The Internet Ecosystem - The Potential for Discrimination

Dirk Grunwald

University of Colorado at Boulder

Department of Computer Science

Technical Report

CU-CS-1062-10

March 2010

University of Colorado at Boulder

Department of Computer Science

430 UCB

Boulder, Colorado 80309-0430

www.cs.colorado.edu

The Internet Ecosystem - The Potential for Discrimination

Abstract

The Federal Communication Commission is considering rules enforcing "network neutrality" and legislation proposing similar goals have been discussed in Congress. The goals of the proposed regulation and legislation are preserve an "open Internet", but are specifically directed toward access networks, or the first link that directly connects users to the Internet. We argue that preserving open competition in a host of "higher level" Internet services is equally if not more important, but since the rate of technology innovation typically out-paces the need for regulation, there is no need to impose regulation at this time. Using specific examples focused on the "visible Internet" as well as new services and applications that enable rapid innovation, we argue that the Internet has fostered a history of technological and business solutions that overcome what seems to be certain market dominance. A key enabler of these changes is the emergence of technologies that lower the barrier for entry in developing and deploying new services. We argue that regulators should be aware of the potential for anti-competitive practices, but should carefully consider the effects of regulation on the full Internet ecosystem. We believe that consumers will be better served through education, maintaining competitive environments and technical forecasting.

The Premise Behind Network Neutrality

The premise behind the current debate in network neutrality was articulated in an FCC policy statement adopted in August, 2005[1] that stated four goals for the Internet:

1. Consumers are entitled to access the lawful content of their choice.
2. Consumers are entitled to run applications and use services of their choice, subject to the needs of law enforcement.
3. Consumers are entitled to connect their choice of legal devices that do not harm the network.
4. Consumers are entitled to competition between network providers, application and service providers, and content providers.

Proposed rules would extend these four core principles with two additional rules:

5. A provider of broadband Internet access service must treat lawful content, applications, and services in a nondiscriminatory manner
6. A provider of broadband Internet access service must disclose such information concerning network management and other practices as is reasonably required for users and content, application, and service providers to enjoy the protections specified in this rulemaking

Broadly speaking, participants in the network neutrality debate use the same term to conflate two issues – accessing content of their choice and, more narrowly, enabling the development of a competitive environment for services, applications and content providers by maintaining "neutral" access to the last link for consumers or the "public" Internet (the "access network").

The two primary concerns have been that access network providers would provide preferential treatment to specific uses of the network and may go so far as to block certain kinds of applications. To support this concern, proponents of regulation point to a small number of documented cases where ISPs have blocked specific services (VOIP and file sharing). There is concern about a lack of transparency in network management and how that might lessen the current Internet or unfairly limit competition. But the ability to limit access to Internet applications is not restricted to access networks. Such restrictions can be enacted by many components used to access Internet content, such as the browser and services or applications within the Internet.

Likewise, there are many ways to enact preferential access. In (Grunwald 2007), we discussed aspects of current Internet access network designs that can lead to higher

---

1 FCC Policy Statement FCC 05-151, issued August 5, 2005, released September 23, 2005.

barriers for innovation and new services or can allow subtle forms of preferential network access. We specifically focused on *asymmetric access links* and *content distribution networks (CDNs)*. Asymmetric access networks[2] make it more difficult for consumers to "self-publish" and commercial content distribution networks[3] can effectively provide "preferential access" to content provisioned on a CDN located within an ISP's network without actually violating "neutral" access network policies.

We argued that these barriers impose as much risk as preferential treatment of access networks but that network neutrality regulation focused solely on access networks would be unlikely to address these barriers. Instead, the proposed regulations may hamper network innovation at the access network as well as the core of the network while still leaving open the door for anti-competitive actions that the regulations are intended to forestall.

In this paper, we explore other parts of the Internet ecosystem and how they affect open and competitive networks. There is broad consensus that layers of the Internet ecosystem other than the access network may impact competition and innovation; the question remains if new rules are needed. In the conclusion of a paper describing the economic history of price discrimination in telecommunications networks, (Odlyzko 2009) wrote:

> "For telecommunications, given current trends in demand and in rate and sources of innovation, it appears to be better for society not to tilt towards the operators, and instead to stimulate innovation on the network by others by enforcing net neutrality. But this would likely open the way for other players, such as Google, that emerge from that open and competitive arena as big winners, to become choke points. So it would be wise to prepare to monitor what happens, and be ready to intervene by imposing neutrality rules on them when necessary."

Odlyzko's point was that what he termed "cloud computing"[4] would become a more important marketplace for innovation than services integrated into access networks; his implication mirrors ours -- focusing on those access networks may distract from anti-competitive behavior in those other markets.

In this paper, we agree with Odlyzko's observation that other parts of the Internet ecosystem are equally powerful in determining the rich competitive environment of the Internet and show for past, current and emerging parts of the Internet. At the same time, we argue that regulation and action, either that proposed for the access network or

---

[2] Most broadband access networks have higher download speeds than upload speeds; e.g. DOCSIS cable models typically support ~8mb/s downloads but 0.4Mb/s uploads and variants of xDSL technology have similar asymmetries. These communication asymmetries make it difficult for consumers to host services in their home or to generate content.

[3] Examples of "Content Distribution Network" (or CDNs) include Akami, Limelight and Amazon Cloudcast. These services make multiple copies of content available at multiple physical locations in the Internet, improving the experience of accessing that content under periods of high demand.

[4] By this term, Odlyzko meant software services hosted on computers not located at a persons home or business; later we'll see that current common usage has two meanings for this term and disambiguate those meanings.

extending beyond those networks (through ambiguity or design), should only be applied when clear harms are shown. The development of specific technologies, coupled with the pace of technology development, the continued innovation of the Internet community and use of existing laws has served the Internet well.

The FCC Notice for Proposed Rule Making released in October 2009 attempts to insure a competitive market place, but it does so through regulating one subset pf providers and certain, specific network characteristics such as traffic priorities[5] and managed services (having multiple services use a single physical transport). This focus ignores the fact that the Internet evolves over time and is far from a finished work. In fact, the National Science Foundation, the national agency that has long funded Internet research, has launched multiple research programs to define the future Internet[6]. Extending the existing Internet is difficult, because it has become essential to society, but there are clear reasons to improve on the current design. Would regulation add yet more friction to process of improving the Internet? Are we doomed to the Internet of today?

Rather than use words like "discrimination", network engineers prefer terms like "network management" and "prioritization". One form of prioritization endemic to the Internet is "congestion control"; congestion occurs in a network when too many packets try to use the same resource (link or router). The Internet Protocol handles congestion by simply discarding packets when resources are limited. But congestion requires that the transmitter slow down, or the network can enter a "congestion collapse" whereby no useful communication takes place[7]. The original Internet design principles emphasized "end-to-end" control (Saltzer 84) and assumed that the computers at each end of a transmission would cooperate to prevent congestion collapse. In 1986, the network experienced a series of congestion collapses that reduced useful throughput by factors of 10-1000. New congestion control methods were introduced then, and have continued to be developed. Different congestion control methods, implemented on devices or working in concert with network routers, affect how competing network flows use the networks to

---

[5] The use of the word "discrimination" in the proposed rules is regrettable. From a technical perspective, discrimination can mean any form of differentiation, including simple traffic prioritization designed to improve performance; however, the word is laden with other meanings by events and history external to network engineering.

[6] The "Future Internet Directions" program (FIND) has funded research to address how parts of the Internet design need to change in response to new demands and technologies. The NSF GENI program (Global Environments for Network Innovation) program is funding the development of test platforms and new technologies for future Internets. Similar efforts are underway in Europe, Japan and other countries as well.

[7] For example, assume two transmitters are trying to use a single common link that has a capacity of 100 packets/second. Both transmitters want all of their data to be received and will re-transmit packets if they are discarded. If one transmitter injects 100 packets per second on to that link while the other injects 10, some packets will have to be discarded. Assuming a random discard policy, 91% of the discarded packets will be from the higher rate transmitter. If the transmitter determines that those packets were dropped, it would retransmit those packets in addition to the existing 100 packets/second, resulting in increased congestion. As more and more packets from the faster transmitter are dropped, it will increase sending rate until its access link capacity is reached. This "congestion collapse" insures that and increasing number of packets are discarded and also negatively affects the slower transmitter, because its packets will make up an increasingly dwindling portion of the packets that traverse the congested link.

improve the overall efficiency of a complex, distributed and decentralized system. Would this research and innovation be possible with the proposed FCC rules in place?

Although the Internet is 40 years old, the commercial Internet is only 15-20 years old. New applications and an increased number of users changes assumptions network engineers have made and exposes the network to new challenges with the concomitant need for new solutions. In an effort to maintain a rich Internet environment, the proposed regulations focus on access networks without considering how anti-competitive pressures can be applied in the remainder of the Internet. It also regulates a mechanism (traffic prioritization) that is used in congestion control, but at the same time is part of the basic Internet design. Likewise, although the FCC NPRM addresses the distinction between the "managed" and "public" Internet, it does so in a limited way that may hamper innovation in "managed" networks or in the Interface between private and public networks.

We argue that there are better ways to maintain a vibrant Internet. These include: having clear standards and methods for measuring what is actually happening in the Internet as well as methods for reporting or disseminating policy to consumers; use existing agencies and policies; encourage innovation and competition for access networks; and, develop "best practices" that can be clearly understood by network operators, regulators and consumers.

Risks To The Internet Ecosystem

The Internet is composed of many parts that make up the "experience" that end users now confront. Just as the phone network is made more useful by 411, white pages, yellow pages, 911 and other services or applications, the Internet is made more useful by domain names, browsers, search engines and services that are integral to the web. Insuring competition and a rich Internet environment by *solely* focusing on the local loop, as is being done with the Internet, clearly misses the mark – the entire "ecosystem" that influences either network experience is important.

To understand how applications and services can foster an anti-competitive environment, we examine a series of past concerns about Internet exclusion and market dominance, starting with the platforms that enabled web access and stretching to services that now generate the most debate. These examples illustrate the rapid pace of innovation and demonstrate that the Internet often innovates its way out of anti-competitive markets; it also shows that even when that doesn't happen, existing laws and regulations enforced by the FTC and the Justice Department can level the playing field.

Access To The Web – The Browser

The web browser is an application that has had almost total market dominance by multiple companies at different times. One of the earliest graphical Internet web browsers was Mosaic, developed by students and staff at the University of Illinois. The Mosaic developers founded Netscape to commercialize the browser. Although other companies, particularly Microsoft, developed other browsers in the mid '90's, Netscape maintained ~80-90% share of the browser market until Microsoft bundled its own product, Internet Explorer, with Windows 98. Netscape's fortunes quickly soured as Internet Explorer

reached a 90% share of the browser market; Internet Explorer now has 63% market share, having lost share to browsers developed in the last 5 years[8].

It's rare for a market to switch from total domination by one product to another so quickly. However, as Netscape discovered, the problem with marketing a browser was how to monetize the product. Most businesses were hoping to use the browser to steer users to specific web properties. Open standards allow rapid substitution of one product for another and but can equally favor the adoption of software that "extends" those standards. Internet Explorer enabled Microsoft to launch protocols that favored other Microsoft products (either Windows desktops or Windows Server). Chief among these were "ActiveX controls", a mechanism to embed software unique to Windows in web page. Many of these "controls" provided mechanisms missing in the Web (such as audio or video); because ActiveX only worked with Microsoft clients, the use of such controls drove many to rely on Microsoft software. The combined control of the most common operating system and the pre-installed browser brought on anti-trust actions and an initial finding of monopoly power.

Although Internet Explorer still dominates the browser market, alternate services, new technologies and standards eliminated much of the threat of Internet Explorer. AOL eventually purchased Netscape and much of the code-base was spun off into the popular open-source "Mozilla" and later "Firefox" browser platform. Additional vendors, primarily Apple, Opera and now Google, produced other competitive browsers. Increased broadband speeds and better software installation and update processes made it easier to install competing browsers. At the same time, browsers became ubiquitous, becoming a universal way to access and control devices ranging from printers to alarm clocks; manufacturers wanted those controls to be universal. A wide-spread "open standards" effort ensued to identify browser techniques that limited users to Windows based computers; lobbying and branding by the World Wide Web Consortium (W3C) led governments and many companies to eschew IE-specific mechanisms to focus on a "works with any browser" standard. At the same time, the development of "Web 2.0" technologies such as AJAX around 2004[9], coupled with increased broadband speeds, meant that many of the Microsoft-specific "ActiveX controls" could be replaced by software that worked across all browsers. The impetus for a standards based browser has become particularly important as web browsers have become an integral part of mobile phones that are unable to use Windows-specific features, such as the iPhone.

Although Internet Explorer still dominates the browser market, that dominance connotes little economic advantage to Microsoft at this point; the majority of Microsoft profits still arise from sales of Windows and Office rather than on-line products. However, without the development of alternative software and open standards by organizations such as

---

[8] Browser adoption rates are highly regional. Gregg Keizer "See Google's Chrome grabs No. 3 browser spot from Safari", Computerworld Jan 2 2010, http://www.computerworld.com/s/article/9142958/Google_s_Chrome_grabs_No._3_browser_spot_from_Safari

[9] AJAX is a term used to describe one way in which "rich" web applications are developed using nothing more than standard web browser protocols. Gmail, released by Google in 2004, was one of the first widely known AJAX applications. Jesse James Garrett (who coined the term AJAX) has a readable description of the technology at http://www.adaptivepath.com/ideas/essays/archives/000385.php.

W3C, the present situation might not have come about and could rapidly change. It's arguable that the anti-trust investigation of Microsoft was what led to the current situation. It's equally plausible that the development of mobile phones, and the demands of that emerging non-Windows ecosystem, or the deployment of broadband and more interactive web pages using AJAX forestalled the dependency on Microsoft-specific features. One thing is certain; competition, innovation and existing legal recourse opened access to the Internet without the next for additional regulation.

Rich Internet applications, Video and the New Content Companies

Less well known than the "browser wars" is the (on-going) battle for "rich Internet applications"[10]. RIA is now a fundamental part of the Internet ecosystem. These environments provide extended usability to systems like Google Mail, NetFlix, Hulu, Microsoft Live, Yahoo News and many other websites – RIA allows conventional "desktop" applications to be replaced by web-based applications. The features that made Internet Explorer indispensible in many areas were for "rich web applications"; RIA environments make that approach work across different operating systems. Microsoft sought to use the Windows infrastructure to allow developers to use existing Windows code in Web applications. The primary alternative approach was Java, developed by Sun Microsystems by which programmers could develop "applets," or programs that ran within a web browser. Although the Java language found extensive use in business software, applets experienced limited success, largely because the process of installing software was relatively complex. Macromedia Flash was introduced in 1996 and rapidly became the primary RIA tool; it's currently installed in more than 90% of browsers and is used to power many video and on-line game sites. Later entrants were Microsoft Silverlight (similar to Flash and Java) and Adobe AIR (developed as an extension to Flash when Adobe acquired Macromedia).

Surprisingly there has been little concern to date that any of these alternatives would preclude effective competition. In large part, this is because there are "open source" implementations of the dominant platform (Flash) and any one system is largely substitutable for the other (although not always on the same device). More importantly, existing and new standards based technologies are replacing many of the functions for which developers turn to RIA frameworks. This point was argued by Microsoft in a 2007[11] response to a motion by the State of California that Microsoft's development of Silverlight should extend the earlier anti-trust actions[12] . Some Microsoft web services (such as Bing! 3-D maps) still require SilverLight and Active-X controls. Others[13] argue that the required use of Silverlight for specific high-profile events (Olympic events,

---

10 See Jim Rapoza, "RIA War Is Brewing", Eweek Emerging Technologies, available at http://etech.eweek.com/content/application_development/ria_war_is_brewing.html

11 See Charles Rule, Bradford Smith, Steven Holley, "Microsoft's Report Concerning The Final Judgments", filed in Civil Action No. 98-1233 (CKK), August 31, 2007

12 An analysis of the filing is available by Todd Bishop, "Antitrust Filing Cites Microsoft Silverlight Concern", Seattle Post Intelligencer Blog, Oct. 17th, 2007, http://blog.seattlepi.com/microsoft/archives/123837.asp

13 See John Markoff, "Microsoft leveraging Silverlight and riling critics", August 11, 2008, http://www.nytimes.com/2008/08/11/technology/11iht-stream11.1.15135139.html?_r=1

Presidential Inauguration) and bundling of Silverlight with Windows 7 will raise the same anti-competitive issues that Netscape faced in the '90s.

The argument that "open" alternatives suffice is compelling; most of the applications by Google rely on JavaScript, a programming language that has long been a standard tool embedded in web browsers[14]. Rather than develop a new programming environment, Google, Apple and Firefox have worked to greatly increase the usefulness of JavaScript, making that standard tool more suitable for many "rich" applications. The web standards community also developed HTML5, the latest combination of the *lingua franca* of web browsers. That standard supplants many of the reasons RIA frameworks were needed, such as high performance video playback, access to geographic location as well as support for storing and accessing data *via* the browser. These individual components allow large changes to applications – for example, using HTML5, Google Gmail can function more like a standard email client allowing access to email even when not connected to the Internet.

This analysis of RIA environments serves to show how regulation decisions are interconnected by past technology. Had Microsoft "won" the browser wars, most of this innovation wouldn't have occurred – developers would have used Microsoft components rather than adopt a new RIA framework. This would have also altered the landscape of devices, such as the iPhone, that are used to access the Web. The competitive alternatives are so diverse and rich that government intervention isn't needed; but the past experience of the "browser war" shows that existing methods for intervention are possible and effective when needed.

Naming and Information Discovery

Names play a central role in the Internet. People need to be able to access websites and services. The Domain Name System (DNS), which translates names to IP addresses, is central to naming in the Internet. With the rise of the commercial Internet and the Internet Corporation for Names and Numbers (ICANN), ownership of domain names clearly related to existing trademarks and properties was disputed and a uniform resolution method was enforced. Naming is one of the clearest cases of regulation applied to Internet services, and a number of national and international laws, rules and bodies have been created to address names, particularly as applied to commercial interests.

Today, search has taken on the importance originally attributed to DNS names. No part of the Internet Ecosystem would appear to be as important as search and search is a now universal way for finding new information, even supplanting the common use of domain names – many of the most common search terms on Google are the names of (often competing) web services, indicating that users rely on search for even trivial or well known information.

Should search be regulated? Recently, there have calls for such regulation (Pasquale 2008), often based on the dominance of a single search engine. While this argument is

---

[14] It should be noted that the development of JavaScript was not without contention. Netscape initially developed JavaScript; Microsoft developed a competing version and submitted that version for standardization. Rather than splintering web standards, Javascript came to unify them.

similar to that of DNS, there is a distinct difference – DNS was a single system essential to the core operation of the Internet, while *e.g.* Google is one of many search services. More over, search services were not originally intended to identify commercial interests – they were intended to "discover information".

Although Google dominates current search services, there have been numerous popular search services over time – AltaVista, Goto, Ask, Yahoo and different Microsoft systems. The current dominance of Google (currently estimated at ~65-85% US market share[15]), coupled with the consolidation of on-line advertising, has led some to call for regulation of search engines and search-based advertising to make it "neutral". The key objection is that search (and Google specifically) is so influential on the way people find information that it constitutes a "gatekeeper" on the Internet. In one New York Times Op-Ed article[16], Adam Raff, founder of a company promoting an alternative search engine, describes how Google has promoted its own products (maps, shopping services, *etc*) over that of other companies in search results. It's difficult to know why a specific Internet tool (*e.g.* MapQuest *vs*. Google Maps) falls from favor. Clearly, advertising a service is one reason, but so are features and usability. It's difficult to simultaneously argue that customers are unlikely to flock to a new search engine but would rapidly switch to new mapping software simply because it is well advertised. Advertising drives the substantial growth of Google; existing anti-trust measures would seem to govern and appear to have been successfully applied in specific instances, such as to counter the proposed joint Yahoo!-Google advertising pact[17] and exclusive licensing of digital books.

Many of the arguments for regulating search are based on the difficulty of effective competition (Pasquale 2008). Search is composed of three main components – crawling, indexing and presentation. Crawling is the traversal of web pages – bringing the content of those pages to be indexed. Indexing records the information in the pages so that specific web pages can be quickly identified. Retrieval and presentation transform search requests into queries that search the indices and present the results to the users. Pasquale and Bracha (Pasquale 2008) argued that creating search engines is costly, but as with much of the infrastructure of the Internet, the software to develop effective and scalable search engines is now free. The Apache Foundation, an organization that manages the development of the free Apache web server also distributes Nutch, an open source search engine, and Lucene, a free indexing mechanism. Yahoo! has also donated Hadoop, software designed to rapidly index large numbers of web pages.

Although the software is free, adoption of new search engines depends on the utility they provide to users. This is usually based on the effectiveness of presenting the results of a search query. Ranking determines the order in which the most important search results

---

[15] As with web browser choices, different search engines are popular in different markets. See Hitslink.com: http://marketshare.hitslink.com/search-engine-market-share.aspx?qprid=4 (85%) or http://blog.searchenginewatch.com/100119-075446 (65%)

[16] Adam Raff, "Search, But You May Not Find", New York Times Op-Ed Dec 27th, 2009 available as http://www.nytimes.com/2009/12/28/opinion/28raff.html?_r=1

[17] US Dept. of Justice Press Release, "*Yahoo! Inc and Google Inc. Abandon Their Advertising Agreement*", Nov. 5, 2008, available at http://www.justice.gov/opa/pr/2008/November/08-at-981.html

are displayed. The "Goto.com" search engine pioneered the "money talks" policy of paid search rankings and Google "AdWords" expanded that base with an auction-based scheme. In many ways, the barriers presented by search engines and ad rankings are similar to Yellowpages. Businesses were at a disadvantage if they did not place paid advertisements in Yellowpages directories. One of the complexities that search companies face is that the variables governing advertisement (placement, frequency, relation to search, *etc*) are more complex than those used in static print media. Defining and communicating those characteristics, and having customers understand them is a complicated task. There is always a need for transparency so that advertisers understand what they are purchasing, particularly when competing "house brands" are also advertised, as Raff argued. This situation is similar to grocery stores that present their own house brand and a diverse array of competing brands whose placement is governed by a combination of consumer demand and "slotting-fees"; slotting fees have received much discussion (Aalberts 1999) as well as government scrutiny (FTC 2003) and enforcement actions at state and national levels (Gundlack 2005). It seems likely that anti-competitive behavior in search would encounter similar scrutiny, and the FTC has already asked companies to disclose paid search results.

Despite the dominance of Google in the search-based advertising market, the search market itself has seen considerable innovation, in part because there are many *corpora* over which to search and many methods to rank or present results. Real-time search, personalized search, social search and peer-to-peer search tools are in active development. OneRiot.com is a start-up that recently partnered with Yahoo! to develop "real time" search (or search about breaking events rather than historical documents) and Lijit.com is a search engine focused on blogs and social networking. Ask.com and Aardvark focus on casting human questions into search queries. It may be that no search engine could compete with Google in the sense of becoming a multi-billion dollar company; many will be acquired by existing search companies. But it's important to recognize that Google, as a company, is little more than ten years old. Given the low barriers to entry (other than customers), there should be continued innovation in search.

It's clear that search has become as important as naming in the Internet; it also influences the experience that users have because they have come to rely on the speed and accuracy of search to locate services. What's not clear is if additional mechanisms beyond current laws are needed to insure a competitive and innovative Internet.

Content Distribution & Cloud Computing - The Invisible Ecosystem

The Internet has visible components, such as the browsers, rich application frameworks and search engines we've discussed. Equally important is the invisible infrastructure that defines how the web and web services are implemented. In this section, we describe services that dramatically lower the barriers for creating new web services. Just as open-source tools such as Nutch, Lucene and Hadoop reduce the technical barrier for developing a new search engine or service, new business models and technology reduce the operational barriers to effectively deploying and scaling those services.

Content distribution networks (CDNs), co-location and peering arrangements are some of the most critical elements of the Internet ecosystem that affect the web as it is used today. A CDN is an organized network of computers that are often placed "close" to Internet

users. Commonly accessed content is then stored on those computers and requests by web users are directed to a "near by" or lightly loaded computers. Content distribution networks can be used to save bandwidth since the content for a popular item does not need to be fetched from a distant location; this was the basis for our original concern (Grunwald 2007) that solely focusing on the access network wouldn't prevent performance discrimination. However, with the drop in price for Internet bandwidth, CDNs have primarily become useful because they provide a way to provide *scalable* service. The canonical example for this is the success that Victoria's Secret (a retailer) had in hosting on-line content before and after using a commercial CDN. In the initial offering, demand for the retailers content exceeded the capabilities of their own web services but successive offerings using a CDN were much more successful[18].

The Web would present a very different experience without CDNs, but the use of a CDN provides as much opportunity to discriminate performance as subtle packet differentiation or "traffic shaping" on an access network. Indeed, comments in (Scherlis 2010) indicate that ISPs in China market their own content networks and hosting services as providing better access to their own clients. In a competitive market place, the difference in performance is less a conspiracy than the result of innovative network architectures. Different combinations of CDNs and network management lead to differing degrees of efficiency (Jiang 2009) but efficient network architectures can still enable competition. At the same time CDNs enhance the ability of a web company or organization to successfully connect with readers without having to invest huge sums in capital infrastructure.

In a Wall Street Journal article in 2008[19], Kumar and Rhoads argued that such "fast track" access violates net neutrality. The fact is that *most commercial content* on websites is distributed using CDNs and that there is significant competition in CDNs in the United States[20]. The proposed FCC rules don't seem to address the importance of this part of the Internet ecosystem. This is arguably good, because no concrete harms have been shown – indeed, the existing "fast track" access has enabled more companies to scale to meet web demand. But this highlights the rather arbitrary nature of the proposed FCC rules. The proposed rules would arguably also prohibit new services or offerings by "network operators" that could achieve the benefits of CDNs using different technical means, thus increasing competition in this segment of the Internet ecosystem.

Peering relationships between different ISPs, application providers and Tier-1 network providers also enable "fast tracks" for information; (Yoo 2010) has a readable description of modern peering architectures. Most of those peering relationships have been historically "settlement free" because they benefit both parties and traffic demands were

---

[18] A case study is available at from Akamai at http://www.akamai.com/html/customers/case_study_victoria.html

[19] See Vishesh Kumar and Christopher Rhoads, "*Google wants it Own Fast Track on the Web*", The Wall Street Journal, Dec. 15, 2008, available at http://online.wsj.com/article/SB122929270127905065.html

[20] Dan Rayburn maintains a list of current CDN vendors; as of July 29th, 2009, his "Updated List of Stand Alone CDNs and Telcos/Carriers Offering CDN Services" listed ~50 companies. See http://www.cdnlist.com for a current list. While this market is currently very competitive, consolidation is expected during 2010.

symmetrical. Increasingly, the line between "backbone", application and edge network provider have blurred. Google and large CDN companies such as Limelight now run some of the largest Internet backbones (Labovitz 2009). At the same time, "edge" network companies such as Comcast, AT&T and Verizon also carry considerable corporate or "non-public" network traffic. Amid the consolidation in networking companies, "paid peering" has emerged as a way to enable content providers or other co-location companies to reduce the *cost* of access while improving performance for their hosted partners[21]. Content distribution networks (and peering) improve performance; being excluded from such interaction would raise costs or limit competition. Reaching a sizable population would be possible but would require significant investment to be "scalable".

The proposed FCC rules don't clearly indicate whether peering and content distribution relationships constitute "neutral" access or in what situations they constitute "discriminatory" access. Again, this is arguably good, because there are few instances in which concrete harms have been demonstrated. In the past, the Internet has been "partitioned" because Internet providers could not agree on pricing for transit or peering relationships[22] and more consumers have experienced network problems from to these business disputes than those affected by the rules in the proposed FCC regulations. Is regulation needed to cover peering? History indicates that existing dispute resolution mechanisms (lawsuits, agreements and contracts) can resolve these problems. This lends credence to the argument that those same mechanisms will insure competition in other Internet services such as CDNs.

Just as CDNs developed out of a need to replicate and distribute "static" content, a new market, "Cloud Computing", has emerged as a technology that subsumes CDNs and facilitates even faster changes in technology. Cloud computing providers such as Amazon EC2, Rackspace, AT&T, IBM, Microsoft and several others run warehouse-sized data centers on which customers can lease and run customized software. Combined with "virtualization technology", which lets users capture the entire configuration of a computer in a form that can be shipped off to a remote data center, cloud computing has changed the economics of establishing Internet services. Cloud computing systems can typically be leased by the hour and new online services can be launched quickly. For example, in early 2010, Amazon's EC2 (a service that popularized the cloud computing model) rents individual "machines" for $0.02/hr to $2.40/hr depending on the machine resources[23]. More importantly, since the leasing is "per-hour" and because machines can be "turned on" quickly, software can be designed to use resources as needed.

---

[21] Despite the rather arcane history of peering arrangements, some access network providers, such as Comcast (http://www.comcast.com/peering/) have clearly articulated rules for how peering relationships are established.

[22] For example, in 2008, Sprint and Cogent networks "de-peered" their networks, causing service disruptions between Sprint and Cogent customers. See http://gigaom.com/2008/10/30/cogent-sprint-un-peer-may-cause-web-slowdown for an analysis and details.

[23] The listed prices are for machine instances, but any practical use of the service requires network bandwidth and storage, which are priced separately -- see http://aws.amazon.com/ec2/#pricing for complete current pricing.

Cloud computing has accelerated the deconstruction of monolithic software systems into components into a "service oriented architecture" that can be used as a service. Examples include Twilio.com, which integrates the legacy telephone network and provides voice-guided phone services. Such services, coupled with the ability to rapid deploy systems using cloud computing, allow developers to innovate in a select part of the software systems. But all these components – CDNs, cloud computing, software as a service systems – are rapidly becoming integral to the way that applications and services are deployed on the Internet. How will they be affected by regulation?

The Risks of Regulation In The Internet Ecosystem

There are several risks to the proposed network neutrality rules. These concerns include whether "neutral" networks even exist or are beneficial, the uncertainty concerning how services and applications should be treated, the risks of mandating monitoring for legal content and innovation in network management. We then address a general concern about the ability or wisdom of applying regulation in an era of fast-paced technology development by examining a particular Internet application regulated by the FCC.

Insensible Neutrality

Proponents of network neutrality legislation assume that people could agree on what a "neutral" network is and that any management other than existing prioritization methods will break applications. Is it possible for consumers to spot a "non-neutral" network? If neutrality cannot be measured or sensed, it's difficult to know when it is being violated or if it is even important. In our earlier work (Grunwald 2007), we detailed how the lack of clearly stated service level agreements for residential service and the multi-party nature of the Internet make it difficult to know what is affecting performance and who is responsible. Studies by networking researchers in (Akella 2003) and also more recently in (Dischinger 2007) have shown through careful measurement that the major performance limitations (latency, bandwidth, jitter) faced by most broadband users occur because of the technologies used in "last mile" access network – the connection to an individual house. At the same time, a study conducted in 2009 of Internet users in the US and Europe (Maier 2009) showed that user's *home networks*, and in particular the use of "WiFi" wireless networks, imposes more latency and variability than the access network itself. These measurement studies were conducted so broadly (across multiple ISP's in multiple countries) that they indicate that latency limitations and variability exist in most access networks. These limitations are caused by pressing existing infrastructure (cable and phone lines) into service for purposes they were never intended to serve, rather than by anti-competitive actions.

Because the Internet is composed of many pieces made by different parties, it's difficult to understand what causes specific problems. This is true even for experts – in a network measurement study, members of our research group initially reported that many types of network sessions were being blocked; upon further analysis (and much embarrassment), we had to retract that report because the problems were caused by a home networking

router[24]. This action only occurred when the home router was overloaded, but if the cause was not immediately clear to networking researchers, it's unlikely that an average consumer could identify similar problems.

As is clear by the success of existing applications, Internet protocol and application designers understand that minor fluctuations in latency and bandwidth go with the territory of the current Internet. Applications and various parts of the broader "Internet Architecture" are designed to accommodate those variations; there's good reason to believe that the design principles used in existing applications could overcome "subtle preferential treatment" just as they overcome the highly variable best-effort characteristics of the Internet. For example, video distribution systems came to rely on "faster than real time" downloads to successfully deliver video on the existing Internet (Odlyzko 2008). Despite the broad success of VOIP companies such as Vonage, Skype and the like, highly interactive applications (voice or video communication and interactive gaming) are usually thought to be sensitive to latency. However, comments submitted to the FCC by interactive game developers (Scherlis 2010) indicate that the current Internet is suitable for those applications.

All of this indicates that improving the speed of Internet access rather than fixing current network designs into law better serves consumers.

Fostering a Competitive Ecosystem

The proposed FCC rules affect only one part of the network, but performance and the user experience are affected by many parts of the network. Both content distribution and cloud computing resources are distributed globally and interconnected by private IP networks; since these are not "public networks", these facilities are free to prioritize traffic for payment without violating the proposed network neutrality rules. Singling out a single part of the Internet for regulation doesn't seem to insure the goal of competitive networks that respond to consumer needs.

There is continued vertical integration of the Internet market wherein "access network" providers also become CDNs or application companies (like Google) or retailers (like Amazon) become cloud computing providers. It's unclear how proposed regulations that distinguish between "public" and "private" networks will apply as those network companies recombine and change form. This either requires greater clarity of when the proposed network neutrality rules apply or, better yet, a "wait and see attitude" with action taken when anti-competitive harms actually occur.

Regulating Legal Content

The proposed neutrality rules focus on *lawful* content, and there have been both calls and proposals for applying "deep packet inspection" to assist in enforcing intellectual property ownership[25]. These efforts pose considerable costs and significant risks, both of

---

[24] See http://www.dslreports.com/shownews/Comcast-Now-Forging-Packets-For-All-TCP-Traffic-93388 for the embarrassing details.

[25] AT&T has stated that it will filter Internet content for such purposes. See "Has AT&T Lost Its Mind: A Baffling Proposal To Filter The Internet" by Tim Wu, Slate Magazine, Jan. 2008, available at http://www.slate.com/id/2182152/. Similar statements have recently been made Comcast CEO Brian Roberts;

misidentifying legal content as illegal and of failing to identify illegal content. Researchers have shown that anyone (including inanimate objects) can be implicated in file sharing (Piatek 2008). Existing file sharing systems are far from "stealthy" and are easy to monitor. Illegal file sharing is already hidden using "anonymity overlays" (McCoy 2008) and simple protocol extensions make it much more difficult to decidedly identify illegal file sharing activity (Bauer 2008).

At the same time, the rapid commoditization of co-location services, cloud computing and content distribution networks are also affecting illegal content. Not only can new companies be launched quickly, less legal Internet services are also possible. One of the many reasons that "peer-to-peer" (P2P) applications are popular is because they allow people to use their own infrastructure for file sharing. With the emergence of inexpensive cloud computing and other leased computing services, there has been a surge in the amount of Internet traffic for "hosted file services" at the expense of P2P services[26], making it easier for file-sharing to use those high performance systems rather than rely on the low-bandwidth uplinks common to the asymmetric network architectures used for access.

The rapid change in infrastructure that drives much of the Internet ecosystem illustrates the challenge to monitoring unlawful content. In two short years, "bandwidth intensive" applications such as video and file sharing have moved to systems using the same protocols and service providers as "legitimate" services. Because those systems use encryption, any mandated monitoring of such traffic will be both expensive and error prone. Stopping illegal content by monitoring traffic requires that *all* traffic be monitored and the costs to implement this will be borne by all users of the Internet. Pushing this requirement on all network providers imposes a significant cost to benefit a different industry.

Curtailing Innovation in Network Management

The propose neutrality rules distinguish between "managed" and "public" services, but the discussion about what constitutes managed services are relatively *ad hoc* and clearly capture the *status quo* rather than what is possible. An existing example would be having distinct network service for latency-sensitive traffic, such as voice. Some existing "competition friendly" networks (Moerman 2005) use a managed network exclusively for one of many possible voice services and relegate "best effort" and streaming video services to other networks all carried on the same fiber; similar capabilities are present to varying degrees in almost all other access networks. Commercial Ethernet uses 802.1Q (VLAN) and 802.1P (CoS) to provide such managed networks. New home network technologies such as Multimedia Over Coax Alliance (MOCA) and HomePlug[27] are

---

see "Comcast Set To Enter Copyright Wars" by Kenneth Corbin, Datamation, Jan. 27th 2010, available at http://itmanagement.earthweb.com/cnews/article.php/3861096/Comcast-Set-to-Enter-Copyright-Wars.htm

[26] This trend has been reported in numerous venues, with one of the more detailed studies being (Labovitz 2009) as presented to the 2009 NANOG network operators meeting; the full report has yet to be published, but presentation materials show a dramatic increase in "hosted HTTP" services rather than the expected increase in P2P services.

[27] See the HomePlug alliance (http://www.homeplug.org/home) and the Multimedia Over Coax Alliance (http://www.mocalliance.org/)

rapidly being developed that allow different managed streams to be carried over the same physical cable.

What have been missing are standards to link the differing streams in access network media to similar capabilities in home networks. A generalized capability to have multiple streams of data for multiple classes of service simplifies the distinction between "managed" and the "public" Internet and would allow additional managed services (*e.g.* video-conferencing could extend current "triple play" networks) or service offerings that let consumers choose between multiple service qualities. Some of these mechanisms are being developed[28], but such innovation will be likely be halted if ambiguous regulation is in place.

Similarly, many existing access network technologies have impediments that limit performance; even seemingly high performance networks such as DOCSIS cable modems benefit from "management" mechanisms[29] to overcome such impediments (Martin 2005). Similarly, the existing congestion control algorithms used to balance the performance of one "flow" *vs.* another at all scales of the Internet is being re-examined by the technical community. Bauer, Clark & Lehr (Bauer 2009c) published a very readable history of congestion control. Internet connections "self-regulate" the bandwidth they use – without such self-regulation, TCP connections would only be limited by the ability of the sender. Those algorithms seek to balance congestion in the network with the ability of the receiver to accept packets. The original algorithms sought to allocate each "flow" a fair share of bandwidth. That decision is one reason why Peer-2-Peer applications exert more pressure on networks than *e.g.* simple host based streaming – p2p applications use many connections to download content, and each is striving for a "fair share" of the access network. There are on-going efforts to evolve network congestion control algorithms to include information from the network in an order to build a more responsive and efficient network; network neutrality legislation seemingly precludes such efforts. These efforts include both the access network and congestion control at routers in the "core" of the Internet.

Technology On Internet Time

The FCC orders affecting the AOL Instant Messaging system during the Time Warner & AOL merge provide a historical lesson about the risks and challenges of predicting the path of technology and the impact that regulation has on that path. Instant Messaging emerged in the mid-1990's as a popular communication system based on a long history of "computer chat" systems in place since the early '70s. Messaging or "talk" applications were initially used on local area networks where the communication latency was sufficiently low. Because "chat" programs allowed users to communicate over long

---

[28] For example, PacketCable uses "Reserved Services Domain" to handle managed services; standards are being developed to bridge that standard with the MOCA home networking standard to preserve different service classes.

[29] DOCSIS cable modem networks tend to have "bursty" uplink connections, and this causes TCP/IP throughput to be lower than what the downlink can support. This particular study examines the effectiveness of "TCP ACK compression" to see if it overcomes the problems in the physical access network. This mechanism monitors TCP/IP connection characteristics and delays specific uplink traffic at the cable modem to eliminate redundant acknowledgement messages. The basic mechanism has been studied in other domains, but is rarely applied.

distances in near-real-time, they became increasingly popular on systems run by companies such as CompuServe, Prodigy, AOL and others. As with much of the on-line content of those systems, chat systems were initially "walled gardens" that only served the members of those services. As the commercial Internet evolved and became popular in the mid-to-late '90's, there was a greater interest in having IM systems operate across multiple services.

Instant messaging is notable because it is one of the few Internet technologies to have been affected by FCC and FTC orders. This occurred during the merger between AOL and Time-Warner; Faulhaber (Faulhaber 2002) has an excellent analysis and history of the reasoning behind orders affecting AOL Instant Messaging (AIM). At the time, Lehman Brothers valued AIM as $5.8B during the merger in 2000. AIM had 130 million members or users[30] and appeared to have considerable market dominance over nascent IM alternatives such as Microsoft MSN Messenger.

Prior to the merger, AOL and Microsoft had engaged in the "IM wars" wherein AOL exploited a security flaw in the AIM software to block inter-operation with competing services, such as Messenger. Microsoft, and other IM companies, lobbied for open access to the AIM service as condition of merger. Faulhaber argues that this was one of the first times that network effects was used as an argument in regulatory oversight in the absence of specific harm. It was thought that if Time Warner were able to block other IM systems from access to their cable modem networks, AIM would have significant advantage. This was thought important because it was clear that as network speeds increased, IM systems would evolve into a series of services (video chat, file transfer) that would expand on the value of the existing systems. The "names and presence directory" (NPD) was seen as being a critical infrastructure for IM services that precluded interoperability with other services. AOL resisted efforts to publish clear protocol standards or allow interoperation between their NPD and other software, asserting concerns of "security" and "privacy" for its users[31].

The FCC conditions for the AOL and Time-Warner merger prohibited the use of new "advanced" videoconference extensions unless standardized server-to-server interoperability mechanisms were implemented. Today, AIM is one of many protocols; although AOL still has the largest number of users, IM has both diminished in importance and multiple competing protocols and systems emerged. Today, it would be fanciful to image that AIM adds $5.8B of value for AOL. What happened?

In a large part, the efforts of AOL to block use of their services spurred development of competing services – this was even apparent at the time the merger conditions were being debated[32]. In addition to the MSN Messenger system, several "open source" efforts were

---

30 See Louise Rosen, *Why IM matters* so much, UPSIDE, September 19, 2000. The original reference is difficult to find, but is included in an *Ex Parte* filing by Gerad Waldron (Covington & Burling law firm) to Michael Powell, re: "Applictions of America Online and Time Warner Inc., For Transfers of Control (CS Docket No. 00-30)"

31 It was noted at the time that AOL's concern about security and privacy was disingenuous given AOL's reliance on a "buffer overflow" attack to block competing services; that same attack could be used to compromise the customer's computer.

32 See Jim Hu, *AOL's Lead in Instant Messaging Arena Dwindles*, CNET NEWS.COM, Nov. 16, 2000 available at http://www.zdnetasia.com/news/communications/0,39044192,13030364,00.htm; this was referenced in footnote

developed to produce scalable messaging platforms with the most successful being Jabber, which produced the XMPP[33] protocol. These multiple implementations allowed companies to launch their own, private and customized IM services because the cost of deploying the technology was greatly reduced. People learned that adopting a new IM system wasn't hard. In part, the plurality of systems and the willing to adopt new IM systems accelerated the use of IM and messaging systems for business applications. One of the complications of using AIM for business purposes was that AIM was often blamed for security lapses and that businesses had poor controls over the identity, security, privacy and logging needed when applying AIM to business applications. In particular, the various financial scandals that precipitated the Sarbanes-Oxley Act of 2002 and other reporting and disclosure rules make it more important to keep accurate records and logs of communication between investors and financial advisors as well as between people in the investment community. This led several companies to stop using public IM networks[34] in favor of "in house" networks. Eventually, those IM systems used web browsers rather than require extra clients to be downloaded. The development of "Web 2.0" technologies such as AJAX changed the IM experience afforded by a browser interface to be equal to that of dedicated software. This allowed businesses to maintain control over "customer chat" and tie the system with customer names or account numbers.

The pace of technology adoption and the peculiar needs of companies seeking to employ IM systems means that although AOL's system is still the largest IM system, there was no strangle hold on innovation or capabilities. The pace of this innovation was addressed in the FCC merger memorandum:

> "Finally, it might be thought that in the rapidly changing technology of the Internet, even network effects and AOL's present position in the market would not prevent successful entry by IM providers other than AOL, that a new breakthrough technology might become available and would be superior enough to AOL's service to overcome the network effects flowing from its NPD, and cause users to shift *en masse* away from AOL. …..We see no evidence at this time, however, of such a new breakthrough technology strong enough to overtake AOL's NPD."

With the benefit of hindsight, we see that within 2-4 years after the merger orders were written, events led to rich IM competition. Customers did not shift *en masse* away from AOL because they didn't need to – they simply used other technologies in concert with AIM.

Hindsight certainly helps in seeing trends, but some trends are only apparent when other technologies arise. One of the FCC's concerns with the AOL and Time-Warner merger

---

435 of the FCC memorandum.

[33] Extensible Messaging and Presence Protocol; see http://xmpp.org for more information.

[34] See for example Thomas Hoffman, *"Sarbanes-Oxley trumps IM at some firms Concerns about security, archiving prompt companies to unplug instant messaging systems"*, Computer World, August 8th, 2005. Available as http://www.computerworld.com/s/article/103752/Sarbanes_Oxley_trumps_IM_at_some_firms

was that it might lead to a new dominant signaling and communication system by the introduction of new services over AIM. We've argued that this didn't come to pass because alternate services became available (and were easy to adopt), that mechanisms existed to work around restrictions and that open standards reduced the barrier for entry. The rapid evolution of technology was in contrast to most of the history of telecommunications and this rapid evolution made it difficult to estimate the impact of regulation.

Maintaining a Vibrant Internet Ecosystem

Technology on the Internet moves both slower and quicker than most technology overseen by traditional regulation. VOIP technologies were in place almost a decade before they became widely adopted. Promising technologies such as AIM arose, peaked and then diminished in value dramatically within that same period of time. The technology for one application was largely a substitute for the other[35] but that wasn't clear at the time.

There are approaches other than or in addition to regulation that would help maintain a vibrant Internet. These include: having clear standards and methods for measuring what is actually happening in the Internet as well as methods for reporting or disseminating policy to consumers; use existing agencies and policies; encourage innovation and competition for access networks; and, developing "best practices" that can be clearly understood by network operators, regulators and consumers.

Measure and Report

Clearly identifying problems in the Internet and apportioning blame is very difficult. Consumers on access networks typically want answers to three questions: can I access a specific service; is the latency or quality of that service acceptable; and, is there a bandwidth problem for a specific service.

Consumers often jump to conclusions when a service or site is blocked or unavailable. Services may be blocked by an ISP – or, the service may actually be down. Alternatively, parts of the Internet protocols not under control of the ISP (such as DNS) may misdirect traffic. In extreme cases, events half way around the world may block services[36].

---

[35] The XMPP protocol used by Jabber and Google Talk has been extended as "Jingle" by Google to enable voice calls. Similarly, voice systems such as Skype added support for basic and "enhanced" IM services.

[36] One example occurred in February 2008 when the government of Pakistan ordered access to YouTube to be blocked within Pakistan. The network operators for Pakistan Telecom implemented that order by issuing a "black hole route". This is a method whereby a network router advertizes that it has an efficient route to the designated host but then actually discards that traffic. That "black hole route" was then published to other ISPs, causing a large part of the world to think that Pakistan had a very good connection to YouTube; this caused broad outages for YouTube. See http://asert.arbornetworks.com/2008/02/internet-routing-insecuritypakistan-nukes-youtube/ for more details.

The debate concerning network neutrality has prompted the development of several measurement tools to determine if application blocking or data modification is occurring. Examples include the "Switzerland" tool (developed by the Electronic Frontier Foundation[37]), the "Glasnost" tool developed by The Max Planck Institute (Dischinger 2010) and the "Measurement Lab" consortium that supports both education and analysis tools[38]. These tools either detect specific problems (e.g. BitTorrent blocking) or identify factors that may delay communication. They are first steps in helping consumers identify what may be wrong and to assist in network monitoring. However, they are still primitive and require considerable sophistication to deploy and interpret.

It would be better for ISPs to be transparent about their network management policies and network conditions. Many ISPs block services that appear to arise from "malware"; sometimes those services are actual but uncommon services. For example, (Scherlis 2010) notes that game developers often need to contact ISPs to remove blocked services that are mis-identified as malware. As the same time, consumers are typically unaware when one of their home computers or devices is launching network attacks on others.

What's missing is a mechanism or protocol for communicating current management and policy information to consumers. Developing standards or protocols for informing customers about "suspicious" traffic would remove much of the confusion when an application stops working. There are existing protocols (such as SNMP and RMON) designed to communicate network performance, but these protocols are designed for network management rather than consumer enlightenment – they provide too much detail for consumers and provide no insight into what steps can be taken to correct problems. Through efforts such as the P4P consortium[39], ISPs have found that it's possible to work with applications to reduce bandwidth demands and costs. Similar tools for communicating with consumers would likely improve customer service and help reduce network security problems. Efforts to inform consumers about broadband capabilities would allow broadband providers to compete based on those different services without consumers complaining about hidden differences. The British regulator, Ofcom, has established a voluntary "Code of Conduct" for ISPs that communicates much of this information to consumers prior to sale and during service[40].

Maintain Competitive Applications, Content and Services

Content distribution and cloud computing services dramatically reduce the infrastructure cost for computing and web applications, allowing non-commercial groups to rapidly scale their efforts. Software innovations and business models that can exploit these new

---

[37] See the EFF website, "Swizterland Network Testing Tool", available at http://www.eff.org/testyourisp/switzerland

[38] Measurement labs (http://www.measurementlab.net/) arose from an effort by a number of companies, university faculty and Internet researchers to determine technical approaches to measuring the network access characteristics.

[39] P4P is a reporting method that allows P2P software to learn the "topology" of ISPs, allowing the P2P software to avoid expensive or congested links. More information is available at http://www.openp4p.net/

[40] The Code for Contact is available at http://www.ofcom.org.uk/telecoms/ioi/copbb/copbb/

platforms are enabling even more rapid innovation. Vertical integration in these markets may or may not lead to anti-competitive behavior; however, these technologies are so new that it's not clear whether they will remain in their current form or if concerns about fair competition will last longer than the technology itself. Rather than enacting preemptive and broad rules to regulate these hybrid "private/public" networks, waiting for harms to emerge, coupled with the possibility of anti-trust laws and enforcement from the FTC and the Department of Justice, will foster more innovation.

Predicting the future of technology is difficult as suggested by the analysis of the likely outcome of the competition surrounding AOL Instant Messaging. That regulation was eclipsed by the reality of rapid technology development, external technology, and changes in business practice and usage patterns. Although there is certain to be consolidation in the "cloud computing" ecosystem, it's remains to be seen if the consolidation will foster anti-competitive behaviors.

Maintain Competitive Networks With Transparency & Clarity

Business networks (primarily Ethernet) have many mechanisms to improve flexibility, control performance or diagnose problems. Consumer access network technology is only beginning to see similar development and there is a real risk that regulation will curtail investment or development of those technologies. At the same time, certain services benefit from separation from general best effort traffic – this is why many businesses use different "virtual private networks" to separate different kinds of traffic. As home users expand the range of services they use, consumers may be better served by technologies that enable multiple network services, each with different qualities.

Likewise, innovations in congestion control will continue and can be implemented in many parts of the networks. Researchers are exploring the tension between enforcing congestion control at the end-pointers (where it may take years to upgrade or replace all the software) *vs*. upgrading specific routers or other parts of the network. Precluding implementation at the access network will simply increase the costs of network management. Rather than exclude specific mechanisms such as congestion control, regulation should be used to foster goals such as competition.

Keep Ahead of the Technology

The Internet is complex, encompassing both traditional communication services as well as computer systems, novel services and rapidly evolving technology. Developing an on-going process for discussing and analyzing the interplay between the different technologies is critical. There are specific actions that can foster more thoughtful review. One would be to have an organization that provides independent and informed council to policy makers about the Internet ecosystem as a whole; this is a difficult charge because some emerging trends aren't apparent until they are established businesses. The other action is to counter specific concerns that have been indicated by prior regulators and develop standards or tools to mitigate those concerns.

There are many bodies that examine and discuss how Internet technology should be developed; other groups discuss business practices and yet others research new techniques or services. It's equally important to have a continued and informed discussion of how technology, business and new services affect future policy so that

policy makers can stay ahead of the technology. It's useful to guide technology before it is widely deployed because that lessens the cost of regulation.

One such example is the "network effects" of systems such as instant messaging systems, or the "stickiness" of specific Email addresses. As an example, although there have been petitions for "email portability", there has been little serious study of the concept. However, "identity" on the Internet is one of the key features that makes "network effects" important. Although AIM wasn't the only messaging tool available, moving to another system entailed rediscovering the identity of your friends. Looking into the present where Instant Messaging has been replaced with Social Networking, the same issues that were raised about AIM "stickiness" may be raised about Facebook or MySpace. Here, the technical community is moving faster than the regularity world -- there have long been Internet standards (*e.g.* DNS) that allow "machine portability" and there are developing standards, such as OpenID, for "people portability". Such identity systems could have significant impact when widely adopted, but it's also important to understand and clarify how such systems will interact with regulation.

Regulation should be a process, not a product

We've argued that regulation or legislation that simply affects control of the access work policies while ignoring the impact of the rest of the Internet ecosystem is a disservice to consumers. At the same time, regulation or legislation that affects the *entire* Internet is over-reaching and also not needed.

To date, most of the network neutrality discussion has been heavily influenced by existing telecommunications regulation – this is natural since most regulation seeks to model new systems after old. This has led regulators to focus on "bits in flight" – *e.g.* the regulation of access networks – while largely ignoring the "bits at rest" (content distribution networks) that make up much of the Internet. That distinction between basic and information services is rapidly being challenged by the development of an integrated Internet ecosystem. Focusing on "bits in flight" also impacts the ability of regulators (or even tech pundits) to predict the evolution of services. We highlighted the example of AOL Instant Messenger in this paper, arguing the comparison between AIM and the existing communications systems missed the rapidity with which new and competing systems could be *deployed* using the existing infrastructure. Standardization and open software and protocols also meant that the cost of *developing* a new system was radically reduced compared to existing telecommunications systems. The rapid evolution of the Internet makes it difficult to insure that regulation is still meaningful by the time it is developed.

True network neutrality is about competition and innovation and any such discussion must involve the full Internet ecosystem. It's clear that narrowly defined rules affecting one part of that ecosystem are not the best solution to maintaining a competitive and responsive Internet. Existing legislation (primarily anti-trust laws in the case of browsers and the threat of similar laws in advertising based search) are being applied and should be able to address future anti-competitive actions. At the same time, consumers would

benefit from competition, innovation and better information about the services available to them.

Bibliography

(Aalberts 1999) – Robert J. Aalberts and Marianne M. Jennings, "*The Ethics of Slotting: Is this Bribery, Facilitation Marketing or Just Plain Competition*", Journal of Business Ethics, Vol. 20 No. 3, July 1999

(Akella 2003) - Akella, Seshan and Shaikah, "*An empirical evaluation of wide-area Internet bottlenecks"*, Proceedings of the 2003 Internet Measurement Conference, Oct. 2003.

(Bauer 2008) -- Kevin Bauer, Damon McCoy, Douglas Sicker, and Dirk Grunwald, " "*BitBlender: Light-Weight Anonymity for BitTorrent"*, In Proceedings of the ACM Workshop on Applications of Private and Anonymous Communications Istanbul, Turkey (September 2008)

(Bauer 2009) - Kevin Bauer, Dirk Grunwald and Douglas Sicker, "The Challenges of Stopping Illegal Peer-to-Peer File Sharing", in Proceedings of National Cable and Telecommunications Association Technical Conference (April 2009).

(Bauer 2009b) -Kevin Bauer, Damon McCoy, Dirk Grunwald, and Douglas Sicker, "*BitStalker: Accurately and Efficiently Monitoring BitTorrent Traffic"*, In Proceedings of the 1st IEEE International Workshop on Information Forensics and Security London, United Kingdom (December 2009)

(Bauer 2009c) Steven Bauer, David Clark, William Lehr, "*The Evolution of Internet Congestion*", Proceedings 2009 The 37th Research Conference On Communication, Information and Internet Policy (TPRC), Oct. 2009

(Dischinger 2007) - Dischinger, Gummadi, Haeberlen and Saroiu, "*Characterizing Residential Broadband Networks"*, ACM Internet Measurement Conference, 2007.

(Dischinger 2010) - Marcel Dischinger, Massimiliano Marcon, Saikat Guha, Krishna P. Gummadi, Ratul Mahajan, and Stefan Saroiu , "*Glasnost: Enabling End Users to Detect Traffic Differentiation*," To Appear in the Proceedings of the 7th Usenix Symposium on Networked Systems Design and Implementation (NSDI), San Jose, CA, April 2010.

(Faulhaber 2002) – Gerry Faulhaber, "*Network Effects and Merger Analysis: Instant Messaging and the AOL-Time Warner Case,"* in Telecommunications Policy, Jun/Jul 2002, 26, 311-333.

(FTC 2003) FTC Staff Study, "Slotting Allowances in the Retail Grocery Industry: Selected Case Studies in Five Produce Categories", November 2003.

(Gundlack 2005) - Gregory T. Gundlack, "Statement on Slotting Fees- Fees Charged By Grocery Retailers for Shelf Space: Are They Stifling Competition?", before the California State Senate Standing Committee on Business, Professions and Economic Development, Feb. 9, 2005.

(Grunwald 2007) - Dirk Grunwald and Douglas Sicker, "*Measuring the Network - Service Level Agreements, Service Level Monitoring, Network Architecture and Network Neutrality*", Intl. Journal on Telecommunications, Feb 2007.

(Jiang 2009) - Jiang, W., Zhang-Shen, R., Rexford, J., and Chiang, M. 2009. "*Cooperative content distribution and traffic engineering in an ISP network*". In Proceedings of the Eleventh international Joint Conference on Measurement and Modeling of Computer Systems (Seattle, WA, USA, June 15 - 19, 2009). SIGMETRICS '09. ACM, New York, NY, 239-250.

(Kare 2004) Rohit Kare, Doug Cutting, Kragen Sitaker, Adam Rifkin, *"Nutch: A Flexible and Scalable Open-Source Web Search Engine"*, CommerceNet Labs Technical Report 04-04, November, 2004.

(Labovitz 2009) - C. Labovitz, S. Lekel-Johnson, D. McPherson, J. Oberheide, F. Jahanian, & M. Karir, "*ATLAS Internet Observatory: 2009 Annual Report*", available at http://www.nanog.org/meetings/nanog47/presentations/Monday/Labovitz_ObserveReport_N47_Mon.pdf.

(Moerman 2005) - K. Moerman, J. Fishburn, M. Lasserre, and D. Ginsburg, "Utah's Utopia: an Ethernet-based mpls/vpls triple play deployment," Communications Magazine, IEEE, vol. 43, no. 11, pp. 142–150, 2005.

(Maier 2009) - Maier, Feldman, Paxson and Allman, "*On Dominant Characteristics of Residential Broadband Internet Traffic",* in Proc. 2009 Internet Measurement Conference, Nov. 2009

(Martin 2005) J. Martin, "*The Impact Of the DOCSIS 1.1/2.0 MAC Protocol on TCP*", Proceedings of the Consumer Communications and Networking Conference (CCNC 2005), Jan 2005.

(Odlyzko 2008) – "*The delusions of net neutrality*", in Telecommunications Policy Research Conference, 2008.

(Odlyzko 2009) - A. Odlyzko, "*Network neutrality, search neutrality, and the never-ending conflict between efficiency and fairness in markets*," in *Review of Network Economics*, vol. 8, no. 1, March 2009, pp. 40-60

(McCoy 2008) - Damon McCoy, Kevin Bauer, Dirk Grunwald, Tadayoshi Kohno, and Douglas Sicker, "*Shining Light in Dark Places: Understanding the Tor Network*", In Proceedings of the 8th Privacy Enhancing Technologies Symposium, Leuven, Belgium (July 2008)

(Pasquale 2008) Frank A. Pasquale and Oren Bracha, "*Federal Search Commision? Access, Fairness and Accountability in the Law of Search"*, Cornell Law Review, September 2008, available as http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1002453

(Piatek 2008) -- Michael Piatek, Tadayoshi Kohno, Arvind Krishnamurthy, "*Challenges and Directions for Monitoring P2P File Sharing Networks –or– Why My Printer Received a DMCA Takedown Notice*". Proceedings of the HotSec 2008 Hot Topics in Security Workshop

(Saltzer 84) -- Saltzer, J. H., Reed, D. P., and Clark, D. D. 1984. End-to-end arguments in system design. *ACM Trans. Comput. Syst.* 2, 4 (Nov. 1984), 277-288.

(Scherlis 2010) - Daniel R. Scherlis, "Re: notice of Ex Parte Communication GN Docket No. 90-191; WC Docket No. 07-52", submitted to Ms. Marlene H. Dortch, Federal Communication Commission, January 2010, available at http://fjallfoss.fcc.gov/ecfs/document/view?id=7020380907

(Sullivan 2009) – Mark Sullivan, "A Day In the Life of 3G", PC World Magazine, June, 2009, available as http://www.pcworld.com/article/167391/a_day_in_the_life_of_3g.html

(Yoo 2010) Chistopher Yoo, "*Innovations In The Internet's Architecture That Challenge The Status Quo*", Journal of Telecommunications and High Technology Law, Jan. 2010

Dirk Grunwald is Professor in the Department of Computer Science, the Department of Electrical and Computer Engineering and the Interdisciplinary Telecommunications Program at the University of Colorado He received his Ph.D. from the Department of Computer Science at the University of Illinois in 1989. Alongside his many gifted students, he has also studied specialties include computer networking, wireless networking, privacy in wireless networks, mechanisms to enforce and counter anonymity in the Internet, advanced techniques to compile languages for emerging hardware devices, computer architecture and design, storage systems and tools for highly parallel computers. He currently has research funding from the National Science Foundation, the Defense Advanced Research Projects Association and Intel. He has graduated 19 Ph.D. students who work at a variety of universities, industrial research labs and various startups. He has served on the technical advisor board of Copan Systems and has consulted on topics ranging from storage systems to mobile computing and high performance systems and well as intellectual property evaluations and expert advice for the Department of Justice.

His current research involves the design and evaluation of *Software Defined Radios* and *Cognitive Radios*, or wireless systems that adapt to their environment and coordinate with one another to achieve high performance. His research group has build advanced software defined radios, deployed adaptive “WiFi” mesh networks using advanced “phase array antennas” and are in the process of deploying an advanced “programmable” WiMAX network.

Dr. Grunwald is thankful for the advice from other members of the ITP faculty while preparing this paper. The author gratefully acknowledges financial support for this work provided by the College of Engineering Faculty Fellowship program and Verizon.