

Network neutrality, the even-handed treatment by intermediate carriers of data streams as they zip through the Internet from one endpoint to another, is a core assumption of current broadband users' online interactions, but it is now threatened by tools that allow Internet Service Providers (ISPs) to discriminate between parties and modes of communication. Network neutrality is a good worth protecting for a variety of reasons, but it is unclear what method is best to assert its continued application on the Internet. Code, economics, and legislation are the three apparent means to this end. Legislation is the least undesirable of these approaches, offering the greatest potential payoffs and the least damage in the worst case.

The broadband users mentioned above include individuals, traditional corporations, and online businesses, each relying on Internet access for different purposes. Individuals take for granted the unfettered performance of their connections, whether using voice over IP, downloading rich media, shopping online, or simply visiting websites to find information on network protocols, legal policies, or movie showings. They also believe that the First Amendment-protected expression of controversial ideas is safe, on the web, from corporate censorship. Nearly all businesses, from Internet powerhouses like Google to energy companies and stock traders, rely on transparent broadband support of their private communications and economic transactions. All of these uses of the Internet are *applications*, the proliferation of which is fundamental to the social and economic benefits that Internet access provides. The primary purpose of network neutrality assurances “is to give users the right to use non-harmful... applications, and give innovators the corresponding freedom to supply them.”¹ Online businesses, themselves users of the Internet and important customers to ISPs, deal in a market for

1 Wu, “Network Neutrality, Broadband Discrimination”, p. 2.

these very Internet applications, and found their business model upon unrestricted access to consumers. Free competition in this market, which today contributes greatly to the American economy, is absolutely dependent upon network neutrality. We see, then, that the application market is intrinsically linked to the *infrastructure* market of ISPs; if the infrastructure does not provide a neutral network on which application providers may fairly compete for customers, the applications market becomes distorted and cannot function properly.²

The network neutrality assumption has been a safe and stable one in the past for all these users, but ISPs now have technological tools available to them to enable “deep packet inspection”, allowing them to examine the headers of every data packet that they handle.³ This in turn opens up the possibility for ISPs to treat each stream of packets differently depending on its source, destination, or content, and thereby increase their revenues through value pricing of large applications companies' access to their online customers or outright competitor lock-out. While they have not yet exercised their ability to practice such policies on a grand scale, these economic motivations guarantee that large ISPs will, in the near future, violate the principle of network neutrality in many different arenas. While the debate over the necessity of network neutrality on the Internet is vigorous and extensive, we leave its thorough discussion to other sources.⁴ Instead, this paper operates under the framing assumption that neutrality is a vital feature of the Internet that must be protected, based upon its importance to all the parties outlined above. What, then, is the best method by which to protect it?

Three apparent approaches to the assertion of network neutrality are code, economics, and legislation. None of these systems is absolutely reliable and each has its own benefits and drawbacks, which will be explored in detail below. The assumption that code-based solutions

2 The “infrastructure/applications” terminology for this distinction is due to Wu and Lessig, *FCC Ex Parte Letter*.

3 Chester, “The End of the Internet?”, in *The Nation*, 13 February 2006 online edition.

4 Many arguments on this topic are outlined in Stern, “The Coming Tug of War Over the Internet”, in *The Washington Post*, 22 January 2006 edition.

exist is founded upon a misunderstanding of the distribution of control among the parties that decide how packets are handled, and may be quickly dismissed. While undue legislative regulation of free markets is to be avoided as much as possible, the Internet infrastructure market has idiosyncratic properties that prevent capitalist economics from behaving properly even in the absence of outside regulation. Furthermore, economic self-regulation at the infrastructure level can make no assurances about the protection of network neutrality, which, as explained above, is vital to the healthy functioning of the Internet applications market. The free functionings of these two markets are mutually exclusive, and legislation sacrifices the former in the interest of the latter. This is an unfortunate but necessary tradeoff and is the best solution to protect the principle of network neutrality and the social and economic freedoms that rely upon it.

Lawrence Lessig makes clear early on in his book Code that, on the Internet, “code is law”.⁵ While social and economic conflicts may play out within networked systems, they do so always over the substrate of the code that underlies that network. Code defines the constraints of the system and therefore the boundaries of conflicts that may occur within it and the solutions to those conflicts. Proper design of this code may allow for mutually beneficial resolution of conflicts. Why not, then, design code that guarantees network neutrality?

The trouble with a code-based approach is that there is no position of power from which parties interested in supporting network neutrality—individual users or businesses—may assert its practice through code. ISPs assert their control over packet routing through code running on their own hardware; in the tussle over network neutrality, it is *their* code that is law. As long as all end users must pass their data through an ISP's router to get it out across the Internet, ISPs exert ultimate control over this data. They may choose to slow its transmission on either end of

5 Lessig, Code, p. 6.

the transaction, or to block it outright; once this choice is made by the ISP, there is no way to undo it at another point in the transaction. No other participants in the data exchange process—neither the sender, the recipient, nor the operators of the central Internet hubs—may affect the ISP's behavior, and therefore the fate of the packets, through code.

Failing to find a solution in code, we now turn to economics. It is possible that economic factors in the infrastructure market might guarantee network neutrality in the long term. Capitalist economics asserts that an unregulated, competitive market will ultimately evolve to meet its customers' desires, and clearly neutrality is a desirable property of the service offered to end users by ISPs. The Internet infrastructure business, however, has intrinsic properties that eliminate economic motivations for behavior that benefits customers.

First, while violation of network neutrality is an undesirable trait in Internet service, customers are likely to misidentify the sources of non-neutral network behavior and respond with economic actions that will not affect ISPs. Customers are accustomed to a transparent Internet: “packets go in, and they come out, and that is all that happens in the network.”⁶ The infrastructure level is therefore invisible to most customers; they often do not even consider the existence of the network that mediates their communication with websites and other users. If an ISP sold preferential treatment of packets, for example, to Yahoo! video, users would find that Yahoo! offered a superior service compared to its competitors and would respond by economically “punishing” those competitors. These users would ordinarily be offended by this sort of manipulation of their economic decisions, but their behavior, by paying those who pay the ISPs, would reinforce such policies.

Second, there are extremely high barriers to entry, and therefore little competition, in the ISP market. In addition to the standard difficulties involved in establishing a business, a startup

6 Clark, et al., *Tussle in Cyberspace: Defining Tomorrow's Internet*, p. 8.

ISP must lay down its own transmission lines or negotiate the use of existing lines owned by another company. In order for data from its customers to reach the rest of the Internet, our hypothetical startup must also negotiate data transmission arrangements with at least one other ISP. The contracts resulting from these negotiations necessarily involve huge amounts of money, and there is no economic recourse for small ISPs who cannot pay the established players. Thus there is a stable oligarchy of powerful entities in the ISP market, the existence of which severely distorts its economic behavior.

The small population of powerful ISPs makes this market ripe for cartel control. In any system with few participants, a strategy of cooperation may be preferable over competition, as it offers unbounded mutual benefit and stronger defense against potential competitors than any individual could muster. This principle holds as true in economics and politics as in evolutionary biology. In economic markets, however, this sort of cooperation amounts to collusion, defending its parties against damage from all sides, whether from competitors or from customers. Allowing this behavior to dominate a market is counter to the laissez-faire capitalism on which our economy is founded, and therefore our government has trade regulations, such as anti-trust laws, to punish its practitioners. In the absence of other strong regulations, however, anti-trust law has a history of spectacularly ineffective application in recent years. One high-profile example is the slow and ultimately meaningless case against Microsoft that was pursued in the late 1990s.⁷ The FTC also failed to serve American consumers through the application of anti-trust law in their settlement with CD publishers over their “minimum advertised price” scheme, which allowed the publishers to take out an interest-free \$143 million loan from the American people and still make \$337 million more over the course of a decade than a fair market would have allowed.⁸

7 *Litigating States' Response to Settlement in United States v. Microsoft Corporation*, http://www.usdoj.gov/atr/cases/ms_tuncom/major/mtc-00030607.htm.

8 FTC Press Release, *Record Companies Settle FTC Charges of Restraining Competition in CD Music Market*, 10 May 2000; Reuters News Service, *Five Music Concerns to Pay \$143.1 Million in Price-Fixing Case*, in *The Wall Street Journal*, 1 October 2002 edition.

Even disregarding the argument above that consumers are unlikely to exert economic pressure on ISPs who violate network neutrality, the absence of disincentives to cartel control in the infrastructure market guarantees a long turnaround time on readjustments to consumer demands. Customers of a cartel have no alternative sources of the service that the cartel provides and therefore no economic means of influencing the properties of the received service. It would take a long time for a startup ISP that provided neutral network access to gain a foothold in the infrastructure market sufficient to provide meaningful consumer choice, especially considering the tools available to an infrastructure cartel to lock out new competition.

If we optimistically assume that customers of an ISP cartel would eventually realize the good of network neutrality, it is nearly guaranteed that a company would eventually force its way into the ISP market and restore balance and competition. The best-case scenario for network neutrality under economic self-regulation of the infrastructure market, then, is a long and steady erosion of consumer rights and the principle of network neutrality, followed by a slow recovery. Is it really necessary to wait for things to get worse before they get better?

Our final option, and the one that I propose offers the greatest promise for protecting network neutrality, is legislation. Free competition within the applications market requires regulation of the infrastructure market; similarly, free competition within the infrastructure market implies restriction of the applications market. Freedom in the two markets is therefore mutually exclusive, and any legislative intervention must choose to protect one at the expense of another. There are many similar regulations in United States trade law, which sacrifice one form of economic freedom to encourage another; this characterizes anti-trust law, trade union law, and international tariffs. The tradeoff in the case of legislating network neutrality is acceptable because, as is argued in the case for the principle of network neutrality itself, the applications market holds more value for the American economy than the infrastructure market does. As long

as the infrastructure would continue to exist, improve, and provide high-quality service under regulatory enforcement of network neutrality, both the economy and all the Internet users introduced at the beginning of this paper would benefit. While this condition is clearly a strong and contentious one, there are reasons to believe it would hold.

Our ideal piece of legislation would explicitly require all ISPs to act as common carriers, transmitting packets to their intended destinations with no unequal treatment. Many opponents of legislating network neutrality insist that there would be no incentives for innovation in the infrastructure market under such a regime. Continued growth of the applications market relies upon infrastructure innovations, as each new generation of data-intensive applications demands more and more bandwidth. However, ISPs have managed to be profitable, competitive, and innovative for the entire history of the Internet up until now without resorting to selling preferential packet treatment. ISPs have for years sold consumers, both residential and corporate, faster and more robust access to the Internet by providing multiple qualities of fringe infrastructure: dial-up, DSL, and all the traditional corporate “fat pipes” like T-3s are available at different prices to customers who are willing to pay for them. As competition within the applications market continues to drive innovation, demand for faster connections and therefore further physical development of the Internet infrastructure will only increase. ISPs may also add value to their services by providing superior customer or corporate support, or any number of packaged services. The infrastructure market would continue to be profitable and a healthy capitalistic system under network neutrality regulation.

One other strong argument against legislation of network neutrality is that it would be unenforceable. After all, if anti-trust law is so ineffective, why should we expect success with additional regulation? While this worst-case scenario is less damaging to the public than the abuses that could arise from relying on economic self-regulation in the infrastructure market, it

merits discussion. Anti-trust law is meant to protect almost every type of economic market that exists in the United States, as monopolies and collusion may infiltrate any market. The language of such laws must therefore be very general, and especially in high-profile cases, great effort and expense must be undertaken in its interpretation. Network neutrality legislation of the sort proposed above would instead be focused on regulating only one market, and in only one specific way. It would be another tool available to the government to assure the free functioning of the Internet applications market, and, more importantly, the most precise tool yet for this purpose. The infrastructure business is high-profile and affects hundreds of millions of consumers every day. Network neutrality violations would be unlikely to go unnoticed for long, and the specific legislation proposed above would allow for the swift and unambiguous legal correction of such violations.

Network neutrality is an essential property of the modern Internet, worth protecting to the best of our abilities. No purely code-based solution exists to assert neutrality, and the best outcome that economic self-regulation of ISPs can offer is a continued decline in service followed by an eventual recovery to the network behavior that we take for granted today. No legislative solution to a problem can be perfect, even in the absence of bureaucracy and corruption. In this case, however, legislation offers the most precise and minimally disruptive mechanism—one that, unlike other potential solutions, suffers from no absolute drawbacks—for the protection of network neutrality on the Internet.