

Workshop on Internet Economics (WIE2014) Report

kc claffy
UCSD/CAIDA
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 It has NOT been peer reviewed.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 content. Comments can be posted through CCR Online.

ABSTRACT

On December 10-11 2014, we hosted the 4th interdisciplinary Workshop on Internet Economics (WIE) at the UC San Diego's Supercomputer Center. This workshop series provides a forum for researchers, Internet facilities and service providers, technologists, economists, theorists, policy makers, and other stakeholders to inform current and emerging regulatory and policy debates. The objective for this year's workshop was a structured consideration of whether and how policy-makers should try to shape the future of the Internet. To structure the discussion about policy, we began the workshop with a list of potential *aspirations* for our future telecommunications infrastructure (a list we had previously collated), and asked participants to articulate an aspiration or fear they had about the future of the Internet, which we summarized and discussed on the second day. The focus on aspirations was motivated by the high-level observation that before discussing regulation, we must agree on the objective of the regulation, and why the intended outcome is justified. In parallel, we used a similar format as in previous years: a series of focused sessions, where 3-4 presenters each prepared 10-minute talks on issues in recent regulatory discourse, followed by in-depth discussions. This report highlights the discussions and presents relevant open research questions identified by participants. Slides presented and a copy of this report are available at

<http://www.caida.org/workshops/wie/1412/>.

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

1. INTRODUCTION

In December 2014, UC San Diego's Center for Advanced Internet Data Analysis and MIT's Computer Science and AI Laboratory co-hosted the 5th interdisciplinary Workshop on Internet Economics (WIE) at the University of California, San Diego. In hosting this workshop series we recognize that the future of the Internet is shaped as much by economic factors as by technical innovations, and our goal is to provide a forum for researchers, commercial Internet facilities and service providers, technologists, economists, theorists, policy makers, and other stakeholders to empirically inform emerging regulatory and policy debates.

The topic for the previous WIE workshop (2013) was "Economic health of the Internet ecosystem". This year (2014) we organized

the workshop slightly differently, motivated by recent events and policy debates in Washington, D.C. Specifically, as the FCC once again tried to impose some regulation on the Internet, while struggling to fit within constraints of the current law, we see growing recognition that ultimately (perhaps within a decade), the current regulatory framework of the law will be deemed unworkable, warranting a re-write. In this context, we are wary of how discussions today tend to center on means, e.g., of crafting a network neutrality regulation, when neutrality is not itself a goal, but rather a means to prevent a set of harms to various players. A clear understanding of what outcomes we are trying to prevent or achieve is a more auspicious starting point for long-range discussion. To this end we attempted to frame a discussion that considered specific goals – aspirations for our future telecommunications infrastructure – as well as what approaches will prove effective in trying to achieve that future. This report tries to capture a flavor of the conversations, while avoiding attribution of previously unpublished viewpoints.

2. HOPES AND FEARS FOR THE FUTURE

Some time ago (2012) the two of us cataloged all the aspirations and fears that we could find articulated in public debates about the future of the Internet. Our objective was to subject the list to critical analysis, and motivate a debate over their desirability and feasibility, and effective means to achieve them. The list is not original to us, nor do we agree with all of them. We tried to comprehensively gather what others have said, for coherent consideration. We began the workshop by presenting this list, acknowledging three high-level conclusions we reached in this collation process, perhaps obvious but often neglected. First, not only are many of the aspirations hard to achieve, but some are incompatible with others. Second, many are under-specified and resist operational definition, leaving it unclear how to translate the aspiration to concrete goals against which to measure progress. Third, most tools society has to shape the future of the Internet seem unequal to the task. In some cases, terms are not yet defined well enough to frame a technology and policy research agenda that could help mitigate, circumvent, or navigate barriers to some of these aspirations.

We used this list to loosely frame workshop discussions. We solicited initial feedback on the list, and asked participants to add their own set of aspirations and fears via an online survey tool.¹ We integrated the survey responses into the list after the first day of the workshop and spent most of the second day discussing competing views and perspectives on the resulting list of aspirations, and whether they can or should be mapped to specific governance instruments. All participants recognized the tension among aspirations, the need to accommodate tradeoffs and navigate diverging

¹Each aspiration can be stated in positive or negative terms: what we want, or the failure to achieve it.

priorities. Our initial list was the following:

1. **Reach:** The Internet should reach every person.
2. **Ubiquity:** The Internet should be available to us everywhere.
3. **Evolution:** The Internet should continue to evolve to match the pace and direction of the larger IT sector.
4. **Uptake:** The Internet should be used by more people.
5. **Affordable:** Cost should not be a barrier to the use of the Internet.
6. **Trustworthy:** The Internet should provide experiences that are sufficiently free of frustration, fears and unpleasant experiences that people are not deterred from using it.
7. **Lawful:** The Internet should not be an effective space for law-breakers.
8. **National security:** The Internet should not raise concerns about national security
9. **Innovation:** The Internet should be a platform for vigorous innovation, and thus a driver of the economy.
10. **Generality:** The Internet should support a wide range of services and applications.
11. **Unblocked:** Internet content should be accessible to all without blocking or censorship.
12. **Choice:** The consumer should have choices in their Internet experience.
13. **Redistribution:** The Internet should serve as a mechanism for the distribution of wealth among different sectors and countries.
14. **Unification:** The Internet (and Internet technology, whether in the public net or not) should become a unified technology platform for communication.
15. **Local values:** For any region of the globe, the behavior of the Internet should be consistent with and reflect its core cultural/political values.
16. **Universal values:** The Internet should be a tool to promote social, cultural, and political values, especially “universal” ones.
17. **Global:** The Internet should be a means of communication between citizens of the world.

We pre-emptively noted that, like network neutrality, some aspirations on the list were not end goals but means to other goals, e.g., *Incentive, Choice*. The aspirations also have a spatial dimension: some are addressed in a national context, e.g., *Reach, Uptake*, others are inherently trans-national, e.g., *Trustworthy*. We invited attendees to consider this list over the course of the day, and also to share their most significant fear(s) about the future of the Internet (either from this list or not) in an anonymous online survey. Since a significant risk of failure is what potentially warrants government intervention, we encouraged pessimism – concerns that might drive the Internet in undesirable directions. We combined these survey responses with our list and used the integrated list as a basis for discussions on Day 2.

In parallel with discussion of fears and aspirations, we followed a format similar to previous years, in which attendees engaged in short presentations and focused discussions on pre-selected topics. This year, our topics were focused on a deeper interdisciplinary understanding of how Internet interconnection works, including dimensions of measurement, modeling, and policy.

3. INTERCONNECTION: NORMS FOR NEGOTIATING

Christopher Yoo (U. Pennsylvania), talked about three often-

conflated roles of pricing in a simple single-sided market: allocating scarce capacity, providing incentives to conserve, and signaling a need (or incentive to sell) to bring markets back into equilibrium. In a two- or multi-sided market, prices are much more complex to interpret, since pricing on one side of the market is determined by elasticity on the other side, and profit potential on one side may drive prices lower on the other side. He also emphasized that bargaining power over prices does not necessarily warrant intervention; it may have positive root causes, such as auspicious foresight in investment and risk-taking, as well as negative, e.g., monopoly power. The challenge for policy is disentangling good from bad sources of bargaining power so as not to inhibit the former. He used the Comcast-Cogent-Netflix dispute to illustrate how multiple sources of strategic behavior in a multi-sided market implies that pricing information alone is insufficient to determine if bargaining power is being used in an abusive way.

Ignacio Castro (IMDEA) distinguished the trend toward *remote peering* – when networks peer via a layer-2 provider that bridges the remote locations – from the trend of Internet flattening in terms of shorter AS paths. Remote peering has emerged as a traffic engineering option to mitigate the cost of supporting 30-40% annual traffic growth rates. Georgios Smaragdakis (MIT/TU Berlin) summarized recent work with Akamai exploring how ISPs and CDNs could collaborate (i.e., share information) to achieve a 10-20% improvement in efficiency and cost of content delivery by coordinating on server selection across a growing ecosystem of CDNs. Operators in the room pointed out that information sharing already supported cooperative delivery of traffic, e.g., peering agreements have requirements to add capacity when traffic levels reach a specified threshold.² David wondered if anyone had an idea how much efficiency could be further extracted from the delivery of content, beyond the 10-20% that Georgios discovered; it seems there are no clear substantial value creation opportunities in the content delivery market.

Discussion following these talks led to several inter-related observations about transit and interconnection economics. First, the opacity of the industry limits the ability to model or even study it. There is not even historical much less current data that could shed light on the value of peering relationships, which is a problem for players as well as policymakers (and researchers). Second, tremendous advances in photonics have created an optical transport and switching industry where transmission costs are largely unrelated to volume of traffic transported. Anyone with a dark fiber facility today can manufacture capacity at incredibly low incremental cost by changing optical endpoints out, putting strong pressure on transit pricing, which in turn can support massive growth in traffic. Third, required financial disclosures reveal that interconnection costs pale in comparison to content acquisition costs, at least in the United States. (In other countries many content owners are state-owned, making the economics different.)

These observations led us to a higher-level discussion of the shifting cost and revenue structure of the industry, and how some aspects differ by region. For example, there is more competition at the retail broadband level in Europe, but large carriers there struggle with profitability compared to content providers, in part due to an antiquated cost structure (too many employees) but more generally (and globally), while bits on the wire today differ widely in their value, they cost the same to transport. For example, providers of ads generate substantially more value per byte than video. Because the Internet originated without commercial operation, the concept of extraction of value from the bits has not historically been

² “When traffic grows 40% annually, you can’t not have cooperation.”

part of the norms of interconnection. The indifference to what sort of traffic is being carried, i.e., a de facto common carriage model, encouraged innovation and diversity in interconnection. Both are at risk with a possible shift toward value-based traffic exchange.

Although price discrimination (i.e., based on value) is known (to economists) as the only way to achieve economic optimality in a perfectly functioning market, broadband is not a perfectly functioning market. Since different users may assign different value to the same sorts of traffic, assigning an interconnection price based on some sort of *average value* may reduce the uptake of content by users with low but non-zero valuation. A similar argument about pricing of cable television channels is used to justify “content bundles”, where a bundle effectively assigns an “average” value to a collection of content, rather than to each separate element of the content.

Most of the discussion in this session was about efficiency gains from long-term investment decisions rather than short-term pricing decisions. Economists emphasized that the most important thing is to get pricing incentives right so that people make the right investments. There was concern that current pricing models for interconnection do not send the right economic signals to make wise long-term investment decisions.

4. MEASUREMENT AND MODELING

The rest of the sessions on the first day focused on measurement, modeling, and regulatory issues around interconnection.

Steve Bauer (CSAIL/MIT) began with a review of different definitions of congestion in operational networks, and some analytic techniques for identifying congestion. He noted that network engineers have several definitions of congestion, not all potentially harmful or avoidable. Even infinite bandwidth networks can experience congestion, due to protocol design and implementation issues. Amogh Dhamdere (CAIDA/UCSD) briefly summarized a recent joint CAIDA/MIT study of a type of persistent congestion that had been in the news lately: points of interconnection between two networks that are persistently congested (for hours per day, over weeks or months) because both networks believe the other peer should fund the costs of the upgrade to alleviate the congestion. Amogh noted a prevailing emphasis on harms to consumers from congested interconnection links, but little attention to potential harms to companies involved in such interconnection disputes, and ways to quantify them. He suggested the need for a metric to quantify the financial impact to various players of having a link congested for a certain period of time.

Srinivas Sakkottai (TAMU) then described some new work he is undertaking on dynamic markets for wireless congestion pricing. He noted that congestion pricing – where users pay more during periods of congestion – has been quite common historically for networked systems: roads, public transit, smart grid, to promote certain kinds of usage. Yet cellular data systems have not thus far used congestion pricing, despite the surge in demand for cellular bandwidth. Instead they degrade performance or access when a usage cap is exceeded, impeding high and low value applications in the same manner. As many have observed, this kind of neutrality is not necessarily what the user would prefer, and it makes sense to explore dynamic market mechanisms that allow for users to give preference to high-value data flows in times of congestion, to better match user preferences to allocated resources. Such mechanisms could improve the efficiency of both primary markets (access to end-users) and secondary markets (end-users sharing bandwidth via hot spots and device-to-device networking.)

Richard Ma (NUS) presented work on theoretical modeling of a practice he calls “subsidization competition”, where in a two-sided

market content providers voluntarily subsidize the usage-based payments users make to their access providers. His premise was that the lack of payment from content and edge providers to access providers is effectively a cross-subsidization of the content and application industry, which has enabled great innovation at the edge of the Internet, but reduced economic incentives for access ISPs to expand capacity. His results indicate that the practice of subsidization competition can increase the competitiveness and welfare of the content industry, although an uncompetitive access market might still warrant access price regulation. Julien Mailland presented a contrasting perspective: what he saw as a natural tendency for two-sided markets to lead to gated and otherwise content-filtered communities, as access providers move to favor content that has paid for access.

The modeling presentations triggered a lively dialogue when we challenged participants to articulate what they had learned so far from them, and one engineer declared, “I learned what a great threat to the Internet economists are! If we have used these models 30 years ago we would have killed the Internet! There would have been no incentive to innovate, no growth, no telephony, no streaming video. If Fedex charged a percent of the value the cargo in the package, what would it do to our idea of transport?” Others echoed Srinivas’ point: if allocation must be identical, and one cannot pay more for high-value applications, this can also hinder innovation.

An economist with an engineering background acknowledged the frustration with economists, but cautioned that economics and engineering are two sides of same coin. One can map any engineering into economic problem and vice-versa, but often how to connect the two is lost in the debate, in part due to discipline-specific terminology. An attorney in the audience joked, “Well, I’m thankful for the economists, they share the space underneath the bus with the attorneys,” but also observed that the lack of definitional agreement on congestion was fascinating. There is no legal definition or policy via the FCC, which leaves much room for semantic gaming on adjudicating disputes. There was general agreement on this point, as well as on the information deficit, since so much is proprietary.

Although this session provided a sample of how academics approach simplified modeling of interconnection, it was clear that a broader analysis of aspirations would require a process for assessing tradeoffs more comprehensively and systematically, e.g., weighing them, assessing costs. In doing so, there are always concerns about the quality of evidence available, and how to relate theoretical models to the real world. There were also concerns that most of the discussion of design and pricing models revolved around streaming video, with no attention to other types of usage, including the trend toward traffic symmetry resulting from symmetric (fiber-based) bandwidth broadband and cloud-based services. Nonetheless, participants saw great opportunities for interdisciplinary work, e.g., industrial engineering theories about cooperative games would benefit from more knowledge about vectors of cooperation between CDN and ISPs.

5. EMERGING REGULATORY CONCERNS

To punctuate as well as frame the broader discussion about fears and aspirations on the second day, we held a session focused on emerging regulatory concerns and the growing gap between regulatory frameworks and the technology they regulate, inevitable given their respective rates of evolution. We drilled down into one specific potential emerging ecosystem dynamic – *cloud interconnection* – that might warrant regulatory attention. We then stepped up a level to contemplate the ambitious goal of telecommunications regulatory reform, and what elements from existing law must be considered.

Bill Lehr (MIT) first led a discussion on his view of what *cloud interconnection* might look like, including the possibility of local community clouds serving the last mile. A cloud service provider noted that cloud infrastructure gravitates where the economics make sense – to regions with cheap land, cheap energy, and cheap dark fiber. Local storage environments would have to compete against the scale of cloud environments today – and not just economies of scale of infrastructure, but also of its management, i.e., centralizing IT talent within one administrative domain. Bandwidth availability also drives or constrains cloud environments – many people maintain large storage facilities at home because the upstream bandwidth is too low to maintain most of their data in the cloud.

Srikanth Sundesaran (ICSI) gave us an appreciation for the difficulty of measurement of cellular networks given their multiple layers of opacity. Using ICSI's Netalyzer tool, his research group has nonetheless discovered many unsettling phenomena (e.g., re-writing of packets) in higher layers of mobile networks.

We reviewed principles of U.S. communications regulation, building on the assumption that eventually Congress will deal with the lack of regulatory convergence in our four communications systems: telephony, cell phone networks, cable tv, and Internet, (governed by Title II, Title III, Title VI, and TBD³). We discussed issues from the existing Act that would likely need to be covered in any rewrite: connection upon reasonable request; just, reasonable, and non-discriminatory pricing and practices; measurement reporting; privacy; unwanted traffic; infrastructure sharing; duty to interconnect; switching costs; interoperability; negotiation; universal service; competition; device attachment; access to public infrastructure; educational use (public broadcasting); and information diversity. When discussing the history and likelihood of future regulatory reform, Christopher suggested that relative to other legislative reforms, such as the 1976 Copyright Act and financial services reform (both of which took over a decade to write from scratch) the 1996 Telecommunications Act was a relatively marginal update.⁴ Today it is hard to envision a complete re-write, because there are so many more stakeholders, divergent interests, legal processes driven by those interests, and (since 2006) partisanship on this topic. There was agreement among participants that “the Internet” is too amorphous and extensive a concept to regulate, so we needed to increase the conceptual separation of broadband from the Internet, and narrow the scope of regulatory focus to bottlenecks that may exist, e.g., conduit/outside structures for wired media and spectrum access/outside structures for mobile.

6. HOPES AND FEARS: REVISITED

At the end of the first day we collated the results of the survey question “What is your biggest fear for the future of the Internet?” with our existing list that we presented the first day, and we spent the first part of day two presenting and discussing the integrated list. For each issue, we considered what tools or approaches could encourage a positive outcome, or prevent a bad one, and what danger signals would justify intervention? One strong concern was that the current political economy may conceptually box us in, limiting the freedom with which we consider redesign of either regulatory or

technology frameworks. Another concern was the risk that government attempts to promote any given set of aspirations could slow down growth of the Internet and/or innovation using the Internet, or impede progress on other aspirations. We did not go into much depth on most of the aspirations, but got far enough to conclude that future workshops drilling down further into the aspirations would be helpful. In support of the continued conversation, the two of us compiled our initial analysis of these aspirations in May 2015 2015, along with this report.⁵

1. Reach/Ubiquity/Uptake/Affordable.

We asked whether there were fears that the Internet would not expand to reach everyone. Many developing countries are supporting or launching projects to build out Internet access, and operators noted that progress was slow but steady in regions with functioning governments that have little or no financial stake in the telecommunications infrastructure in the region. Problems arise in developing nations with monopoly network infrastructures that are run by kleptocrats with no incentive to optimize costs. People refer to the great promise of mobile competition, but the reality is that wireless networks are a little bit of wireless and a lot of wired infrastructure, so kleptocracy over the wires fundamentally constrains the ability to expand reach. A shared perspective was that ITU rules prevent use of the most valuable resource available to achieve Reach by limiting use of spectrum to its original allocation, e.g., broadcast, which inhibits its use for extending Reach. (The developed world has been known to ignore ITU frequency allocations to support better uses for the spectrum.) In developed regions, one ISP noted that when building out new infrastructure, people already connected to the Internet seem to care more about getting the unconnected people connected than the unconnected people care themselves. Access providers have found it sometimes difficult to convince unconnected people to subscribe, even at a subsidized price. We discussed an FCC survey from 2010⁶ which revealed that the biggest trigger for adoption of broadband is family and friends pushing one to get online.

There were also concerns about whether concerns about high performance, e.g., access bandwidth to support on-demand video streaming, obscured our focus on connecting the 4 billion world citizens who are not connected. We had previously discussed zero-rating services as a popular mode of expanding *Reach* in the developing regions, although it implies that people in those regions pay for access with their privacy (i.e., personal information being used to target ads) rather than a subscription fee. A risk of zero-rating services is that users may not be aware of this trade-off, and they may also never learn of other services. Improved transparency could inform consumers of the (non-monetary) costs of “zero-rating” services, including how data is used.⁷

2. Equity/Local Values/Universal Values/Global

There were several fears about the future character of the Internet. One concern mentioned earlier was how much video-on-demand frames the evolution of architecture and pursuit

³The 26 Feb 2015 FCC ruling on Title II followed this workshop but pre-dated publication of this report.

⁴The 1976 Copyright was motivated by a 1964 Supreme Court ruling that cable retransmission of local signals was not performance and thus did not require compensating broadcasters for retransmission. Broadcasters were then motivated to help get the Copyright Act re-written to clarify the opposite, which was passed 12 years later. Notably, the need for an update for the 1934 Telecommunications Act was first proposed in 1973 (25 years before it happened).

⁵http://www.caida.org/publications/papers/2015/inventory_aspirations_internets_future/

⁶“Broadband Adoption and Use in America, Results from an FCC Survey”, <http://transition.fcc.gov/DiversityFAC/032410/consumer-survey-horrikan.pdf>.

⁷One participant noted an analogous transparency requirement the FCC imposed on the wireless industry: requiring the wireless carriers to identify the cost of the handset alone vs. the subsidized leasing of it in the contract.

of innovation, rather than the ability for every user to be a content producer and distributor. Other attendees agreed there were higher priorities in the public interest, including affordable access, and high quality educational content.

Another fear was a global trend toward fewer open and more closed platforms, and similar dynamics that would turn the internet into a bunch of fragmented gated communities. Another attendee observed that the Internet is already fragmented and will fragment further no matter what we do, for both economic and political reasons. Certainly geo-specific differences will force the Internet to evolve in divergent manners across the globe, which some feared would lead to persistent and growing inequity, globally and domestically, and an inability of low income communities to meaningfully and productively use the Internet.

A related suggested aspiration was *empowerment*. We discussed where applications come from that create opportunities for people, e.g., what opportunities Facebook vs. Wikipedia provide, and at what cost.

3. Evolution/Innovation/Generality.

of the cloud), operators might have considered the bottleneck *Evolution* is a means to a goal rather than a goal in itself – perhaps *Innovation* or *Generality* are the ultimate aspirations and infrastructure evolution promotes them. It is also not obvious which aspect of the ecosystem most urgently needs to evolve at a given time. The discussion of infrastructure evolution to support innovation naturally led to the topic of investment, and the fact that the the global Internet including the cellular infrastructure for the last few decades have been funded mostly by revenues from voice telephony. The Internet itself exists due to a relatively small U.S. government investment in an operational research capability, along with a massive capital investment previously made in the (monopoly-profitable) telephony infrastructure. At this point, the future growth of the Internet cannot depend on infrastructure initially developed for other purposes. At the same time, the government is not strategically focusing research funding on areas (such as optical networking) that are key to growth. The future is in the hands of the private sector, both for deployment and for R&D. This fact raised concerns about the (im)balance between application and infrastructure innovation, specifically the over-investment on highly visible and profitable applications, and under-investment in research and technology innovation that can support infrastructure development.

There was concern about capturing the full potential of the Internet. As the Internet becomes the de facto “control plane” for the economy, it also is shifting in architecture to include not only traditional networks but computing and storage facilities. As these these different components co-evolve, the composition and location of tensions and problems will shift over time.

4. Trustworthy/Lawful/National Security.

There were fears that loss of trust in the Internet would limit its uptake and utility, since users are only likely to value a service highly if they trust it. Another fear was that the Internet is not stable and secure enough to withstand the attacks that will come given how embedded it is in so many aspects of our lives. Attendees acknowledged the existence of ways to improve the trustworthy character of the Internet that depend on larger ecosystem context, e.g., norms, institutions, markets, information.

Many feared that economic incentives in an Internet-of-Things

future would lead to systems with extremely insecure components, much like today’s home WiFi routers, which are typically not field upgradable and stunningly insecure. The reality of security today lends credibility to the argument for walled garden. A more abstract but deep problem is that Western countries’ perspectives on national security or trustworthiness tend to co-exist with Local Values, specifically free flow of information, which some countries see as a direct threat to their own national security since it facilitates organized dissent. As in meatspace, security-related aspirations vividly illustrate the need for a discussion of tradeoffs.

5. Unblocked

Privacy and invasive surveillance were a concern, although a counterpoint fear was that privacy concerns might kill the vibrant world of advertising-supported (free) apps. A recurring theme was that zero-rating services might promote some aspirations at the expense of others, e.g., using pricing to inflict censorship (of everything not zero-rated).

6. Choice and Competition.

There was much concern with inadequate competition and lack of choice in broadband access, and the implications for (high) prices for consumers. There was also concern about the inauspicious interplay of competition, regulation, and innovation. One outcome might be adverse implications for the character of the Internet: loss of access to content, and freedom of expression online.

One participant was thoroughly disappointed that our society (or at least people in this workshop) had apparently given up on enhanced competition as an important tool for achieving our many aspirations. He reminded us why we should keep competition in the toolbox: the variety of conflicting objectives, since consumers value services differently; and rapid evolution of technology, business models, and edge services. In this situation, heavy-handed regulation might achieve short term objectives (although we are all skeptical), but it cannot reliably achieve long term (dynamic) objectives: investment efficiency, innovation, risk-taking incentives. Conversely light regulation might achieve long-term objectives, but consumers will be forced into situations where they have little or no recourse, resulting in higher prices and reduced innovation and entry of new end uses. Competition can potentially achieve both short-term and long term objectives.

And although competition may be a lost cause in the U.S., this participant felt strongly that we (especially academics!) still have an obligation to state clearly that the first best solution would be to have meaningful competition at all levels of the stack, to the extent that it can be achieved. Policy in the U.S. and elsewhere should keep a long-term goal of enhancing competition, and pursuing it at smaller scales where it is more feasible – cities, states. Meaningful retail competition requires either three meaningful competitors, preferably using different technologies, or a complete disruption of the existing landscape. Promoting competition first requires understanding the importance of legacy business revenue flows for funding investment in shared infrastructure that is essential to growth of new markets and services. We should be concerned as to where the value creation is in over-the-top entertainment TV services, and how to shrink capital requirement to sustain optimal investment? As one attendee noted, if all of this is just about the TV industry then it is really not an issue for major national industrial policy.

7. Making Progress: Incentives and Transparency

We asked if there were fears about making progress toward the future. One fear was that regulation (or business models) would ossify current roles and practices, thereby foreclosing on future flexibility. Another worried that we tend to solve important issues like congestion and poor performance in a rather ad-hoc way without any fundamental changes to their root causes. So they repeat.

There were comments about measuring progress: whether the right things can be measured and monitored, such as economic behavior, business practices, alternative architectures. This concern escalates as networks become more closed, since it becomes harder to measure behavior (and it is worse in mobile networks).

The issue of incentives received attention. One attendee asked how we can align the incentives of different players, e.g., ISPs and CPs, so as to achieve social goals, e.g., user utility. Another worried that current approaches to regulation may be counter-productive: treating all bits in an identical fashion might stifle innovation in unpredictable ways. One attendee feared regulation in general, since “in the quest to stay busy, policymakers will take action that causes more harm than good”. Another worried that we might end up with “management of Internet driven not by sound engineering principles and/or economic analysis, but by public opinions, which may be ill-conceived”.

We had a brief discussion about how transparency can help users make better choices? One attendee challenged us to articulate exactly what information about performance characteristics or network practices or pricing and terms will drive innovation.⁸

7. CONCLUSIONS

The participants in this workshop, with their diversity of views, concerns and beliefs, are probably a good reflection of the larger population of informed players in this space. In some respects, the diversity of opinions was a familiar landscape. The larger conclusion is that the Internet ecosystem is complex, with many players, many dependencies, and a number of what might be called “cross-player subsidies”. The economics are embedded in what many would agree is an imperfect market, but most markets are imperfect, and the overall vitality of the Internet suggests that the system has, up to this point, sustained amazing progress. Most of our fears are not about what has already happened, but what might happen in the future—that somehow one set of actors or another will either acquire and abuse market power, or on the other hand lack the incentive to invest at the level that the overall ecosystem might seek. In fact, there is nothing new in the observation that investment in one sector (e.g., broadband access capacity) can gate progress in other sectors (e.g., bringing high quality video to the market). At a high level, the concern about regulation is that it might lock one set of players into a set of constraints that might eliminate their ability or incentive to adapt as the overall system evolves. At the same time, many participants articulated strong fears of abuse of market power, and an erosion of the positive aspects of the Internet as commercial interests shape the future. We are depending for the future we want on a set of actors we do not necessarily trust to bring us that future.

8. WORKSHOP PARTICIPANTS

- Co-Host: kc claffy, (CAIDA/UC, San Diego)

- Co-Host: Dave Clark (MIT)
- Rebecca Arbogast (Comcast)
- Johannes Bauer (Michigan State)
- Steve Bauer (MIT)
- Roger Bohn (UC San Diego / IRPS)
- Ignacio Castro (Institute IMDEA Networks/ICSI)
- David Clark (MIT/CSAIL)
- Phillipa Gill (Stony Brook U)
- Scott Jordan (UC, Irvine)
- Craig Labovitz (DeepField)
- William Lehr (MIT)
- Jonathan Liebenau (London School of Economics)
- Richard Ma (National University of Singapore)
- Julien Mailland (Indiana University)
- Kevin McElearney (Comcast)
- Milo Medin (Google)
- Debasis Mitra (Columbia University)
- Patrick Ryan (Google / University of Colorado)
- Srinivas Shakkottai (Texas A&M University)
- Georgios Smaragdakis (MIT / TU Berlin)
- Andrea Soppera (BT)
- Srikanth Sundaresan (ICSI)
- Tony Tauber (Comcast)
- Sanjay Udani (Verizon)
- Qiong Wang (University of Illinois at Urbana-Champaign)
- Christopher Yoo (University of Pennsylvania)

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-1414177: “Mapping Interconnection in the Internet: Colocation, Connectivity and Congestion.” This report reflects the views of the authors and some workshop participants, but not necessarily the National Science Foundation.*

⁸ “If you cannot come up with any such information, transparency will not solve the problem.”