue, and protect the Internet from competing private platforms that would otherwise threaten the Internet's long-term viability.

VI. A REGULATORY REQUIREMENT FOR TRANSPARENT CONSIDERATION OF HARMS AND BENEFITS WILL SUPPORT ANY DIRECTION THE FCC CHOOSES

Today, ISPs are required to disclose network management practices, which must (for wireline networks) be “reasonable”, a vague criterion that is not subject to any objective review. An alternative is to require that any proposals for network management and discriminatory treatment of traffic, including outright blocking (e.g., for spam filtering [1]) disclose a harms-benefit analysis.

We emphasize four essential challenges to this approach. First, this published analysis should be sufficiently detailed and measurable so as to allow an independent third-party to independently estimate the potential for benefit and harm. For example, if an ISP gives priority to game and VoIP traffic (which would benefit jitter-sensitive flows) and limits this class to a small fraction of the overall traffic (say 5%), its negative consequences is spread across the other 95%. The consequential harms can probably be quantified, and will probably be minimal.

The second challenge in relying on a harm-benefit analysis is that some harms may not be detectable either by edge providers or sophisticated end users (e.g., researchers). But measurement as a tool of enforcement will be a challenge regardless of which path the FCC takes, including how to detect prohibited or “unreasonable” behaviors. This challenge reinforces the importance of our recommendation that the FCC work with other agencies to promote research, tools and capabilities to measure, quantify, and characterize QoE (Section II).

Third, an approach that describes harms to the end user in terms of impaired QoE must account for what constitutes reasonable expectations about QoE at any given time. In the 1980s, when the research community first demonstrated IPTV, it was not reasonably expected to work over the deployed Internet. Similarly, high-speed HD video is probably not a reasonable expectation over most of today's cellular infrastructure; users should and do have different expectations about what is a reasonable QoE in wireline and wireless networks. But, as an alternative to the prevailing notion of “reasonable network management”, the concept of “reasonable expectation of QoE” can also resolve the awkward current distinction between wireline and wireless discrimination rules (wireless providers today have no obligations to limit themselves to “reasonable network management”). Specifically, the FCC could subject both wireline and wireless providers to requirements with respect to discrimination, but allow a wider latitude for reasonable QoE based on capacity constraints and unpredictable demands of wireless infrastructure.

Fourth, the harm-benefit analysis will have to balance the potential for harms to specific edge providers due to inappropriate discrimination with the potential of more general harms such as those due to under-investment in capacity. Application developers and edge providers depend on the capability of the infrastructure to serve their needs; performance limits due to under-investment in infrastructure can be a real barrier. Thus far, broadband access speed has been the only quality metric – the FCC set a target of 100M homes at 100 mb/s by 2020 [9]. Another increasingly important metric is the overall cap on data download per billing period. Today, some providers in the U.S. are setting usage caps of 300 GB/month or less. An HD movie today uses about 3 GB/hour [13], which means that a Netflix-subscribing household can watch no more than 100 hours a month or about 3 hours a day of HD video. For a family with multiple displays and heterogeneous watching habits, 3 hours a day is already too constraining. So inadequate usage caps as well as inadequate access speeds represent real harms both to end user and some edge providers. These harms do not arise because of discriminatory treatment but just an overall level of under-investment and/or objective of maximizing profit margin on current capital investment.

In taking note of this sort of harm, which does not arise from traffic discrimination or network management, we are not arguing that using a harm-benefit analysis is a justification to broaden the scope of the FCC's authority—we recognize that the NPRM primarily concerns itself with issues of discrimination. However, the FCC has been willing at least to opine on issues that seem to be primarily about pricing: the Open Internet Report and Order [5, sec. 72] speculates that differently priced usage tiers would probably not attract regulatory attention. We urge the FCC to identify the full range of harms that might befall both end users and edge providers, of which we believe usage caps will increasingly be important, and use all means at their disposal, both formal and informal, to mitigate these harms. Do not be overly distracted by a focus on “neutrality”.

REFERENCES

- [1] AT&T. ATT spam filtering, March 2012. <http://www.att.com/gen/public-affairs?pid=20879> “[C]onsistent with recommendations from the Federal Trade Commission (FTC) for guarding against SPAM, AT&T prevents the use of Port 25 for sending email on our wired consumer broadband Internet access services.” AT&T cites the authority of the FTC to argue that the benefit of allowing the user to have access to port 25 is outweighed by the potential for harm.
- [2] Steven Bauer, David D. Clark, and William Lehr. The Evolution of Internet Congestion. In *Telecommunications Policy Research Conference (TPRC)*, 2009. Provides background on different definitions of Internet congestion that are used within the community.
- [3] David D. Clark, William Lehr, and Steven Bauer. Interconnection in the Internet: The Policy Challenge. In *Telecommunications Policy Research Conference (TPRC)*, 2011. Describes constraints on bilateral negotiation between networks contemplating interconnection. In most cases (but not large providers such as Netflix), commodity transit prices serve as a cap on the price of paid peering interconnection.
- [4] Federal Communications Commission. Fcc 05-151, policy statement, 2005. In the matter of Appropriate Framework for Broadband Access to the Internet over Wireline Facilities, CC Docket No. 02-33, etc., http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-05-151A1.doc.
- [5] Federal Communications Commission. Report and order, preserving the open internet, fcc10-201, 2010.
- [6] United States Federal Communications Commission. Notice of Proposed Rulemaking, May 2014. <http://www.fcc.gov/document/protecting-and-promoting-open-internet-nprm>.
- [7] David D. Clark and KC Claffy. Anchoring policy development around stable points: an approach to regulating the co-evolving ICT ecosystem. Technical report, 2014. (submitted for publication to *Telecommunications Policy* but also at <http://www.caida.org/publications/papers/2014/>).
- [8] David D. Clark, Steve Bauer, Amogh Dhamdhere, Bradley Huffaker, William Lehr, Matthew Luckie and kc claffy. Measurement and Analysis of Internet Interconnection and Congestion. In *Telecommunications Policy Research Conference (TPRC)*, 2014. Reports preliminary results from an experimental method that probes links in the interior of the Internet for evidence of persistent congestion, specifically applied to interconnection links. We found several (but not widespread) examples of links where congestion persisted for many hours each day, and correlated them with well-publicized unresolved business issues, e.g., between Netflix and Comcast.
- [9] Federal Communications Commission. National Broadband Plan, 2010.
- [10] Jean-Philippe Lang. Bufferbloat, 2014. <http://www.bufferbloat.net>.
- [11] kc claffy and David D. Clark. Platform Models for Sustainable Internet Regulation. *Journal of Information Policy*, 4, 2014. (to appear, but available at <http://www.caida.org/outreach/papers/2014/>). Uses multi-sided platform theory in conjunction with layered characterization of IP networked ecosystem to model key technical and business aspects of today's industry. We used the model to illustrate the implications of different ways that ISPs use their single-firm IP platforms today: selling private “specialized” services (carrier IPTV), multi-firm private services (carrier-grade VoIP), or the global Internet, a service made by interconnecting single-firm IP platforms.
- [12] Matthew Luckie and Amogh Dhamdhere and David D. Clark and Bradley Huffaker and kc Claffy. Challenges in Inferring Internet Interdomain Congestion. In *Proceedings of ACM SIGCOMM Internet Measurement Conference (IMC)*, 2014.
- [13] Netflix. Internet Connection Speed Recommendations, 2014.
- [14] European Parliament. European Parliament legislative resolution on the proposal for a regulation of the European Parliament and of the Council laying down measures concerning the European single market for electronic communications, April 2014. <http://www.europarl.europa.eu/sides/getDoc.do?type=TA&reference=P7-TA-2014-0281&language=EN>.