Scenarios for the Network Neutrality Arms Race

William H. Lehr
Sharon E. Gillett
Massachusetts Institute of Technology

Marvin A. Sirbu
Jon M. Peha
Carnegie Mellon University

Presented at the
34th Research Conference on Communication, Information, and Internet Policy (TPRC)
September 29-October 1, 2006
Arlington, VA

-- August 31, 2006--

ABSTRACT

Several factors suggest that meaningful network neutrality rules will not be enshrined in near-term U.S. telecommunications policy. These include disagreements over the need for such rules as well as their definition, efficacy and enforceability. However, as van Schewick (2005)[1] has demonstrated in the context of the Internet, network providers may have economic incentives to discriminate in welfare-reducing ways; in addition, network operators may continue to possess market power, particularly with respect to a terminating monopoly.[2] On the other hand, the literature on two-sided markets,[3] the challenge of cost-recovery in the presence of significant fixed and sunk costs, and the changing nature of Internet traffic all provide efficiency-enhancing rationales for discriminatory pricing and traffic management. Thus, policy-makers face a daunting challenge: discriminatory behavior is likely to occur and distinguishing between good and bad discriminatory behavior is difficult.

Assuming that various forms of network-based discrimination are likely to occur, broadband end-users may employ a variety of technical and non-technical strategies to counteract its effects, which in turn, will likely elicit

[1] See van Schewick, Barbara (2005), "Towards an Economic Framework for Network Neutrality Regulation," Paper presented at the 33rd Research Conference on Communication, Information and Internet Policy (TPRC 2005), George Mason Law School, Arlington VA, September 2005, available at: http://web.si.umich.edu/tprc/papers/2005/483/van%20Schewick%20Network%20Neutrality%20TPRC%202005.pdf

[2] These problems arise even if the access provider is not dominant in other respects. See DeGraba, P., "Bill and Keep at the Central Office as the Efficient Interconnection Regime," U.S., FCC OPP Working Paper 33, (2000).

[3] See Rochet, J. and J. Tirole (2004), "Two-sided markets: An Overview," working paper, IDEI, available at: http://faculty.haas.berkeley.edu/hermalin/rochet_tirole.pdf

further responses from the network operators.[4] The goal of this paper is to characterize the resulting arms race by examining scenarios for how downstream end-users of broadband, sometimes in conjunction with upstream players (e.g. content providers), might respond to limit the potential harm from network-based discrimination. We identify three classes of end-user responses: (1) infrastructure-based bypass (e.g., municipal open access networks, mesh networking, or multi-homing); (2) technical and non-technical counter-measures (e.g., letter writing campaigns, end-to-end encryption, or onion routing); and (3) living with the differentiation (e.g., time-shifting and DVR buffering to use low-grade transport to view high-quality content).

Our analysis suggests several implications for policy-makers. First, even in the absence of network neutrality regulation, end-users (and upstream providers) have a range of technical and market-based strategies for responding to discrimination. Second, providers may find it difficult to maintain forms of discrimination that are associated with positive externalities, such as an expanded user base or less congested networks. Thus, a priori, the welfare implications of end-user responses are ambiguous. Third, the availability and effectiveness of responses to discrimination are likely to vary not only by geography but also by the level of skill and economic resources available to particular customers, raising potential equity issues. Moreover, the effectiveness of end-user strategies depends critically on the mode of behavior adopted by the operator. We conclude that end-user responses are not sufficient in themselves to render concerns of non-neutral operator behavior mute. Finally, the outcome of the resulting network neutrality arms race is uncertain and reflects the dynamic nature of the Internet. Where or if this race will end, whether regulatory intervention to steer its progress is desirable, and if so, how to intervene efficiently remain complex questions that require further research and discussion.

I. Introduction

The U.S. is in the midst of a transition to broadband access as the dominant mode for mass-market connectivity to the Internet. At the same time, the FCC and the Courts have largely dismantled the legislative framework, put in place by the Telecommunications Act of 1996, that sought to ensure regulatory-protected competitive access to last-mile access facilities controlled by the Incumbent Local Exchange (Telephone) Companies (ILECs). The ILECs have argued that the transition to facilities-based competition make such non-discriminatory access rules unnecessary and detrimental to investment incentives. Google, Microsoft, AOL, and others who depend on broadband access to reach end-users have called for "network neutrality" protection through legislative amendments to the Communications Act of 1934. The precise framing of these rules differs across the numerous bills and amendments sponsored in each house, but their

[4] The primary focus of this paper is on responses. See the companion paper by Peha discussing possible forms of network-based discrimination in more detail.

general intent is to restrict broadband providers' scope for offering discriminatory broadband services.[5]

While much of the debate over network neutrality revolves around whether discrimination is acceptable, there are conflicting definitions of discrimination. For the purposes of this paper, we define discrimination from an *economic perspective*, so discrimination occurs when differences in prices are inconsistent with differences in the costs of providing service. Where networks treat different information streams differently with respect to price or quality of service, we will refer to this as *differentiation*, regardless of whether the differences are related to cost or not.[6] Network neutrality policies are generally intended to deter practices that constitute both discrimination and differentiation, but it is possible to have one without the other. More reliable packet delivery, different burstiness characteristics (peak to average bit rates), or other specialized transport services (e.g., diverse physical routing, security filtering) may require additional resources. With these definitions, setting prices in accordance with the resources consumed would be differentiation, but would not be discrimination.

The arguments for and against "network neutrality" are complex, with merit on both sides. It is unclear whether any such rules will be adopted in the near term. To judge whether such rules would be useful, we must consider what could occur in their absence. Thus, in this paper, we assume that no network neutrality provisions have been adopted through legislation or regulation, without offering an opinion on the desirability or feasibility of such rules. In a companion paper, Peha[7] discusses a number of price and non-price strategies that may be employed by a facilities-based network provider with market power to discriminate, or act in

---

[5] For example, Rep. Markey (D-MA) sponsored legislation that would have added a new section to the Communications Act with explicit protection for network neutrality (see, H.Amdt.987, the text is available at: http://www.rules.house.gov/109_2nd/specialrules2nd109/hr5252/109hr5252_markey20.pdf).

Thus far rules to protect network neutrality have been excluded from the legislation that passed in the House in May (H R 5252, the Communications Opportunity Promotion and Enhancement Act of 2006 or so-called "COPE Act", see http://thomas.loc.gov/cgi-bin/bdquery/z?d109:HR05252:) and the Senate in June (S 2686, the Communications, Consumer's Choice, and Broadband Deployment Act of 2006, see http://thomas.loc.gov/cgi-bin/bdquery/z?d109:s.02686:), but the fight continues.

The defeat of pro-neutrality amendments was hailed in the trade press as a victory for broadband access providers like Comcast, Verizon, and AT&T; and a defeat for Internet companies like Google, eBay, and Amazon.com (see, for example, Declan McCullagh ("House rejects Net neutrality rules," CNET News.com, June 8, 2006, available at: http://news.zdnet.com/2100-9588_22-6081882.html).

[6] In common parlance, discrimination and differentiation may be used interchangeably to signal behavior or treatment that is "different" but that need not have any normative connotation, or if it does have a normative connotation, it may be good (as in, "to use good judgement") or, perhaps more commonly bad (as in, "prejudiced treatment" or "racial discrimination"). Since it seems that "differentiation" is a less loaded term, we will prefer using it when we want to refer to discrimination or differentiation that may not reflect economic discrimination. Furthermore, as will be explained further below, end-users may employ the strategies we discuss whether the operator is engaging in economic discrimination or not, and economic discrimination – if it occurs – may be welfare enhancing or not.

[7] See Jon Peha (2006), "The Benefits and Risks of Mandating Network Neutrality, and the Quest for a Balanced Policy," paper prepared for the 34th Research Conference on Communication, Information, and Internet Policy (www.tprc.org), George Mason University, September 29-October 1, 2006. www.ece.cmu.edu/~peha/balanced_net_neutrality_policy.pdf.

contradiction to the intent of "net neutrality." Here, we investigate how end-users may respond in the face of efforts by a facilities-based provider to differentiate in some of the ways described by Peha. For this investigation, we further assume that network operators have sufficient market power so that concerns over harmful economic discrimination occurring remain relevant, although as we will explain further below, the welfare implications of end-user responses remain ambiguous in any case, even when used in response to cost-based differentiation.

We understand that the viability of workable end-user strategies depends on a number of factors that are inherently uncertain such as the progress of technical and business model innovation, future regulatory reforms, and the evolution of market competition. We offer an analysis of prospective strategies in light of our understanding of current technology, industry economics, and the implications for regulatory policy. Moreover, since end-user responses are likely to elicit further responses from broadband access providers, this analysis is partial, examining only the first round in an on-going arms race. Nevertheless, our analysis does provide some interesting insights for policy-makers seeking to dissect the net neutrality debates.

First, even in the absence of additional network neutrality regulation, end-users have a variety of strategies available for responding to discriminatory practices from carriers. This suggests that the appropriate counterfactual case to consider (in the absence of net neutrality regulation) is a world in which end-users are confronted with discrimination and respond to it. Second, providers may find it difficult to maintain forms of discrimination or differentiated treatment that are welfare enhancing. Third, that the availability and effectiveness of responses to discrimination are likely to vary not only by geography but also by the level of skill and economic resources available to particular customers, raising potential equity issues. Moreover, as we explain further below, the effectiveness of end-user responses depends critically on the mode of discrimination employed by the operator. We conclude that although there are a range of end-user responses possible, that the concerns that motivated calls for net neutrality regulation remain valid. Finally, that the Internet is a dynamic system. Behaviors by operators that certain end-users find undesirable will elicit user responses that will elicit further responses from operators. This arms race has its own costs associated with it. Whether or how the resulting "arms race" may conclude and whether policy interventions might represent an efficient response are complex questions that have been poorly addressed in the network neutrality debates thus far.

The balance of this paper is organized into four sections. In Section II, we offer further clarification on the network neutrality debate and what we mean when we say an operator is engaging in differentiated/discriminatory behavior. In Section III, under the assumptions that there are no new "network neutrality" rules and broadband access providers engage in discriminatory practices like those outlined, we present a three-part taxonomy of end-user responses to such discrimination. In Section IV, we discuss the implications of such strategies for policy and the evolution of broadband markets. Section V concludes.

II. Setting the Stage: the Network Neutrality Debate

The future of Internet access is broadband. Putting the requisite infrastructure in place to support next generation broadband platform services requires substantial investment in new facilities. Crafting a regulatory framework to promote effective competition and provide appropriate investment incentives has posed a daunting challenge for policymakers. Traditionally, access to last-mile telecommunications facilities have been regulated as a natural

monopoly and have been subject to substantial common carrier regulation. In light of the emergence of alternative facilities-based broadband access platforms such as those provided by cable television companies, 3G mobile service providers, and a variety of other next-generation providers such as municipal networks, power companies, and WISPs, proponents of further deregulation argue that facilities-based competition is sufficiently robust to effectively eliminate the last-mile bottleneck, and therefore, render continued open access regulations of such facilities unnecessary.[8] However, the extent of facilities-based competition varies significantly by locale, and even where available, two wired platforms in a market may not offer sufficient competitive discipline to eliminate concerns about an abuse of market power.

If facilities-based providers have market power, they may use that market power to extract surplus monopoly rents, and perhaps more importantly, to distort the evolution of broadband applications and services. With sufficient market power, such facilities-based providers could engage in price and non-price discrimination to pursue anticompetitive goals, and thus adversely impact the interests of end-users, nascent competitors, and other participants in the industry value chain.[9] Because of such concerns, a number of consumer groups and industry participants such as Google, Microsoft, and AOL that rely on last-mile broadband access to reach their end-consumers have called for "network neutrality" legislation. The goal of such legislation is to protect against last-mile broadband access providers engaging in harmful discriminatory practices.

Unfortunately, while it is plausible to believe that broadband access providers with market power may seek to engage in harmful discrimination,[10] it is also reasonable and probable that there are many contexts in which discrimination is welfare enhancing and differentiation is consistent with robust competition. For example, end-users are likely to applaud operator efforts to block *malware* traffic such as viruses or distributed denial of service (DDoS) attacks or from malfunctioning devices that may be congesting network resources unintentionally (e.g., a workstation that has gone haywire). Furthermore, the ability of the Internet to support diverse applications efficiently is enhanced by the ability of operators to selectively offer differentiated Quality of Service (QoS) handling on a bit-by-bit, flow-by-flow basis. For example, packet-prioritization can allow real-time (delay-sensitive) traffic such as voice telephony to be carried over congestible networks with delay-insensitive traffic (e.g., email). Such QoS is key to enabling the transition to Voice-over-IP (VoIP) that is viewed as offering benefits in enhancing convergence and expanding the scope for competition.

Additionally, the observation that different users or bit streams are charged different prices does not mean that *economic discrimination* is occurring. As already noted, from an economic perspective, discrimination is when differences in prices are inconsistent with

---

[8] See FCC (2005b), "In the Matter of Appropriate Framework for Broadband Access to the Internet over Wireline Facilities," *Report and Order*, FCC 05-150, Released September 23, 2005.

[9] Operators with market power may seek to protect their market power in order to bypass regulations (leveraging market power from regulated into unregulated markets) or by raising rivals costs. Such behavior may be implemented using a variety of price and non-price strategies (see for example, Economides, Nicholas (1998), "The Incentive for Non-Price Discrimination by an Input Monopolist," *International Journal of Industrial Organization*, vol 16 (May 1998) 271-284).

[10] See Peha (2006), note **Error! Bookmark not defined.** *supra*; or, van Schewick (2005), note 1 *supra*.

differences in the costs of providing service.[11] For example, charging the *same* price for services that cost different amounts to provide is a form of economic discrimination (*e.g.*, charging flat rate access for users that use different amounts of a congestible resource).[12] More reliable packet delivery, different burstiness characteristics (peak to average bit rates), or other specialized transport services (e.g., diverse physical routing, security filtering) may require additional resources that justify higher prices. Few object to the notion that users who consume more costly resources (e.g., require higher capacity links or send substantially more traffic) should pay higher prices. Furthermore, as the literature on two-sided markets suggests, identifying when prices are appropriately cost-based is complicated by the presence of network externalities and shared/common costs.[13] Whether end-users or content providers who communicate over a shared access platform should pay above or below incremental cost may depend on the flow of externality benefits. For example, an ISP may find it advantageous to subsidize access to certain classes of end-users to build an attractive market for advertisers; or alternatively, may charge access to end-users at above-cost rates to facilitate subsidizing free content that makes access more attractive for all. Indeed, in the presence of robust facilities-based competition, we should expect to see carriers seeking to differentiate themselves by offering a variety of what may appear to be discriminatory services as part of the normal competitive process. Of course, the ISP may also have market power and be using discriminatory practices to protect that market power or to effectively extract monopoly rents from both sides of the market. Moreover, even if the ISP has market power and is price discriminating, such behavior may be welfare enhancing. For example, Ramsey pricing represents a welfare-maximizing strategy for recovering fixed and shared/common costs using discriminatory pricing.

One might consider behavior as non-neutral when an operator treats VoIP traffic from one provider as different from another (discriminating among like streams based on source or destination differences), or VoIP traffic as different from other broadband traffic (discriminating among applications), or some bits as different from other bits (even more complex methods of per packet or flow-based discrimination). The non-neutral treatment may be reflected in different prices or in different ways in which the traffic is handled (how packets are routed, the end-to-end delay experienced by packets, or the probability that packets are dropped).

In light of the difficulty of identifying when discrimination is occurring and the many good reasons for implementing discriminatory practices, it is difficult to craft legislation that

---

[11] Metering costs associated with charging differentiated rates should be included. In the past, Andrew Odlyszko has argued that it is lower cost for carriers to over-provision the network instead of introduce QoS-differentiated prices which explains, in part, why technical QoS mechanisms that have been available for many years have not been widely used in the general Internet (see, Odyzko,, Andrew (1999), "Data Networks are Mostly Empty and For Good Reason," *IT Professional* 1 (no. 2) (March/April 1999), pp. 67-69). However, new work suggests that there may be renewed interest in such QoS techniques as broadband access expands. See, for example, Briscoe, Bob and Steve Rudkin (2005), "Commercial Models for IP Quality of Service Interconnect," *BTTJ Special Edition on IP Quality of Service*, 23 (2) (April 2005) (available at http://www.cs.ucl.ac.uk/staff/B.Briscoe/projects/ipe2eqos/gqs/papers/ixqos_bttj05.pdf) or Broadband Working Group (2005), "The Broadband Incentive Problem," MIT Communications Futures Program White Paper, available at: http://cfp.mit.edu/groups/broadband/docs/2005/Incentive_Whitepaper_09-28-05.pdf.

[12] During the 1970s and 1980s, there was extensive debate over whether the volume-discounted "WATS" service offered by AT&T to business customers reflected justifiable cost-based discounts to large users or was "unreasonably discriminatory" and therefore in violation of the Communications Act.

[13] See Rochet and Tirole (2004) footnote 3 *supra*.

selectively restricts only bad forms of discrimination. Thus, it is worthwhile asking what might happen in the absence of any regulatory rules to protect network neutrality.

For the purposes of this paper, we assume that there is a facilities-based provider (perhaps more than one) who is engaging in a mix of price and non-price strategies for differentiating among user traffic. We call all such behavior "non-neutral." It includes charging different prices for users (uses), applications, or bit handling which may (or may not be) justified by differences in costs. In the face of such behavior, certain users may find themselves receiving lower quality service (increased delay, higher proportion of dropped packets, or restricted access to content or applications) or paying higher prices (whether cost-justified or not) and may seek to offset such behavior. In the next section, we discuss the strategies that end-users may employ in response to "non-neutral" behavior by a network operator.

Furthermore, to focus attention on end-user responses, we assume that the choice of facilities-based alternatives is limited and ignore the potential that the market for broadband access services is already or will be effectively competitive. Thus, we discount the ability of an end-user to bypass discrimination by simply choosing another facilities-based provider. Even when there are multiple facilities-based broadband providers in a community (e.g., Comcast cable modem and Verizon DSL service), we assume that individual end-users are unlikely to subscribe simultaneously to multiple providers (at each point in time, each subscriber has only one facilities-based broadband provider) and that there are switching costs which preclude highly dynamic switching among facilities-based providers.[14] We note that even in the absence of such switching costs, there is likely to be a terminating monopoly if the costs of terminating traffic to end-users are not fully reflected in the prices paid by end-users. Such a terminating monopoly, can give rise to the market power necessary to support certain forms of harmful discrimination even in the presence of effective competition in originating access.[15]

As Peha explains in his companion paper, currently deployed or soon to be deployed industry technology make it feasible to implement a diverse array of discriminatory strategies based on the source/destination address of individual packets, the protocols being used, the characteristics of individual streams, or aggregate packet flows.[16] Network operators may base such discrimination solely on information included in the traffic streams or may supplement it with non-traffic information (customer billing and other third-party databases that may include indicators of a customer's willingness-to-pay). Moreover, such discrimination may be implemented at different layers (link, network, or transport layers) or places in the network (individual subscriber link or some higher back-haul aggregation point) which may also influence the information-basis available to the carrier to implement discriminatory treatment. A number of the end-user responses we consider take the form of attempts to obscure the information that the provider may use to discriminate. Consequently, the attractiveness of alternative end-user responses will depend, in part, on the strategy employed by the carrier. Thus,

[14] Under current business models, broadband access services require specialized equipment (cable vs. DSL modems) and are offered with fixed monthly fees or annual subscriptions such that it is unlikely to be economic for an end-user to subscribe simultaneously to both cable and DSL broadband services.

[15] See DeGraba (2000), note 2 *supra*.

[16] See Peha, note 8 *supra*.

we will discuss the responses in relation to specific carrier strategies for implementing discrimination in the next section.

III. Revenge of the Edge: end user responses to discrimination

In this section we describe three classes of strategies that end-users (perhaps in conjunction with upstream providers) may adopt to bypass, counter, or learn to live with discrimination or differentiation by the network operator. The first of these are attempts by end-users to bypass the broadband access bottleneck by taking advantage of alternative infrastructure that does not differentiate in the same way, including those that may be deployed by end-users, with little or no involvement of traditional service providers. This class of strategies render discrimination and differentiation less effective by increasing the attractiveness of outside options.

The second class of strategies are more complex and varied. These include direct technical and non-technical countermeasures intended to render the carrier attempts at differentiation more costly, or equivalently, less effective.

The third class of responses we consider are analogous to the first in so far as they represent improvements to the outside options available to end-users. We term these "learning to live with differentiation" in that they render the impact of the differentiation and discrimination less harmful to the user without either bypassing the network provider bottleneck or attacking the provider differentiation directly. These strategies work through complementary investments that make the poor quality of service resulting from discrimination or differentiation less relevant to end-users.

We discuss each of these in turn in the following sub-sections.

A. Bypassing Differentiation

The first class of responses arises from the observation that the ability of end-users to bypass differentiated treatment by an operator (whether it is economic discrimination or not) will depend on the feasibility and ease by which an end-user may select an alternative physical path for routing his or her traffic.[17] To the extent such options for switching to another bit-path exist, the end-user may be able to evade the effects of undesirable operator behavior.

We focus on access bypass because given the high fixed and sunk costs of access networks (particularly wired ones), access is typically the least competitive element in the Internet connectivity value chain. For example, the most recent FCC status report on broadband deployment found that over 40% of U.S. zip codes had neither or only one cable modem or

[17] Such strategies may be considered directly analogous to developments in the 1970's and 80's that enabled users to bypass monopoly telephone access networks to reach one's choice of competitive long-distance suppliers. Now, the access network provides broadband Internet connectivity, and the "long-distance" suppliers are the competitive backbone suppliers (Level 3, Broadwing, etc.).

ADSL access provider.[18] Moreover, very few areas of the U.S. have more than two viable facilities-based commercial broadband ISPs[19] and the extent to which duopoly competition will eliminate concerns over market power remain unresolved (although as we discuss further below, some of the end-user responses can exploit the fact that there are multiple facilities-based bit-paths).

Furthermore, even in areas where consumers can choose among multiple facilities-based providers, access plays a unique role in conferring a terminating monopoly. As long as a customer subscribes to only one access provider, there is only one way to deliver traffic to that customer.[20] Access providers could exploit this fact by favoring or degrading the delivery of traffic to their customers from particular sources, according to whether those sources have or have not paid additional fees to the provider. Such terminating monopoly power may exist even when there is ample competition for originating access connections if there are sufficient switching costs. Indeed, switching costs may work to mitigate the effectiveness of even originating access competition if it is costly for users to multi-home or dynamically switch between access providers.[21] Commercial access providers typically try to make such switching costly through the use of long-term contracts and bundled service offerings (telephone, television, and data offered as a package).

Therefore, concerns about harmful economic discrimination resulting from an abuse of market power have traditionally been focused on last-mile access facilities. Moreover, it is worth noting that the concerns over harmful discrimination differ for large business and mass-market (consumer and small business) because the large business customer may face lower switching costs (be more likely to avail themselves of bypass alternatives). Thus, for example, multi-homing is more likely to be a viable end-user response to differentiated treatment for a business end-user than for a mass market end-user. A multi-homed user who has multiple connections could switch dynamically among them if discrimination is experienced on one but not the other. For a large business, the opportunity cost of outages may be high enough to justify paying for multiple simultaneous connections, however, most consumers would seem unlikely to be willing

---

[18] Table 16 of the FCC's most recent broadband deployment status report shows that 13.5% of zip codes have neither ADSL or Cable Modem, and 27% of zip codes have only one of the two; at http://www.fcc.gov/wcb/iatd/comp.html, see the 7/06 release reporting on data as of 12/31/05.

[19] The FCC report (note 18 *supra*) states that alternatives to cable modem and ADSL (fiber, satellite, terrestrial fixed or mobile wireless (licensed and unlicensed), and electric power line) make up a negligible portion (1-2%) of the lines in service in residential areas. Given that the FCC also estimates satellite availability within at least 88% of U.S. zip codes, its lack of significant uptake suggests that consumers consider it a poor competitive alternative; we assume its adoption occurs primarily where it is the only option.

[20] The monopoly conferred by termination is a property of any networked industry, as true for broadband as for telephony. As an FCC report (DeGraba, 2000) noted in the telephony context: "The current requirement that carriers pay the called party's network to terminate calls confers monopoly power on the called party's network with respect to terminating access. This market power arises from the fact that the calling party's carrier, whether a local carrier or an IXC, has no alternative carrier that can terminate a call to a particular called party. Thus, the calling party's carrier must pay the terminating network whatever price it demands in order to reach the called party. In effect, each terminating carrier, no matter how small, has a monopoly over termination to its own customers. Recently in fact, IXCs have begun to complain that certain CLECs have exploited their monopoly power in termination by setting access charges that far exceed those charged by major incumbent LECs,"

[21] Multihoming involves subscribing to more than one provider simultaneously and is usually discussed in the context of enhancing reliability and performance (Habib and Chuang, 2005).

to pay for two long-term subscriptions simultaneously. Further, large businesses may find it attractive to lease circuit connections direct to a competitive backbone ISP, thus bypassing the less competitive access market. The use of alternative access networks to bypass discrimination imposed by a broadband ISP is a useful consumer response only in certain market circumstances. Most obviously, alternative access networks or ISPs must exist; the alternate providers must offer more attractive options; and the net benefits of switching to an alternative bitpath have to exceed the net benefits of following an alternate strategy.[22]

In the rest of this section, we consider whether trends toward alternative technical and institutional access arrangements might expand the range of scenarios in which bypass of the access network is a viable option. We consider whether alternative technical and institutional forms of access can lower the switching costs among networks that differ with respect to their discriminatory practices, for example by allowing access to alternative networks on a low- or no-cost basis, or by enabling the purchase of access on a more dynamic basis than the long-term subscription model. These possibilities could be enabled by cooperative arrangements among users to share each others' access connectivity, by broadband resale arrangements (typically utilizing WiFi) with business models closer to micropayments, and by municipally operated open access networks. As we discuss in the examples below, each of these trends presents different challenges as a means to bypass discrimination.

1. Cooperative Access Sharing

By *cooperative access sharing*, we mean groups of end-users who band together and, in effect, share commercially provided broadband access among themselves. Such cooperatives range widely in form, depending on the size and geographic dispersion of the user group as well as their technical, social, and economic organizing principles.[23] Common to all of them, however, is the feature that the broadband access purchased by some member(s) of the group can be utilized by other members of the group without having to pay the full long-term subscription price. Such arrangements, where feasible, have great potential for enabling access network bypass.

Cooperatives at the smallest and most informal end of the spectrum typically involve free sharing among neighbors who live within close proximity, know each other, and are reasonably technically savvy. In areas where those circumstances apply (primarily dense urban neighborhoods featuring MDU housing), such arrangements could become relatively common, but it is difficult to judge since they are rarely documented or measured. One documented example is provided by Jon Crowcroft, a professor of computer science (specializing in networking) at Cambridge University in the U.K, who describes the access sharing arrangements

[22] That is, the service offered on the best alternative network has to be sufficiently more preferred (because of price or quality differences) to justify incurring the costs of switching. Moreover, the net benefit of switching (bypass) also needs to be compared with alternative strategies of countering or living with the differentiation. These latter strategies are discussed in subsequent sections.

[23] Sandvig (2004) discusses early examples of cooperatives formed around WiFi; in addition to access sharing, he discusses cooperative actions to map the location of WiFi nodes as well as to develop WiFi-related software.

that he and three neighbors, who live within 500 yards of each other in a London neighborhood, have put in place (Figure 1).[24]

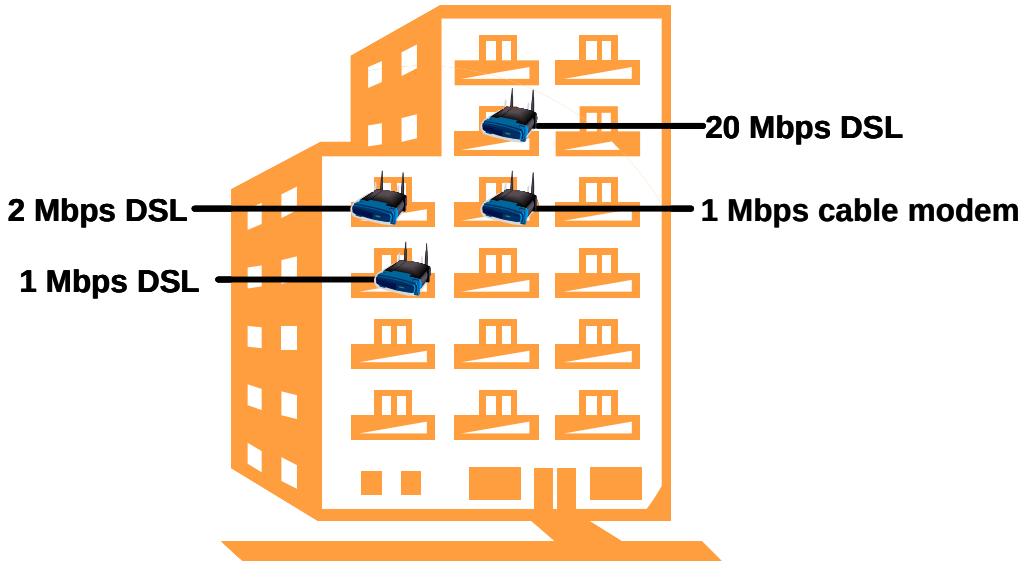

Figure 1: Example of cooperative access sharing among neighbors

As illustrated in Figure 1, each neighbor subscribes to a different DSL or cable modem access service and attaches an open WiFi router to it, enabling anyone within WiFi range to utilize their broadband connection.[25] Each neighbor is also technically sophisticated enough to know how to encrypt their traffic at higher layers (obviating the need to rely on a closed, encrypted WiFi router for privacy of user data), and to configure or otherwise modify their WiFi routers to prioritize the use of the broadband connection (to the owner first, then to the neighbors, then to others who might be passing by).

Although service outages were the primary motivation for setting up this particular cooperative, such an arrangement could easily be used to bypass differentiation, since each user in effect gains consistent and free access to another broadband service provider, albeit prioritized relative to his neighbor's own use. Such arrangements, however, are completely contingent on broadband access providers' tolerance of consumers' open WiFi access points; were that policy to change, it could represent a counter-counter-reaction in the network neutrality arms race.

[24] See Prof. Crowcroft's presentation from January 2006, "Two Open and Shut Case Studies," available at http://cfp.mit.edu/events/slides/jan06/Jon-Crowcroft-Open.pdf. In particular see slide 2, "Open Systems Interconnection 21st century style."

[25] Recall that the U.K. access market features DSL unbundling, widening the range of DSL services users might want to switch among. If this were a U.S. example, the value of this type of arrangement would only be in gaining more dynamic access to the duopoly competitor (assuming the duopolists have distinguished themselves with different approaches to discrimination), or in evading volume limits by spreading traffic among multiple connections. See the next section on technical countermeasures.

Cooperative sharing on a larger scale – both geographic and among more people – requires more formalization of technical and/or institutional organizing principles. Examples of projects that address the technical dimension include CUWin and Roofnet, while FON (discussed further below) addresses the institutional dimension.[26] CUWin (originating in Urbana-Champaign, IL, home of the University of Illinois) and Roofnet (originating at MIT in Cambridge, MA) both involve experimental city-scale deployments of wireless mesh technologies that extend university research into the design and scalability of these networks. Both projects make their technology available as open source software, facilitating experimental deployments that are also taking place beyond their home communities. Roofnet has also engendered a commercial spinoff, Meraki Networks, aiming to produce very low-cost equipment suitable for end-users wishing to organize their own wireless mesh networks.[27]

Figure 2: Wireless mesh network
Source: adapted from http://www.freepress.net/wifi/wireless_practices.pdf

Figure 2 illustrates the technical principle of a wireless mesh network, in which wireless routers relay connectivity across multiple network nodes, enabling coverage of a larger geographic area than can be reached by the individual wireless access points shown in Figure 1.[28] A wireless mesh can be self-contained (i.e. have no connection to other networks), providing only localized communications among users directly connected to the mesh. Such localized connectivity represents its own form of access bypass, in that users can talk to each other without utilizing the facilities of any traditional broadband access providers.[29] Should differentiation be

[26] See, respectively, http://cuwireless.net, http://pdos.csail.mit.edu/roofnet/doku.php, and http://en.fon.com/info/whats_fon.php.

[27] See http://meraki.net

[28] Figure 2 illustrates a *static* mesh, in which the nodes are fixed. Much of the technical research in this domain focuses on more complex *dynamic* meshes, in which nodes may be mobile and come in and out of the network.

[29] The same observation applies to the inter-neighbor connectivity shown in Figure 1. However, the number of endpoints that can be accessed in the inter-neighbor case is obviously much smaller, rendering it a less effective form of bypass. Such an arrangement is analogous to the use of a PBX for within-building or within-campus routing of corporate telephone calls.

practiced in such a way that it interferes with localized communications (e.g. imagine a provider that charges for QoS in such a way that it raises the cost of gaming among friends in a neighborhood), then end-user construction of a providerless mesh would constitute an effective bypass mechanism.[30]

If users wish to access the broader Internet, however, then at least some of the nodes in the mesh must have traditional broadband connectivity. In contrast to the one-to-one mapping of broadband connections to wireless routers shown in Figure 1, the wireless mesh's relay capability enables one-to-many sharing.[31] With this architecture, it is feasible for a group of mesh users to band together to share the purchase or use of "business class" broadband connectivity as a means to bypass discrimination. By business class, we mean broadband access that is more expensive than consumer-grade connections, but clearly allows resale and is less discriminatory.

Making this model work, however, depends on the existence of a viable financial arrangement for supporting the cost of these connections. Existing cooperatives have typically relied on the generosity of institutional users willing to share their business-class broadband connectivity for free. For example, CUWin's broadband "uplinks are paid for by the city, the university, local businesses and non-profits, and individual residents -- all of whom have decided (individually) to share their lines."[32] While this solution is relatively simple to implement and has served existing deployments well, for such projects to scale beyond university-based trials, a richer range of solutions and more operational experience are needed. Indeed, the structuring of incentives to share resources with non-trivial costs in mesh networks at scale remains a topic of active academic research.[33]

2. Broadband Resale

While research-inspired networks have focused primarily on the technical issues involved in wireless meshes that can share broadband uplinks, at least one business has been founded to

---

[30] Alternatively, localized communications could be viewed as a form of "living with discrimination" – if the provider discriminates only on traffic that goes beyond the local area, the user responds by confining their communications only to the local area.

[31] CUWin is also working on adding "multi-gateway and multi-pathing support -- which would allow users to aggregate bandwidth across multiple uplinks." Personal email from Sascha Meinrath of CUWin, August 14, 2006.

[32] Personal email from Sascha Meinrath of CUWin, August 14, 2006. In another model, NYCWireless, a cooperative in New York City that uses CUWin's software, has established partnerships with non-profit entities to support the broadband uplink required to provide free access in public parks (see http://www.nycwireless.net/about). This type of access is a less effective mechanism for bypassing discrimination, at least in most geographic areas that do not feature inviting outdoor weather year-round.

[33] There is a large literature on inducing cooperative resource sharing in networks. For example, the recent ACM-sponsored NetEcon conference, June 11, 2006 (see http://www.cs.duke.edu/nicl/netecon06/), included a number of such papers. See, for example, Pai, Vinay and Alexander Mohr (2006), "Improving Robustness of Peer-to-Peer Streaming with Incentives," or Bauer, Steven, Peyman Faratin, and Robert Beverly (2006), "Assessing the assumptions underlying mechanism design for the Internet," . This paper was representative of a number of papers at the NetEcon06 workshop on incentives to share resources in mesh/grid networks. Or, see, V. Srivastava, J. Neel, A. MacKenzie, R. Menon, L.A. DaSilva, J. Hicks, J.H. Reed and R. Gilles, "Using Game Theory to Analyze Wireless Ad Hoc Networks," *IEEE Communications Surveys and Tutorials*, 4th Quarter 2005; or Akyildiz, Ian, Xudong Wang, and Weilin Wang (2005), "Wireless Mesh Networks: A Survey," Computer Networks 47 (2005) 445-487.

experiment with new institutional models for access sharing. FON is a UK-registered startup with marquee venture capital funding and close to 80,000 registered users worldwide at the time of this writing.[34] Rather than build a mesh or in fact any new infrastructure, FON creates a community of registered "FONeros," each of whom contributes their own broadband connection and wireless router, suitably configured with FON-provided software that enables users to function as either *Linuses* or *Bills* (Figure 3). *Linuses* elect to share their access, i.e. participate in a barter-based exchange with all other registered FONeros who come within range of their wireless router. *Bills* have the additional capability of providing paid access to *Aliens*, essentially non-FONeros who access FON routers using a "pre-paid" business model. FON then shares revenues from *Aliens* 50-50 with its *Bills.*

Figure 3: FON Business Model
Source: http://en.fon.com/info/whats_fon.php

While FON's *Linus* model offers cooperative infrastructure sharing for free, its *Bill* and *Alien* models represent broadband resale with a payment style (pre-paid) that is much closer to micropayments than to long-term subscriptions. Thus, for users who have cleared the hurdle of becoming a FONero or purchasing pre-paid *Alien* access, using FON to bypass differentiation presents low transaction costs. However, for this solution to be feasible on an ongoing basis, a user must be able to depend on being able to find FON access where needed. Although FON would seem to have many users, they are spread all over the world, leading to low density of FON access points currently.[35] In addition, as with the Crowcroft model, FON's model depends critically on the willingness of broadband access providers to tolerate the connection of FON-configured routers to their network. In particular, it is not difficult to imagine that broadband providers would detect the presence of *Bills* and require the purchase of more expensive business-class broadband in return for the privilege of allowing resale.

[34] FON's investors include Google and Sequoia Capital; see http://en.fon.com/info/our_investor_partners.php. Also see http://en.fon.com/; on August 15, 2006 this page lists 76,736 "Foneros."

[35] For example, as of August 11, 2006, the maps at http://maps.fon.com/ showed no FON users in the two Boston-area towns that are home to the MIT-based authors of this paper.

Complementing FON's *Bill/Alien* model, WiFi-based access in coffee shops and the like provides another form of broadband resale that could enable users to bypass differentiation.[36] Although in many cases such access is nominally free to users, we consider it a form of resale, given that associated costs are typically embedded in the prices of the food or other merchandise that users pay for. Free access in a coffee shop obviously presents low transaction costs, but only for users within range. Like access in public parks, this kind of "hotspot-only" access enables bypass of differentiation only at the margins, and only to the extent that the coffee shop purchases wired access links that are themselves free of differentiation..

3. Municipal Open Access

How does municipal entry into broadband markets affect the potential for users to bypass differentiation or discrimination? The answer depends on how the municipal entry is structured. Much municipal broadband activity simply represents the addition of another traditionally structured facilities-based competitor to a local marketplace, whether via a Municipal Electric Utility (as is the case with most municipal wired broadband deployments) or a private-sector ISP partner (common in municipal wireless deployments).[37] With this structure, the fact of municipal involvement is largely irrelevant to the question of whether the new entrant will engage in discriminatory practices; more relevant is simply the fact that an additional facilities-based competitor has been added to the local market.

On the other hand, some municipal entry follows an "open access" model, in which the public sector operates the physical network facility and provides only wholesale transport, while multiple ISPs sell retail services to consumers.[38] In this case, the public mission of the entity operating the physical network would discourage discriminatory practices at that layer, which makes differentiation less likely. And while the retail ISPs may indeed have the incentive to discriminate, the open access structure can enable the presence of many more competitors at that layer than are typically present when all ISPs must also operate their own physical network. In this case, the municipal entry did not add one competitor, but possibly many.

4. Prognosis

Bypass is not possible in monopoly broadband markets (there has to be an alternative bit-path to switch to) and not necessary in vigorously competitive ones (where consumers have lots of choices they can easily switch to). The alternative infrastructure trends discussed in this section have the potential to enable more dynamic switching among providers and are valuable

---

[36] See, for example, the description of publicly accessible wireless hot spots in Austin, TX in (Fuentes-Fuentes-Bautista and Inagaki, 2005), and the list of cafes etc. providing free wireless connectivity (not necessarily via the CUWin mesh) in Urbana-Champaign, IL at http://cuwireless.net/hotspots?PHPSESSID=119398262b20abf90f6df7b95de0328b .

[37] See Gillett, Sharon, William Lehr, and Carlos Osorio, "Municipal Electric Utilities' Role in Telecommunications Services," forthcoming in *Telecommunications Policy;* Gillett, Sharon, "Municipal Wireless Broadband: Hype or Harbinger?" 79 *Southern California Law Review* 561, 2006; and Sirbu, Lehr and Gillett, note 38 *infra*.

[38] See, Sirbu, Marvin, William Lehr, and Sharon Gillett (2004), "Broadband Open Access: Lessons from Municipal Network Case Studies," paper presented to the 32nd Annual Telecommunications Research Conference, October 1-3, 2004, Arlington, VA (available at: http://cfp.mit.edu/groups/broadband/muni_bb_pp.html).

for lowering switching costs in areas with intermediate levels of competition. However, each trend is also associated with significant limitations and caveats, suggesting that the viability of infrastructure-based bypass will depend a lot on local context and should not be regarded as a strong or generally applicable response to harmful discrimination, if it occurs.

Although municipal open access would ensure a neutral facilities provider while considerably lower switching costs among ISPs, it is only common in two states in the U.S. (Utah and Washington), where this structure is required by law. Wireless mesh networks might provide a mechanism for users to bypass consumer-oriented ISPs who may discriminate, but remain experimental both technically and institutionally. Options for broadband resale, such as the FON system or WiFi hotspots in coffee shops, are geographically limited -- a user encountering discrimination on her home connection would have to get up and go somewhere else to bypass it (an unlikely proposition in most cases).

In addition, several of the scenarios discussed above rely critically on the contractual and technical ability of users to attach a WiFi device to their consumer-grade broadband connection and open it up either generally or selectively to other users. This capability cannot be taken for granted in the consumer market, where less technically advanced users are generally happy to accept their broadband ISP's offer to set up their in-home WiFi network for them, thus giving the provider control over access to the WiFi router (i.e. the opportunity to close up what might have been an open network or a customizable device). For example, some broadband providers supply their customers with wifi bridge/router devices that allow the network to monitor all MAC addresses. Customers may be required to register a limited number of MAC addresses in advance, and traffic from unregistered devices can be blocked, or the network can simply add extra charges. This strategy can be used to combat sharing and resale.

On the other hand, at least some of the mechanisms that providers use to discourage WiFi sharing and resale may be amenable to technical workarounds. Although development of such workarounds is definitely a marginal phenomenon (confined largely to the "geek" community), their use need not be. The Internet's success at fostering rapid adoption of innovation suggests that "geek" workarounds can migrate into the mainstream much more readily now than in the past. We now turn our attention to the broader range of technical countermeasures to discrimination.

B. Countering Differentiation

Rather than seeking to bypass the discriminatory bit pipe via the strategies described in the prior section, an end-user may seek to attack differentiation directly. Such countermeasures may involve technical or non-technical responses, may require sophisticated expertise or may be relatively simple, may be effective if employed individually (virally) or may require coordinated action, and may be in response to economic discriminatory behavior (whether welfare enhancing or not) or to differentiated treatment (e.g., to circumvent higher charges or reduced QoS for more intensive uses of scarce resources).

In this section, we focus principally on technical responses since these are likely to be less well-understood by the general reader. However, it is worth noting that end-users also have a number of non-technical responses available for opposing discriminatory treatment. One of the most potent responses of end-users in a market is to shine a light on bad behavior. Letter writing campaigns, empowered by the Web and email, can be used to call attention to practices by

operators that harm end users. These campaigns can provide a powerful inducement for operators to behave in welfare-enhancing ways. Even a relatively small number of dedicated technical sophisticates can have a big impact on industry behavior and policy. Sharing information on alternatives, prices, and practices can improve market efficiency and enhance competition by making consumers more aware of their options.[39] This can include better information regarding bypass alternatives or technical countermeasures. Of course, the fact that everyone can have a voice in the new Web-enabled marketplace means that the signal-to-noise ratio may be woefully low.

Another non-technical strategy for countering the impact of discriminatory or differentiated treatment is to simply lie. End-users may misrepresent information in subscription forms. For example, service providers commonly charge business users higher prices than residential users, motivated in part by the higher traffic loads and service quality requirements -- and hence costs -- anticipated from business customers. Users can self-select into the wrong (from the operator's perspective) usage categories or can disregard acceptable use policies in service agreements (e.g., prohibitions against running a server). Of course, if the discrimination is based on actual traffic characteristics, the operator may be able to identify liars and enforce penalties (including throttling or discarding the traffic from disallowed uses).

Some of the non-technical strategies are more effective if coordinated. For example, letter writing campaigns or service boycotts to discourage discriminatory practices are likely to be more effective the larger the number of participants. Other strategies such as lying are conversely more likely to be effective if uncoordinated since mass lying is likely to induce carriers to adopt better verification techniques (which may include basing discrimination on actual traffic patterns).

In the balance of this section, we focus on technical countermeasures. These can take a variety of forms. Many of these take the form of attempts to obscure or hide the information used by the operator to differentiate, so are akin in some respects to "lying."

As described by Peha (2006),[40] deviations from network neutrality are realized when an operator uses various forms of information about a traffic stream to provide differentiated treatment to that traffic. Modern routers contain software modules described as *packet classifiers* whose job is to determine the appropriate *class* to which to assign a packet. Packets within that class are then given a particular treatment with respect to *packet scheduling*, *routing* or *packet dropping* from buffer queues. For example, in order to discriminate against or in favor of traffic from a particular content provider, the stream of packets from that provider may be classified using the packet's IP source address which generally indicates the computer from which the traffic originated. If that computer can be associated with a particular provider, than traffic so classified could be prioritized, dropped, or rate limited, by the ISP's routers providing treatment which is different from that accorded other classes of traffic from different sources, and therefore not neutral.

---

[39] In this regard, software for measuring and quantifying discrimination will be increasingly important. See McMillan, R., "Black Hat: Researcher creates Net Neutrality test," Computerworld, Aug 2, 2006, available at http://www.computerworld.com/action/article.do?command=viewArticleBasic&taxonomyName=government&articleId=9002154&taxonomyId=13.

[40] See Peha (2006) note 10 *supra*.

Peha lists a variety of characteristics that could be utilized to classify traffic, including:

1. link layer protocol
2. IP address of source and/or destination
3. Upper layer protocol field in the IP header (*e.g.* indication of whether the upper layer protocol is TCP or UDP
4. Type of Service (TOS) field
5. Packet length
6. Interpacket spacing
7. Transport layer protocol source port or destination port. Because many applications use *well-known ports*, application type can frequently be inferred from port numbers
8. Application header and content information determined from *deep packet inspection*

Figure 4. shows the format of an IP v4 header and the header of a contained TCP protocol data unit (PDU). In effect, virtually any of the fields in these headers could be used to classify traffic and thus provide the basis for differentiation.

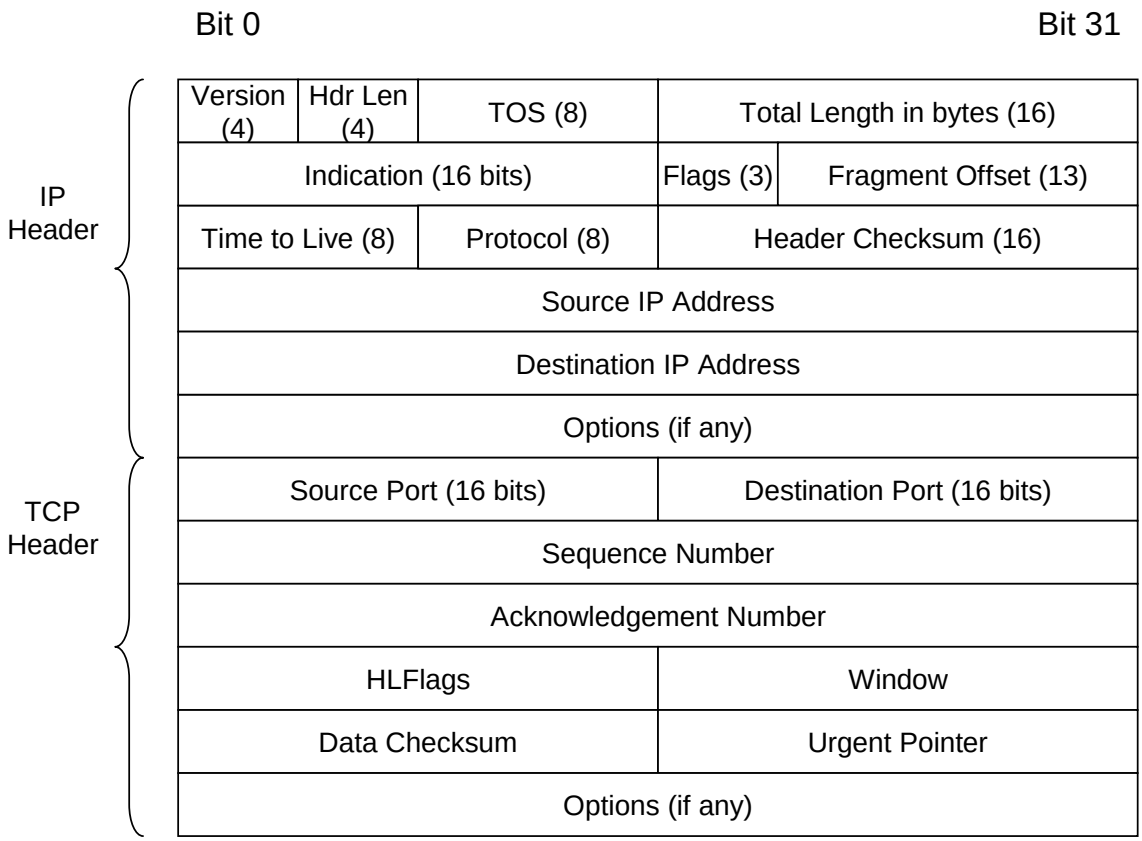

Figure 4. TCP/IP Headers

Suppose that a residential customer or service provider is aware that a particular broadband access provider is engaging in differentiation involving classification using one or

more of these fields. In this section we elaborate on the steps users can take to impede various forms of packet classification and consequent discrimination.

Performance enhancing vs performance degrading differentiation

It is useful to distinguish between two forms of differentiation. Suppose for the moment that the default is to treat all traffic on a uniform, first-come, first-served, "best effort" basis. An operator can either provide classified traffic with enhanced performance, relative to the default (for example by giving such traffic priority with respect to packet scheduling or reduced probability of packet dropping) or by deliberately degrading the performance of classified traffic (either by dropping packets, or scheduling them infrequently, and thus rate limiting the classified traffic). Note that the laws of queuing theory[41] say that changes from first-come, first-served packet scheduling policy cannot shorten the overall average per-byte queuing delay at routers. Thus, giving priority scheduling to some packets—and thus shortening their queuing delay—automatically implies that the remaining, or default traffic experiences degraded performance relative to the case when all traffic is treated equally. It is thus disingenuous of operators to state that giving priority to some traffic has no effect on the remaining traffic. The converse is also true: providing reduced performance to some traffic (e.g. P2P) will shorten the delays experienced by the remaining traffic. In general, impeding classification, as a strategy, makes it impossible to selectively degrade performance of the classified traffic: if traffic cannot be classified, than all traffic must be given the same default treatment. Conversely, impeding classification is not an effective response to performance enhancing differentiation, as the enhanced performance is only provided to classified traffic.

Port Numbers

From a network operator's perspective, the easiest way to differentiate based on application is to use well-known port numbers to infer the application and to classify the resulting traffic. Thus, when Madison River blocked Vonage's VoIP traffic, it is alleged to have done so by blocking traffic on port 5060, the well-known port assigned to the Session Initiation Protocol (SIP) used to setup VoIP calls.[42]

In response to performance degrading classification based on port number, users (or rather application writers) may rewrite software to use random ports, or to use well-known ports typically associated with other applications. Thus, many P2P applications are able to use port 80—normally associated with world-wide web traffic—as the port for downloading shared files. Indeed, the powerful capabilities of the world-wide web's HTTP protocol means that port 80 can be used for generic remote procedure calls, allowing virtually any application type to use port 80 for communications between clients and servers.[43] If P2P traffic uses the same port number as

---

[41] More precisely, this result holds only for infinite buffer queues. See Prabhu, N. (1997) *Foundations of Queuing Theory*, (Kluwer, Boston).

[42] FCC (2005) "In the Matter of Madison River Communications, LLC and affiliated companies," FCC Order DA 05-543, March 6, 2005. Because SIP is used for more than VoIP (*e.g.* video conferencing, instant messaging, buddy lists) blocking port 5060 caused collateral damage to applications other than VoIP.

[43] Temming, R. and H. Meet (2002) "SETIRI—Advances in Trojan Technology" presented at the 2002 Blackhat Conference, available at http://www.blackhat.com/presentations/bh-usa-02/sensepost/bh-us-02-sensepost-notes.pdf.

WWW traffic, discrimination based on classification by port number will end up encompassing WWW traffic as well as P2P, an undesirable result for the network operator.

Encryption as a Response

Packet classification based on header information derived from layers above the IP layer can be completely obscured by the use of encryption on all data following the IP header. The Internet IP Security protocol (IPSec) provides a standard means for senders and receivers to encrypt all of the contents of an IP packet, other than the IP header itself. Another standard, the Secure Sockets Layer protocol (SSL) provides a means for encrypting all traffic on a TCP connection. Thus one response by end users to classification based on port numbers, or any form of deep packet inspection, would be to begin to use IPSec more widely, thus obscuring port numbers and higher level packet information. Some ISPs are alleged to rate limit traffic associated with P2P applications, such as BitTorrent, using deep packet inspection.[44] A response has been to create new versions of P2P software which encrypt the traffic, and thus defeat the packet classification based on port number, and all forms of deep packet inspection.[45]

Note that, encryption at the application level, while sufficient to provide privacy, does not serve to disguise the application from a packet classifier. Thus a VoIP call, set up using SIP running on well-known port 5060, is readily identifiable as SIP, even though, following an initial exchange of packets, subsequent application packets are encrypted at the application level.

There is a cost, of course, to using encryption, including the burden of increased CPU loads on both clients and servers. However, for applications such as P2P, this burden is widely distributed, and thus not problematic given modern processors. For client server applications, the burden on servers is more significant, but specialized hardware has been developed to support the use of encryption for electronic commerce applications, and this can reduce the economic cost of the additional processing burden for service providers.

IP Address

Because the IP destination and source address must be available so that routers can route a packet, standard IPSec leaves these fields in the clear. However, there are several methods that can be used to disguise either the ultimate source or destination address of an IP packet, which can impede packet classification based on these values. Common approaches include proxies, VPNs, and various types of routing anonymizers.

Web proxies are a common form of application proxy, frequently used to reduce the volume of web pages fetched from distant locations. For example, a firm might use a web proxy

---

[44] Azureuswiki (2006), "Bad ISPs" List of ISPs believed to be limiting BitTorrent traffic. Available at: http://www.azureuswiki.com/index.php/Bad_ISPs.

[45] Torrentfreak (2006) "Encrypting Bittorrent to take out traffic shapers" Available at http://torrentfreak.com/encrypting-bittorrent-to-take-out-traffic-shapers/.

at a branch office to cache pages fetched on behalf of user A at the branch, in order to respond to a similar request by users B and C from the cache, and thus avoid using expensive access bandwidth to retransmit commonly viewed pages. Use of a web proxy has two impacts on packet classification. First, the pages viewed by B and C do not cross the ISP's network, and thus are not subject to any form of classification or discrimination. Second, requests to any web server outside the branch always appear to come from the IP address of the proxy, not of the ultimate end user. Thus, it is impossible to treat users A, B and C differently, based on their IP address, as only the IP address of the proxy is visible outside the branch office.

Of course, the network operator can provide uniform differentiation to all traffic destined for the branch by classifying based on the proxy's IP address. However, proxies can be operated by third parties on behalf of millions of users. For example, all web requests from customers of AOL appear to originate from one of a small number of AOL proxies, making it impossible to differentiate among the millions of AOL customers.[46] In Portugal, third party proxies serving unrelated consumers are common as they serve to convert requests for internationally served web pages into domestic requests, thus reducing the volume charges of Portuguese ISPs which are higher for packets sourced internationally than for those sourced domestically.[47]

A common service on the Internet is layer 3 Virtual Private Network service (VPN). In a layer 3 VPN, implemented using either IPSec or Point-to-Point Tunneling Protocol (PPTP), a source host creates an IP packet, possibly from an IP address space different from the address space assigned by his ISP. This packet is then encapsulated, and possibly encrypted, inside another IP packet. This outer packet has a source IP address from a space recognized by the broadband access provider, and a destination address of a VPN server. All traffic, for whatever destination, and using whatever higher layer protocols, is carried from the source to the VPN server using this encapsulation. At the VPN server, the encapsulation is removed, and the packet is reinjected by the VPN server into the Internet where it travels to the original destination IP address.

The use of a VPN server has several consequences.[48] First, it hides from the local broadband access provider the true destination (source) of originating (returning) traffic. From the perspective of the broadband access provider, all traffic appears to be going between the end user and the VPN server. Thus, it is impossible for the access provider to classify based on the IP address of the remote service, because only the address of the VPN server is ever visible to the access provider. The access provider cannot know from the IP address if the ultimate correspondent is a VoIP company, Movielink, or a corporate web site. Moreover, if the VPN tunnel is encrypted, deep packet inspection is not possible either. The use of a VPN service also eliminates one of the barriers to wider use of IPSec on the Internet: clients need only establish shared keys with the VPN server, not with every possible correspondent.

---

[46] AOL (2006) "AOL Proxy Info," AOL Website, Available at http://webmaster.info.aol.com/proxyinfo.html.

[47] Portugal Telecom (2006), "SapoADSL," Available at: http://adsl.sapo.pt/prodtarif.html, Visited August 16, 2006.

[48] Felton, Edward, (2006), "Nuts and Bolts of Network Neutrality," version of July 6, 2006, Available at http://itpolicy.princeton.edu/pub/neutrality.pdf.

In practice, the use of a third party VPN service by a residential subscriber can hide the identity of the parties with which a subscriber is communicating or from which the subscriber is receiving content. It cannot hide the identity of the subscriber herself, when the access network terminates at the premises of an individual subscriber. Only where there is a premises network, shared by multiple end users--such as a corporate office, or a University campus--can the identity of the individual end users be partially concealed, though they still can be associated with the institution.

The network access provider can still discriminate against *all* traffic sent to the VPN service, regardless of the application type or provider identity. It just can't selectively discriminate.[49]

To be effective as a means of disguising a subscriber's correspondents, the VPN service must be provided by a party unrelated to the network provider, and must provide similar service to many unrelated end users, so that packets reinjected into the network from the VPN service cannot be distinguished by subscriber.

In an effort to prevent the use of these techniques for impeding classification, some network access providers have banned the use of VPNs or IPSec, or charge an additional fee for their use.[50] A side effect is that these measures impede the use of encryption for privacy protection.

The link between traffic origin and destination can also be obscured using routing anonymizers.[51] To realize routing anonymity, the sender encrypts a packet, including the destination address, and sends it to a routing anonymizer which removes the outer encryption and forwards the packet, much as with a secure VPN. By going through multiple hops of decryption and forwarding, it becomes virtually impossible to link origin with destination. Onion routing will not, however, prevent an access provider from identifying the ultimate source or destination of a packet if that end point is on its network.

We are already witnessing the rise of non-carrier VPN and anonymizer services which can provide the functions described above. The Tor[52] onion routing network claims more than

---

[49] *Ibid*.

[50] Fusco, P. (2000), "Comcast Cuts Home Coax Connections to VPNs**,"** *Internetnews.com*, Sept 8, Available at http://www.internetnews.com/xSP/article.php/455741, Visited Aug 28. This policy has since been rescinded. See Hearn, T. (2003), "Comcast to FCC: Virtual Private Nets are OK," *Multichannel News*, May 19, Available at http://www.multichannel.com/article/CA299451.html?display=Policy, Visited Aug 28, 2006.

[51] Goldschlag, D., Reed, M. and Syverson, P. (1999), "Onion Routing for Anonymous and Private Internet Connections," Communications of the ACM, vol. 42, No. 2, February 1999l

[52] Dingledine, R., Mathewson, N. and Syverson, P. (2004), "Tor: The Second-Generation Onion Router," in Proceedings of the 13th USENIX Security Symposium, August 2004.

800 anonymizing routers as of August, 2006,[53] and Piratepartiet recently introduced a commercial VPN "Darknet".[54]

Packet length and inter-packet spacing

Packet encryption alone may not be sufficient to disguise the application behind a stream or *flow* of packets. A classifier observing a flow of short packets between two addresses might be able to deduce from packet length and interpacket spacing that the flow is part of a VoIP call. Several router vendors claim to have such capability.[55] In response, VoIP application software could attempt to disguise a call by padding packets with useless information, or adding null packets to the flow in order to defeat such a classifier. If all packets are carried across the access network through an encrypted VPN tunnel, one consequence is to mix packets from multiple flows within the same tunnel, which can also interfere with attempts to infer application information from packet length and inter-packet spacing.

Speed tiers and performance enhancing discrimination

The maximum speed possible on a given access link is determined by the type and length of cable (twisted pair, coaxial, fiber) and the capabilities of the electronics deployed at either end. To reduce inventory costs, operators typically deploy the same technology to all subscribers. Nevertheless, in an effort to extract greater revenues from those with higher willingness to pay, broadband access providers, especially in the U.S., offer multiple service tiers which are differentiated by the maximum speed at which the subscriber may transmit or receive. These tiers do not reflect the inherent limits of the access technology deployed, but rather are implemented via rate limiting at a network switch, typically the DSLAM or the CMTS. For example, Verizon offers the following menu of DSL service and prices:[56]

768 kbps down/128 kbps upstream	$17.95/mo (1 yr. commitment)
3Mbps down/768 kbps up	$29.99/mo (1 yr. commitment)

As noted above, these limits are enforced at the switch, even when the access link might be capable of 8 Mbps or more.

---

[53] "Number of running Tor routers," Available at http://www.noreply.org/tor-running-routers/ Visited Aug. 28, 2006.

[54] Paul, R. (2006) "Swedish political party offers commercial darknet access," *Ars technica*, Aug. 15, Available at: http://arstechnica.com/news.ars/post/20060815-7502.html, Visited Aug. 28, 2006.

[55] Caspian (2006) *Caspian Media Controller QOS Operation*, April 2006, http://www.caspiannetworks.com/PDF/QoS_Overview.pdf.

[56] Verizon, "Verizon Online DSL Pricing Plans," Available at http://www22.verizon.com/ForHomeDSL/channels/dsl/popups/verizononlinedslpricingplans.asp, Visited August 16, 2006.

The lowest speed tier is not sufficient to transmit NTSC quality digital video in real time, and certainly not HDTV. Of course, the network owner could decide to suspend this artificial limit when transmitting its own video service offering. Where does this leave an unaffiliated competing Video Service Provider (VSP) which might want to offer real time video content? The broadband access provider might propose to this VSP the following: in return for a premium payment by the VSP, the access network will ignore the normal rate limit when forwarding the VSP's traffic. Rather than limit its customer base to those who have purchased a higher speed tier sufficient to carry the video stream, the VSP would be able to market and deliver even to those who have selected a lower speed tier, albeit at a price—a price which the network operator need not levy against its own VSP.

Obviously, such discriminatory treatment provides a significant economic advantage to the affiliated VSP. Is there anything the network operator or the end user can do to circumvent this discrimination? Unfortunately, there is very little that can be done to get around such behavior. Since the default speed limits impede any source of real-time video delivery, the network access provider's intervention is required to make any such delivery possible. Offering the VSP an exception to the default limits is a form of performance enhancing discrimination which will only be provided to traffic classified as coming from the VSP and paid for. Impeding classification using the mechanisms described above would deprive the VSP of the performance enhancement.

If the rate limiting is implemented at a more distant level in the access provider's network, such as at a regional router, delivery using P2P technology, such as BitTorrent, from other subscribers within the region, would avoid the rate limiting imposed by the regional router.[57] Unfortunately, as noted above, such rate limits are generally imposed at the very first concentrator switch, *e.g.* the DSLAM or the CMTS, and thus cannot be evaded in this manner.

The only other alternative for the VSP is to focus on pre-fetching and non-real-time video delivery as discussed in Section C below.

Differentiation at layer 2.

The physical link between the subscriber and the first network switch can be subdivided into several logical channels using link layer protocols (*e.g.* separate ATM Virtual Circuits, separate virtual LANs, separate DOCSIS service flows). These logical channels share the capacity of the link, and network hardware can control the share of channel capacity, on a static or dynamic basis, that is allocated to each logical channel. One strategy by which a network access provider can favor its own services is to use one logical channel for providing best effort access to the Internet and third party applications, while reserving a separate logical channel for the provider's own equivalent services. If the provider's own logical channel is prioritized over the access link relative to the Internet access logical channel, the result is the same as if the prioritization were occurring at the IP layer. For example, in providing its Digital Phone™ VoIP service, Time Warner reserves a separate logical channel for its own VoIP traffic, while traffic

[57] AOL's In2TV service uses P2P software to deliver HD movies to viewers. See http://television.aol.com/in2tv. Visited Aug. 30, 2006

directed at a third-party VoIP supplier, such as Vonage, must travel over the lower priority Internet access channel. Thus Time Warner provides differentiated access to its own phone service relative to competing services.

Control of logical channel assignments is implemented both at the network switch and in the network terminating device (cable modem, DSL modem, optical network termination). Given operator control over the hardware at both ends, there is little that an end user can do to override these controls.

Costs of countermeasures.

End user efforts to defeat packet or flow classification by carriers are not without their costs. Encryption chews up processor cycles, may introduce start-up delays while encryption keys are exchanged, and increases packet size (or reduces usable data per maximum-length packet). Users may also need to pay service fees to a commercial darknet provider. There is also the cost in user time to research alternatives and configure encryption capabilities.

Impeding flow classification by padding packets, or introducing null packets, adds to overall traffic per usable bit of information. VPNs and onion routers cause packets to follow longer routes, thus increasing bit miles and consumption of network capacity. In short, attempts by operators to limit some forms of traffic (such as P2P) to reduce network congestion may have the effect of actually increasing congestion as users adopt countermeasures to disguise their traffic. User specified routing, as occurs with onion-routing, may also increase an access provider's costs, if it results in greater use of expensive transit networks in place of cheaper direct peering. Onion routing can also defeat efforts by operators to use local caching to reduce network traffic. In the absence of distance and volume-based pricing, consumers may have little disincentive to use these techniques, despite the costs they impose on operators.

C. Living with differentiation

The final class of strategies we call "living with differentiation" because they encompass strategies that end-users may adopt to mitigate the pain from operator efforts to differentiate and/or discriminate. These strategies are designed such that the user can tolerate inferior quality of service (QoS), i.e. higher delays, higher packet loss rates, or lower data rates. When a network operator charges more for better QoS, the user can decide whether to pay for the QOS, or to invest resources in the ability to tolerate poor QoS. This applies whether the additional charges reflect actual costs to the network operator or not.

Three general strategies are particularly useful for end users, either alone or in combination. We discuss all three here at a high level, and then present examples of each. One is to download information earlier and store it locally. Information retrieved locally is immune from unexpected delays or losses due to differentiation, although by preloading, the user gives up some degree of interactivity. In the second strategy, users download and store information that they are not sure they will need later. When they guess right, they have avoided the risk of degraded quality of service without giving up interactivity, although when they guess wrong, they have wasted communications and storage capacity and may suffer degraded service while retrieving the correct information. In the third general strategy, processing is performed at the

sender and/or receiver to increase the value to the application of each bit that passes through the network, or more often, each bit that passes through in a timely manner.

One example of downloading information earlier and storing it longer is running streaming video through a large buffer before displaying it. For example, with a two-second buffer, the user must wait an extra two seconds before her video begins, but subsequent network delays of up to two seconds will go unnoticed. In a more extreme example, she might download all of the content well before it is used, as in podcasting. The user downloads an entire audio or video program to play back later, rather than stream the program as needed.

The viability of using time-shifting buffers may depend on the application and end-use model. This can affect the extent to which buffering may be unnoticeable (not materially impact the end-user's experience) or require a significant shift in behavior. For example, applications like VoIP, interactive gaming, and sports programming may be very sensitive to delays and become effectively unusable except with relatively short time delays (e.g., on the order of 10s of milliseconds for voice or interactive gaming, but perhaps on the order of seconds for sports programming). Other applications like video-on-demand or background file transfers may tolerate much longer delays (larger buffers) and still be consistent with the original use model.

Alternatively, time-shifting that requires significant modifications in end-user behavior (downloading programming to watch later) may not be viewed by the end-user as significantly less valuable. For example, pre-loading content to my DVR that lets me more conveniently schedule my viewing, screen out offensive content, or more flexibly (portably) view my content may be preferred to more limited versions of "real-time"[58] streaming delivery that may be available at the wrong time, include commercials, or be associated with a display in a fixed location.

Of course, we can't always preload what we will need. Another strategy is to preload many things that we *might* need. For example, a box on the consumer premises might maintain a local store of fifty movies available for immediate viewing. This box would periodically replace currently stored moves with new alternatives, but these downloads could be done over long periods when the consumer is not waiting for the transfer to complete. Thus, the consumer can choose from among fifty movies with no advance planning, and the network's quality of service is not a problem. The same technique can be used by sophisticated web proxies, which spend their idle time preloading web pages that users might want to view in the future, judging from previous experience. As a variation of this strategy, redundant information might be sent out, in case the network loses or greatly delays some of it. Why wait for data to be lost before retransmitting it? For example, a network operator could make voice over IP (VOIP) very unattractive by significantly delaying 10% of the VOIP packets. If the application simply sends every packet twice, much higher loss rates can be tolerated. (In practice, there are more effective forms of redundancy.) Ironically, this strategy of collecting information that may or may not be needed later greatly increases the amount of traffic that the network must support. If prices depend only on peak data rate, as is typical today, that is of no concern to the consumer, but it could be a problem for the network.

[58] In a packet network, all traffic is buffered so never real-time. We use "real-time" to refer to the end-user experience of no noticeable delay.

As with the first strategy, the viability of downloading excess (contingent) content may depend on the application and end-use model. For example, it may be relatively easy to predict some content needs (top box office movies but more difficult for eclectic choices) while much more difficult to predict others (source/destination of future VoIP calls). Some content may be shared among a community of users which would make buffering less costly (and may be combined with local multicast as discussed earlier), whereas other content may be end-user specific (idiosyncratic).

A third strategy involves greater processing at sender or receiver. For example, for voice or video applications, users may change their encodings to tolerate lower data rates or sporadic periods of high packet loss. Other applications could be designed quite differently to operate more effectively over networks with questionable quality of service. For example, in a multiplayer video game, a server could do all the processing and send a video signal to all players, or the server might send background images once and then send video only of the objects that move, or it might merely send information about the rate at which objects move and the client software would reconstruct a video image. Each design approach places a different load on the network, and requires a different quality of service from the network.

From the above examples of mitigating the effects of discrimination, we observe the following. First, in order to require less of the network, it is generally necessary for end users to purchase hardware with greater storage, processing, or both. Second, these techniques may require greater technical sophistication, although often that sophistication is required of the application designer rather than the user. Finally, some of these techniques may put an unexpected burden on the network.

IV. Implications for Regulatory Policy, Welfare, and the Evolution of Broadband

In the preceding section, we identified three classes of end-user responses to discriminatory or differentiated treatment by operators. On the one hand, policymakers concerned that operators may use discriminatory behavior to leverage their market power may be comforted in knowing that end-users have options available to make such behavior less effective and more costly. *Ceteris paribus*, the threat of such responses makes discriminatory operator behavior less likely in the first instance. On the other hand, as discussed earlier, the operator may be employing welfare-enhancing discrimination or cost-based differentiated pricing. In either case, the threat that end-users may effectively counter welfare-improving discrimination (e.g., Ramsey pricing) or differentiated treatment (cost-based prioritization) may pose a threat to cost recovery and continuing investment in broadband.

For example, deep packet inspection may facilitate many forms of discrimination that end users would avoid, but it also facilitates the detection and eradication of malware, and other security threats. If end users begin using encryption more often to avoid discrimination, they may make the network less secure. This may also make it harder for law enforcement to employ socially valuable wiretapping.

Even as broadband access penetration saturates, broadband traffic will continue to grow, and hence, aggregate costs also will continue to grow even if Moore's Law-like cost efficiencies act to slow such cost growth. Adopting usage-based pricing (whether some form of Ramsey or cost-based) may prove an efficient industry response, but it is unclear how the movement to such

a response ought to be coordinated since it may involve more complex and coordinated (across the industry value chain) pricing than we have seen before.[59] For example, two-sided market theory suggests that bill-and-keep, sender-only, or receiver-only payment models for traffic may each be inefficient under plausible traffic and market scenarios. Thus, at a high-level, the availability of viable customer responses to discrimination has ambiguous welfare impacts. Even in the absence of any market power, end-users may adopt strategies to counter even justified cost-based pricing or traffic management strategies.

In the extreme form, such strategies have a "beggar thy neighbor" effect. For example, high traffic users may seek to avoid paying for even the incremental costs that their traffic imposes on the network. While economics does not provide unambiguous guidance on how to best allocate shared costs, it is clear that traffic which does not recover at least its incremental costs (after properly accounting for any externalities) should not be carried. If used widely, the resulting failure to adequately recover aggregate costs will result in a "Tragedy of the Commons" which will threaten on-going investment in expanding and maintaining broadband infrastructure.

Additionally, some strategies are less efficient than others. For example, onion routing or bit-packing to obscure the true nature of traffic consumes more network resources than are necessary to move the traffic. If network resources are indeed scarce, such wasteful behavior represents a further deadweight loss to the overall system. These techniques also make it harder for the network to route around malfunctioning devices and congestion, so end users are likely to experience more disruptions.

Furthermore, the adoption and effectiveness of strategies may depend to varying degrees on the extent to which end-user responses need to be coordinated. For example, it may seem less plausible that welfare-reducing responses would be adopted *en masse* (the Lemming phenomena), while *beggar-thy-neighbor* strategies by a minority remain a credible threat to an open and flexible Internet. Nevertheless, the fact that some strategies may be more accessible or the benefits privately realized (cost shifting), implies that the viability of end-user countermeasures may raise equity concerns, in addition to overall efficiency concerns.

Because the Internet makes it feasible to rapidly disseminate information and technical fixes embodied in easy-to-use software, it is possible that even unsophisticated end-users may be able to participate in technically-sophisticated countermeasures, thus rendering the distinction between coordinated and uncoordinated adoption somewhat less clear. On the plus side, this may reduce equity concerns since the implications of end-user countermeasures may be more widely shared; on the negative side, it may make the Lemming phenomenon of movement to a collectively unfavorable outcome more likely.

In terms of their overall impact on welfare, the implications of the three classes of strategies are ambiguous, but perhaps the first and third (bypass and living with differentiation) are more likely to be welfare-enhancing. We say that because when appropriately applied, these can expand the universe of business and end-use models available in the market and expanded consumer choice and business model flexibility are generally regarded as good things. However, this need not be the case. For example, investments in bypass capacity by end-users may fail to

[59] See Broadband Incentive Problem, note 11 *supra*.

realize scale and scope economies that more coordinated, centrally-planned investment would achieve (e.g., by a traditionally regulated public utility), or one might suspect a municipality of being less efficient than an investor-owned service provider, although there is no *a priori* reason to expect municipal networks to be any less efficient (Lehr, Sirbu, and Gillette, 2006). And, some of the living with discrimination strategies may result in generating excess traffic that uses scarce network resources.

With respect to the second class of strategies – direct and indirect countermeasures – the overall welfare implications seem even more ambiguous. Since this is true about real "arms races," perhaps we should not be surprised. The evolution of enhanced end-user capabilities at the edge (more intelligent network edges) in an open Internet (one not tightly controlled by the network/operator-controlled center) offers important advantages in terms of more flexible service and business models and more competitive markets, while also posing threats in terms of collective security (viruses, SPAM, DDoS) and leading to less than optimal cost recovery.

Another lens through which to consider the policy implications of end-user strategies is their likely effectiveness. Not surprisingly, the effectiveness of end-user responses to differentiated or discriminatory treatment by an operator depends critically on the strategy employed by the operator, and on the ease with which the operator may implement a counterstrategy. For example, obscuring the source or class of traffic (using onion routing or hiding VoIP traffic by application port hopping) will not work if the nature of the discrimination is to throttle or otherwise degrade the transport of all traffic which does not pay for premium service. And, strategies that depend on end-users deploying WiFi access points (bypass via end-user mesh networks) or buffering (living with differentiation through time shifting or pre-loading contingent content) assume that the end-user controls the WiFi access node or buffer (DVR). Operators that provide (control) the end-user WiFi access point or DVR buffer effectively preclude such end-user strategies. Or, buffering strategies that may be effective against peak-limiting strategies are not effective against volume-based pricing (i.e., pre-downloading a movie over a slow or fast link is the same if the charge is based on the GB of traffic downloaded per month).

Within the three classes of strategies we identify, there are hierarchies of difficulty/likely effectiveness. For example, the cost of using end-user buffering or multi-homing increases with the size of the buffers or the multi-homing (multiple access subscriptions) maintained. With respect to hiding traffic information, discussed in the section on counterstrategies, certain types of information are easier to hide than others. For example, discrimination that is based on deep packet inspection of the content may be countered effectively with end-to-end encryption; whereas obscuring the source and destination IP addresses in packet headers is more difficult for an end-user to do and generally requires relying on new types of intermediaries (via onion routing or use of a "dark-net"), and if not done, leaves the operator with a very useful hook for implementing discrimination.

Summing up, it seems that the only end-user strategy that would reliably allow an end-user to escape service provider discrimination or differentiation is to physically bypass the provider's access network (e.g., municipal network[60]). The other strategies are all very dependent

[60] According to a recent news report, Culver City (CA) has installed filters on its muni-network to filter out pornography and P2P traffic (see, "When muni-WiFi becomes a vehicle for Muni-Censorship," Techdirt.com,

on the mode of discrimination imposed by the provider, and in some cases, the operator has relatively simple alternatives available to counter end-user responses to discrimination. For example, if users move to disguise P2P traffic by encryption, in order to avoid carrier rate limiting, the operator can introduce volume-based pricing, which will give end-users a financial incentive to self-limit their P2P traffic. In other cases, it may be more difficult for the operator to respond. An operator bent on blocking P2P traffic altogether will find it hard to overcome the use of encryption, padding, and random ports to disguise the traffic.

True physical bypass, however, is likely to be resource-intensive since the bypass network has to be substantial enough to extend to an aggregation point that allows interconnection with competitive backbone providers. For that to be the case, end-user networking has to mature significantly beyond where it is today. It seems more likely to hope that effective access competition will emerge to make the risk of adverse discrimination less likely than to hope that such large-scale end-user bypass will become generally available.

Thus, we believe that while end-users have a variety of strategies that may be used to counter provider-based discrimination, these are not sufficient to render the risk of harmful discrimination mute. Moreover, the existence of such strategies which may find use in opposing even socially beneficial discrimination and differentiation (employed to ensure legitimate cost recovery or to block malware) pose a new set of challenges for policy-makers. The prospect of an on-going "arms race" between operators and end-users makes the already difficult challenge of forecasting how broadband markets will evolve even harder. One obvious outcome of such an arms race is to have more intelligence distributed at the edges and working potentially in conjunction with but also in opposition to network-centric efforts to manage the evolution of our collective end-to-end broadband infrastructure.

V. Conclusions

In the U.S. and elsewhere, industry participants, academics, and policymakers have been debating the appropriateness and need for regulatory or legislative protections against the potential for harmful discriminatory behavior by broadband access providers. In the U.S., this has resulted in calls from certain consumer advocates, Google, Microsoft, eBay and others for new legislation that would protect network neutrality. The proposed legislation would restrict broadband access providers' ability to engage in discriminatory or differentiated traffic handling.

In this paper, we consider what might happen in the absence of any regulatory or legislative efforts to protect network neutrality. We examine the range of responses that may be adopted by end-users if they are confronted with discriminatory or differentiated (aka "non-neutral") treatment that the end-users find objectionable. We organize these responses into three classes (Section III). The first, "bypass," addresses the discriminatory treatment by seeking to avoid using the physical bit-path offered by the offending operator. This class of strategies include municipal networking, multi-homing (to facilitate carrier switching), and mesh ("carrier-

---

August 23, 2006, available at: http://techdirt.com/article.php?sid=20060823/1054224). Obviously, bypassing discrimination via an "open access" municipal network will not offer end-users relief if the municipal network is the one discriminating. As some of us have argued elsewhere (Sirbu, Lehr, and Gillette, 2004), it matters *how* the municipality implements open access.

less") networking. These work by using end-user controlled equipment to route traffic locally (downstream of the discrimination as a PBX routes within-building calls and thereby avoids hitting the local telephone switch), by bypassing the last-mile access network and connecting directly to competitive backbone providers, and by sufficiently aggregating end-user traffic to allow more flexible (and hopefully less discriminatory) access to broadband access providers (community mesh networking). While such bypass strategies offer the best assurance of successfully countering operator discrimination or differentiation, they are costly and remain speculative because the technologies and institutional details of how these might work still need more work and market-testing/development.

The second class of responses include a complex array of technical and non-technical countermeasures to discriminatory or differentiated traffic handling. Many of these operate by trying to obscure the basis on which the carrier might seek to discriminate. This class of strategies include things like end-to-end encryption, onion routing, and application port-hopping. Some of these (application port hopping) are relatively easy for application programmers to implement while others (onion routing or darknets) depend on third-party intermediation or a fairly advanced degree of end-user coordination. Furthermore, the different strategies identified are more or less effective against particular types of operator behavior. None of the technical countermeasures we identify are likely to be very effective against carrier discrimination/differentiation based on premium charges for premium service. They are most effective against relatively simple forms of operator discrimination (e.g., based on application use of a particular port).

Finally, the last class of end-user strategies we consider ("living with differentiation") include ways in which end-users can effectively make do with inferior quality access service. Through buffering to time-shift and pre-loading of contingent content or through additional end-node processing, end-users may be able to realize the benefits of premium service over lower quality (less expensive) access connections.

The fact that end-users have a range of potential responses means that the analysis of potential net neutrality legislation needs to consider an appropriate counterfactual. In the absence of net neutrality rules, end-users will have a range of options to counter discriminatory behavior. The extent to which this is likely to mitigate the risk of harm from discriminatory behavior depends in part on the effectiveness of such end-user responses, and the resources end users must invest to engage in these responses. We conclude that although there are a number of responses, there are no silver bullets. Even taking into account the existence of end-user responses, the concerns about non-neutral treatment from operators with market power remain.

Moreover, the existence of end-user responses further complicates the policy analysis. As we explain and others have also pointed out, there are many situations in which operators engaging in differentiated or discriminatory practices would be socially beneficial (e.g., to ensure appropriate cost-recovery, to sustain on-going investment or to block malware). End-user responses that render such good discrimination/differentiation more costly or less effective may be welfare reducing. Additionally, employing such end-user strategies imposes additional resource costs and may have negative spillover implications for other end-users.

The analysis presented here highlights the complexity of analyzing market dynamics. Our discussion has shown how some types of discriminatory behavior may be relatively easy for end-

users to counter (e.g., application port-switching or encryption to address outright blocking of port 5060 traffic), while other modes of discrimination will be much more difficult for end-users to address (e.g., operators charging premium prices for prioritization or layer 2 discrimination). And, some types of discrimination (e.g., based on source/destination IP addresses that implement discrimination based on the identity of the content provider) may require intervention of third party intermediaries to address (e.g., a darknet). The complexity of responses and counter-responses could result in an on-going arms-race. Such an arms-race, while offering a potential alternative to regulation (market forces address that which regulation fails to), could prove quite costly for all concerned and enhance market uncertainty, which imposes its own costs on infrastructure investment.

Policymakers confronting the broadband future need to be concerned about the potential for an abuse of market power on the one hand (realized through adverse discrimination or differentiation) and the potential for an on-going arms race disrupting future investment on the other. While the lack of clear answers may be frustrating for policymakers, it suggests that further research is needed. Furthermore, policymakers should not expect to be able to adopt a first-best regulatory solution in light of the inherent complexity of forecasting how the arms race will play out. This does not mean, however, that no regulation is necessarily the best response. Some sort of net neutrality rules may make sense to protect against obvious abuses of market power, to discipline the arms race, and to provide uncertainty-reducing guidance to market participants (broadband access providers, upstream content/application providers, *and* end-users alike) as to what the rules of the road will be. As of our writing, however, the authors of this paper have not found a recipe for such rules that we can mutually endorse.

VI. Bibliography

AOL (2006) "AOL Proxy Info," AOL Website, Available at http://webmaster.info.aol.com/proxyinfo.html.

Akyildiz, Ian, Xudong Wang, and Weilin Wang (2005), "Wireless Mesh Networks: A Survey," Computer Networks 47 (2005) 445-487.

Azureuswiki (2006), "Bad ISPs" List of ISPs believed to be limiting BitTorrent traffic. Available at: http://www.azureuswiki.com/index.php/Bad_ISPs

Bauer, Steven, Peyman Faratin, and Robert Beverly (2006), "Assessing the assumptions underlying mechanism design for the Internet," paper prepared for NetEcon 2006, ACM Conference on Electronic Commerce, June 11, 2006, paper available at: http://www.cs.duke.edu/nicl/netecon06/papers/ne06-streaming.pdf.

Briscoe, Bob and Steve Rudkin (2005), "Commercial Models for IP Quality of Service Interconnect," *BTTJ Special Edition on IP Quality of Service*, 23 (2) (April 2005) (available at http://www.cs.ucl.ac.uk/staff/B.Briscoe/projects/ipe2eqos/gqs/papers/ixqos_bttj05.pdf)

Broadband Working Group (2005), "The Broadband Incentive Problem," MIT Communications Futures Program White Paper, available at: http://cfp.mit.edu/groups/broadband/docs/2005/Incentive_Whitepaper_09-28-05.pdf.

Crowcroft, Jon (2006), “Two Open and Shut Case Studies,” presentation to MIT-Cambridge CFP-CRN Joint Research Meeting, June 2006 (available at http://cfp.mit.edu/events/slides/jan06/Jon-Crowcroft-Open.pdf).

Declan McCullagh (2006), "House rejects Net neutrality rules," CNET News.com, June 8, 2006, available at: http://news.zdnet.com/2100-9588_22-6081882.html.

DeGraba, P. (2000), “Bill and Keep at the Central Office as the Efficient Interconnection Regime,” U.S., FCC OPP Working Paper 33.

Economides, Nicholas (1998), "The Incentive for Non-Price Discrimination by an Input Monopolist," *International Journal of Industrial Organization*, vol 16 (May 1998) 271-284

FCC (2005) “In the Matter of Madison River Communications, LLC and affiliated companies,” FCC Order DA 05-543, March 6, 2005.

FCC (2005b), “In the Matter of Appropriate Framework for Broadband Access to the Internet over Wireline Facilities,” *Report and Order*, FCC 05-150, Released September 23, 2005.

Felton, Edward, (2006), “Nuts and Bolts of Network Neutrality,” version of July 6, 2006, Available at http://itpolicy.princeton.edu/pub/neutrality.pdf

Fuentes-Bautista, M. and N. Inagaki (2005), “Wi-Fi’s Promise and Broadband Divides: Reconfiguring Public Internet Access in Austin, Texas,” Telecommunications Policy Research Conference, September.

Fusco, P. (2000), “Comcast Cuts Home Coax Connections to VPNs**,”** *Internetnews.com*, Sept 8, Available at http://www.internetnews.com/xSP/article.php/455741, Visited Aug 28. This policy has since been rescinded.

Gillett, Sharon (2006), “Municipal Wireless Broadband: Hype or Harbinger?” 79 *Southern California Law Review* 561, 2006.

Gillett, Sharon, William Lehr, and Carlos Osorio (2006), “Municipal Electric Utilities’ Role in Telecommunications Services,” forthcoming in *Telecommunications Policy*.

Habib, A. and J. Chuang, “A Measurement-based Analysis of Residential Multihoming,” *Infocom 2005*, (2005)

Hearn, T. (2003), “Comcast to FCC: Virtual Private Nets are OK,” *Multichannel News*, May 19, Available at http://www.multichannel.com/article/CA299451.html?display=Policy, Visited Aug 28, 2006.

Lehr, William, Sirbu, Marvin, and Sharon Gillette (2006), "Wireless is changing the policy calculus for broadband," forthcoming in *Government Information Quarterly*.

McMillan, R., “Black Hat: Researcher creates Net Neutrality test,” Computerworld, Aug 2, 2006, available at http://www.computerworld.com/action/article.do?command=viewArticleBasic&taxonomyName=government&articleId=9002154&taxonomyId=13.

Odyzko,, Andrew (1999), "Data Networks are Mostly Empty and For Good Reason," *IT Professional* 1 (no. 2) (March/April 1999), pp. 67-69

Pai, Vinay and Alexander Mohr (2006), "Improving Robustness of Peer-to-Peer Streaming with Incentives," paper prepared for NetEcon 2006, ACM Conference on Electronic Commerce, June 11, 2006, available at: http://www.cs.duke.edu/nicl/netecon06/papers/ne06-streaming.pdf.

Peha, Jon (2006), "The Benefits and Risks of Mandating Network Neutrality, and the Quest for a Balanced Policy," paper prepared for the 34th Research Conference on Communication, Information, and Internet Policy (www.tprc.org), George Mason University, September 29-October 1, 2006 (available at: www.ece.cmu.edu/~peha/balanced_net_neutrality_policy.pdf).

Portugal Telecom (2006), "SapoADSL," Available at: http://adsl.sapo.pt/prodtarif.html, Visited August 16, 2006

Prabhu, N. (1997) *Foundations of Queuing Theory*, (Kluwer, Boston).

Rochet, J. and J. Tirole (2004), "Two-sided markets: An Overview," working paper, IDEI, available at: http://faculty.haas.berkeley.edu/hermalin/rochet_tirole.pdf

Sandvig, C. (2004), "An Initial Assessment of Cooperative Action in Wi-Fi Networking," *Telecommunications Policy* 28 (7/8), pp. 579–602.

van Schewick, Barbara (2005), "Towards an Economic Framework for Network Neutrality Regulation," Paper presented at the 33rd Research Conference on Communication, Information and Internet Policy (TPRC 2005), George Mason Law School, Arlington VA, September 2005, available at: http://web.si.umich.edu/tprc/papers/2005/483/van%20Schewick%20Network%20Neutrality%20TPRC%202005.pdf

Sirbu, Marvin, William Lehr, and Sharon Gillett (2004), "Broadband Open Access: Lessons from Municipal Network Case Studies," paper presented to the 32nd Annual Telecommunications Research Conference, October 1-3, 2004, Arlington, VA (available at: http://cfp.mit.edu/groups/broadband/muni_bb_pp.html)

Srivastava, V., J. Neel, A. MacKenzie, R. Menon, L.A. DaSilva, J. Hicks, J.H. Reed and R. Gilles (2005), "Using Game Theory to Analyze Wireless Ad Hoc Networks," *IEEE Communications Surveys and Tutorials*, 4th Quarter 2005.

Temming, R. and H. Meet (2002) "SETIRI—Advances in Trojan Technology" presented at the 2002 Blackhat Conference, available at http://www.blackhat.com/presentations/bh-usa-02/sensepost/bh-us-02-sensepost-notes.pdf

Torrentfreak (2006) "Encrypting Bittorrent to take out traffic shapers" Available at http://torrentfreak.com/encrypting-bittorrent-to-take-out-traffic-shapers/

Verizon, "Verizon Online DSL Pricing Plans," Available at http://www22.verizon.com/ForHomeDSL/channels/dsl/popups/verizononlinedslpricingplans.asp, Visited August 16, 2006.

Wu, Timothy and Lawrence Lessig (2003), *Ex Parte Submission in CS Docket No. 02*-52, Before the Federal Communications Commission, In the Matter of the Inquiry Concerning High-Speed Access to the Internet over Cable and Other Facilities, INternet over Cable Declaratory Ruling, Appropriate Regulatory Treatment for Broadband Access to the Internet over Cable Facilities, August 22, 2003 (available at: http://www.freepress.net/docs/wu_lessig_fcc.pdf).