

Network Neutrality & Network Transparency¹

I. Introduction

Significant disagreement concerning the basic technical features of internet interconnection and protocols plague debates about internet governance. For example, network neutrality, though existing only in draft legislation and a handful of Federal Communications Commission (FCC) statements, has inspired a voluminous debate, a debate in which theory and ideology has outpaced established fact. Defined quite vaguely as “a bit is a bit” or “like content is treated the same,” network neutrality appears to constitute a type of non-discrimination requirement, similar to the FCC’s other open access regimes, like open network architecture (ONA),¹ Section 68,² or local competition under the 1996 Telecommunications Act.³

Network neutrality, however, leaves key distinctions unexamined. How can (and who will) determine what content is alike?⁴ Is different treatment for spam and other malware “reasonable traffic management”? Does the discrimination involve intranet internet (the “internal internet”), *i.e.*, discrimination of traffic within one ISP or backbone, or does it involve inter-internet (the “external internet”), *i.e.*, discriminatory peering and transit agreements among ISPs and backbones? These questions are particularly nagging as very little about internet management is publically known. As the recent BitTorrent controversy in the United States reminds us, very few know very much about internal internet management protocols. Similarly, few details of peering and transit agreements are public.

We maintain that (i) internet policy cannot be made, if it is to be made at all, in an informational vacuum and (ii) disclosure can serve the dual purposes of informing any potential policy and curbing anticompetitive behavior with minimal regulatory cost to either the government or firms. Calls for disclosure, however, have been ad hoc, with various internet controversies like BitTorrent prompting specific disclosures.⁵ More important, these calls have failed to describe disclosure’s purpose or justification. They have not shown *how* disclosure would better inform consumers, curb anti-competitive abuses, or further other policy goals.

Identifying internet disclosure’s nature, purpose, and goals is particularly important given that the traditional economic and regulatory justifications for telecommunications disclosure seem inadequate in the current internet market and regulatory environment. First, disclosure’s economic justification—that disclosure allows for better consumer choice—seems inapplicable to the internet for the simple reason that given the lack of competition in the last mile connection, consumers lack choice, informed or otherwise.

¹ We thank Joachim Täg and other commenters at the 2008 Telecommunications Policy Research Conference held at George Mason University School of Law.

Second, historically, the basis for telecommunication disclosure has been regulatory. This purpose also seems inappropriate given the agency's lack of meaningful process to address internet abuses as well as the lack of any generally accepted regulatory goals that would allow for non-controversial FCC regulation—a point the Comcast-BiTorrent order discussed below makes clear.

Third, the typically “static” type of disclosure, like cost data for ratemaking, seems to be an inappropriate form for the internet. Unlike telephony, in which routes and carriers were largely fixed, identifiable, and relatively few, internet traffic travels, at least ex ante, through unknown numbers of carriers on largely unspecified routes. Internet connection quality is a probabilistic function (mathematically sketched-out below) not meaningful captured by a handful of metrics, as with telephony.

Last, markets face a collective action hurdle to providing useful information about internet access quality because the “quality” of a consumer's internet experience depends not merely on his or her relationship with the ISP, but with the ISP's interconnection relationship to other networks. No one ISP has the ability or incentive to provide information about other networks so as to provide a more complete picture of internet quality.

We argue that mandated disclosure should create “network transparency.” We look to network engineering to propose a dynamic disclosure regime that enables consumers to react and even circumvent anticompetitive network managements. Rather than reproduce a static, regulatory disclosure, we encourage a “bottom-up,” flexible type of network disclosure which reacts quickly to changes in network technology and provides the specialized, even individualized, revelations that traditional regulatory disclosure cannot. Such disclosure should provide better informed consumers; and, perhaps more important, a sort of netroots antitrust enforcement. It also would legitimate and empower third parties to take part in the project of internet governance.⁶

To achieve this end, we would require public disclosure of (i) each network's QoS policy, with reference to some industry standard like the widely used DiffServ architecture and (ii) all peering and transit agreements as well as (iii) the implementation of these interconnection agreements by mandating network's registration in the Internet Routing Register (IRR) or equivalent public registration.

Our disclosure's *immediate* goal is neither improved consumer choice nor more effective regulatory oversight. Instead, tracking the incentives that function in open source software collaborations, our disclosure would be aimed at high level network users, what we term with no doubt flattering vainglory, the “internet vanguard.”⁷ Our disclosure is aimed at those with pecuniary, personal and reputational interests to track, react to, and defend open networks, strengthening a balance, even an ecosystem, between network managers' centralizing tendencies and users' decentralizing tendencies.

II. The Challenge of Internet Disclosure

Understanding how the internet market works requires disclosure of network management practices and interconnection agreements. Yet, to improve consumer welfare and increase the amount of usable information for market decisions, as opposed to satisfy policy researchers' curiosity, it is far from clear what must be disclosed. Unlike the telephone network's quality of service, internet quality cannot be reduced into a few simple, easily understood metrics. Second, the disclosure's intended audience is not clear given its highly technical nature. Third, mandated disclosure typically aims to produce better informed consumers to create more efficient markets or better informed bureaucrats to better enforce the law. Neither justification seems compelling given that consumers right now have so little choice as to ISP. Improved information seems unlikely to improve efficiency in those highly concentrated markets. Further, the current status of regulatory oversight is haphazard and lacks any clear policy rules or goals which would allow a regulator to distinguish, in a non-arbitrary way, good from bad discrimination.

In short, it is unclear (i) what such disclosure should consist of, (ii) to whom it should be aimed, and (iii) what it should achieve. The following examines these questions.

a. Disclosure's Content

Effective disclosure must actually aid markets in developing information in assessing the value of the product. But, determining this quality or value of internet access presents difficult, even intractable, questions.

First, the scattered infrastructure owners each have their own discrimination policies and only know the identity of their set of adjacent peers. While one network can produce the information that it knows about—its local traffic policies and the peers to whom it connects—this is not enough. To fix the value of access, even the most technically competent analyst needs to know more; she needs to know about the adjacent networks.

Second, even though one could estimate an ISP's value using a recursive function, as the sum of the value of their peers, reduced by the aggregate traffic discrimination among those peers; the ISP cannot be expected to disclose the traffic policies of all its peers, and of all their peer's peers, and on and on. This would amount to a single ISP disclosing the network policies of the entire accessible internet. It is not clear that any one ISP would have the ability, let alone incentive, to acquire the information to do so.

Third, even this complete recursive function would not accurately define ISP's service value. Consumer value depends upon the types of traffic, scaled by the treatment each class receives by the networks on which it travels. This value does not lend itself to ex ante measurement. The road analogy is helpful here: what is the value to a property owner of having ready access to the road system? This would be a function of the classes

of car traffic they would have, scaled by the speed or tolls that each class would face by all roads. This kind of inquiry requires great speculation, with costly uncertainty.⁸

To demonstrate this uncertainty, it may help to draw out symbolically the value V of access to a single user u . This requires summing the Session values of all classes (or types) of traffic C , for all destinations (or sources) of that traffic D , that u would contact. The *Session* value to u of access to that destination d using just that class of traffic c , is a complicated function of (1) the destination d , (2) the class c , and (3) the treatment the traffic receives from each network between u and the d :

$$V_u \approx \sum_{c \in C_u} \left(\sum_{d \in D_u(c)} [Session_u(d, c, \{T_n(c) \mid n \in N(u, d)\})] \right)$$

V_u	the value of access to a single user u
C_u	the set of all classes (or types) of u 's traffic
$D_u(c)$	the set of all destinations (or sources) of u 's traffic of type c
$Session_u(d, c, treatment)$	the value to u of c traffic receiving <i>treatment</i> when transmitted between u and d
$T_n(c)$	the treatment network n gives to the traffic of class c
$N(u, d)$	the set of networks along the path between u and d

And this complicated formula is woefully reductive.⁹ The *Session* function is complex and opaque. Even if we knew the class of traffic, we knew the destination, the networks along the way and the treatment each gives to the traffic; we still must wonder about the infinite possible communications that may be contained within that traffic. Are they reading Wikipedia, or editing it; watching videos of riots in Tibet¹⁰, or of a dog on a skateboard¹¹? And how does the user's perceived value change, when a network along the way treats the traffic a little differently?

Despite these difficulties, observe that, of all the unknown variables in this function, the user cannot possibly know $T_n(c)$ or $N(u, d)$ —they are the infrastructure details known only to the network owners. This keeps the user from performing this utility calculation. Similarly, the network cannot know V_u , C_u or $D_u(c)$; since they all depend on user preferences—so the network owner cannot make the calculation either. Furthermore, this whole formula ignores the timing of the choices. And it tells of a fixed set of choices, which is absurd in the context of a rapidly changing internet where new network configurations and new networked applications are introduced and removed all the time.

Even if we just take the *average* value of some class of traffic, the dimensions used to approximate this calculation (class, treatment, direction) are interdependent and resist reduction into simple terms. For example, the value to a consumer of a certain traffic class—say traffic with streaming video characteristics—greatly depends upon the content of the video (or of the traffic flow that appears to the network to be a video stream; recall, the network can only make statistical inferences about the actual content of traffic). Further, this particular video stream may be encoded in such a way to combat some prevalent network delays or disruptions but the same encoding may leave it more vulnerable to still other disruptions.

All of which suggests that networks have little incentive to provide disclosure, and it is unclear that any regulatory mandate could provide perfect, or even useful, disclosure. The strategy we employ gives those necessary bits of information—the raw instrumentalities—to those who have an incentive and capacity to tease meaning out of this raw data, *i.e.* the “internet vanguard,” a group discussed below. They will craft the necessary disclosure, tailored to the exigencies and needs that no regulation could possibly foresee.

b. Disclosure’s Audience

Many advocates of internet disclosure appear to believe that disclosure will allow for efficiency-enhancing consumer behaviors. In other words, consumers will choose ISPs with network management disclosure policies closer to their desired tastes—and those tastes would include non-discriminatory networks.¹² There are several problems with this scenario.

The primary problem is that most people do not really have a choice among ISPs.¹³ If there really is no choice, if there really is no competitive broadband market, information will not do much in making the market perform any better. People will only gain a better understanding of the extent of their ISP’s market power and their own consumer impotence.

Another problem with targeting consumers generally is the technical complexity of the disclosures. The disclosure will be of little use to the typical consumer. Like the disclosure on prescription drugs—those inserts with detailed information about the composition and chemical attributes of the accompanying drug for which you need an advanced degree in pharmacology to understand¹⁴—technical disclosures of traffic management and peering relations cannot provide guidance to the average consumer.¹⁵

c. Disclosure’s Purpose

When a producer can provide, convey, or obtain the pertinent information at a lower cost than can the consumer, mandatory disclosure may be appropriate.¹⁶ Such disclosure provides the market, in the cheapest way, the information that will allow consumers to buy things that best match their ideal preferences, resulting in economic efficiency.¹⁷ By diminishing a firm’s ability to exploit consumers’ ignorance (by

demanding higher prices or by delivering lower quality goods), mandatory disclosure increases market efficiency—or so the story goes. Not surprisingly, mandated disclosure exists in situations like food labeling, in which the consumer cannot easily ascertain an important product characteristic,¹⁸ *e.g.*, does my cupcake have trans fats?

But mandated internet disclosure (alone) does not seem likely to improve consumer surplus or efficiency. This is true primarily because consumers have so little choices in their broadband providers. It does no good to learn about the deficiencies about your provider if you cannot switch to another. Further, disclosure must be easily understood by consumers in order to aid their choices. As the above discussion shows, no such metrics exist now.

Telecommunications disclosure historically has not served to inform consumers: it has served to inform regulators, primarily for rate making functions. For instance, for almost a century, long-distance companies submitted cost information to substantiate their tariffs. The FCC knew what to do with this disclosure because rate-of-return or price cap regulation reflected coherent regulatory goals with an established implementation. Regulators had goals: set prices that give adequate returns to producers and would emerge in the absence of market power.

In contrast, it is not quite clear what the FCC would do with proposed network disclosure because there is no accepted standard as to what constitutes legitimate internet discrimination. The FCC's recent BitTorrent order demonstrates not an attempt at coherent network neutrality regulatory regime but, regrettably, a willingness to create a standard vague enough to accommodate myriad political exigencies.

In the BitTorrent order, the FCC ruled Comcast's deep-packet inspection violated in a "*prima facie*" way the FCC's two internet "principles" of allowing consumers to run applications they wish to and maximizing consumer control.¹⁹ To reach this conclusion, the FCC found that Comcast "blocked," or at least "delayed" traffic in a discriminatory fashion.²⁰ Recognizing that traffic management may justify blockage or delay, the FCC created a test to determine the acceptability of an ISP's management: whether its blocking or delaying methods are "carefully tailored to its interest in easing network congestion" and/or constitute "reasonable network management."²¹

This test, borrowed for no apparent reason, from First Amendment constitutional strict scrutiny analysis, fails to provide meaningful guidance as to acceptable traffic shaping techniques. Under traditional First Amendment strict scrutiny analysis, a court upholding a government speech restriction must find that it (i) furthers a compelling state interest and (ii) constitutes a "narrowly tailored" means of doing so.²² For instance, the Supreme Court, in upholding the Federal Election Campaign Act's restriction on federal campaign contributions, found that such restrictions were "marginal" but the state's interest in prohibiting vote buying compelling.²³

The FCC's problem in adopting this First Amendment test to traffic shaping is *all* network management and policies can limit consumers' ability to run applications and/or

minimize their control. The FCC assumes that an ISP has a menu of options to achieve a given degree of network efficiency and that it can pick the least restrictive or “narrowly tailored” option. That simply may not be the case. Rather, a given level of efficiency might require certain types of discrimination, delay, or even blockage—creating an inevitable trade-off between efficiency and user “freedom.” The FCC gives no metric, or even conceptual rule, to balance this trade-off. The only rule BitTorrent Order announced is that the most egregious practices of Comcast, coupled with their misleading public statements, are unacceptable for now.

III. What is Currently Disclosed

Those with pecuniary, personal and reputational interests already attempt to track network management techniques. A varied group of academics, professional programmers and engineers, hobbyists, and others, the aforementioned “internet vanguard” already have developed ways to estimate what the network is doing. A skilled hacker²⁴ can poke and prod at the network in an effort to estimate some of the data that we suggest disclosing. In this section we explore these limits. We review efforts to detect (i) network’s, like an ISP’s, traffic management practices to control and shape traffic *within* its network and (ii) efforts to determine the interconnection relationships between networks.

If these estimates were revealing enough, our disclosure proposal would be redundant. But, as discussed, the techniques are limited and underscore the need for a regulatory response.

a. Uncovering Internal Traffic Management

Dan Kaminsky frequently gives a talk entitled “Black Ops of TCP/IP.”²⁵ Kaminsky—a network security researcher: a hacker with a business card—discusses his latest efforts to do “bad” things to TCP/IP. In his 2006 and 2007 talks, he described two ways for a consumer to detect network neutrality. He clarified the network neutrality problem using the language of internet exceptionalism: rather than defending network neutrality, the focus should be on detecting and defending against “provider hostility.”²⁶

While his techniques are clever, providing some blurry-eyed detection of “provider hostility,” they are severely limited. The 2006 technique relies on error-prone inferences made worse by conflating factors.²⁷ And the 2007 technique has to make use of unrelated security bugs, is deeply inefficient, and still cannot detect many forms of traffic management.²⁸

The Electronic Frontier Foundation recently made the first release of its “switzerland” software.²⁹ Instead of squeezing through a security bug (as did Kaminsky’s 2007 method), switzerland lets users explicitly coordinate with a trusted third party to validate traffic. It detects any drop, forgery, or modification between any two computers, when both are running the software. It also tracks a variety of meta-data

about the quality of the connection, in hopes of someday detecting more subtle interference.³⁰

This software is still in its infancy. But it already lets consumers detect the most blatant practices by ISPs, like forging packets.³¹ Despite this accomplishment, Switzerland is unable to reliably detect more subtle (and common) techniques of ISPs-- unless and until a broad community forms to coordinate its use by aggregating and analyzing the meta-data. Even then the policies can only be roughly inferred. Like Kaminsky's first detection tool, Switzerland must make an error-prone inference in the face of conflating factors (like unintended network performance fluctuations). Switzerland also cannot be certain which network between the two connected computers is actually causing the discrimination.

b. Inferring External Interconnection

Moving away from efforts to detect the network's internal traffic management, we look now at efforts to infer how the networks are externally interconnected. The topology of the internet remains the focus of many network researchers. But the data required for this research is scarce and not widely available.³² Furthermore there are many conceptual "maps" of the internet. We are interested specifically in how separately administered networks are interconnected.

One way to map the internet is by tracing the way data is routed between IP addresses—IP topology. Most computers come with a "traceroute" program that gives a user the path of IP addresses along the way to a specified destination. The traceroute software attempts to plot out the sequence of intermediate hosts between the source and destination. It sends a series of packets addressed to the destination so that the network will send them along the path to the destination. The trick is that the packets are crafted to *not* make it all the way to the destination and to reveal the IP address of the last hop that they managed to reach. By crafting each packet to go one-hop further than the last, the traceroute tool receives the sequence of IP addresses between the source and destination.³³

The Cooperative Association for Internet Data Analysis (CAIDA) developed a tool named "skitter" to collect this traceroute data from 25 strategically placed locations around the internet.³⁴ At each of the 25 places they perform traces to as many addresses on the internet as they can (they keep a list of destination IP addresses, the goal is to hit every nook and cranny of the net).

But when traceroute reveals the sequence of IP addresses, it does not quite reveal how any networks are interconnected. Instead we have the interconnections between IP addresses—we have the IP topology—and some of these IP addresses may reside within the same network. To get network topology (the set of interconnections between separately administered *networks*) each IP address must be translated into the network on which it resides. This translation requires a mapping between IP addresses and the

network that administers that address. It's like knowing the cities when you want a map of countries—you must be able to match city (IP) to country (network).

To understand how networks are identified, we must explore inter-network protocols (that is, protocols that handle routing between not just two devices but between two separately administered networks). One problem that internet protocols are designed to address is the coordination of separately administered networks as a single utility.³⁵ When a computer sends out a packet destined for a particular IP address, how does a router know where to send it next? This routing problem is solved in two parts: internally (using “interior gateway protocols”) and externally (using “exterior gateway protocols”).

Interior gateway protocols (IGPs) handle routing traffic within a set of routers that are commonly controlled.³⁶ A network administrator has complete control over these routes since the traffic is only traversing his own network. These internal routes are not important for this interconnection discussion.³⁷

Exterior gateway protocols (EGPs) handle routing traffic among separately administered networks.³⁸ On the public internet, a separately administered network is called an “autonomous system” (AS) and each AS is assigned an identifying number (called an ASN).³⁹ The “Border Gateway Protocol” (BGP) is the EGP used on the public internet to exchange routing information among these ASes.⁴⁰ So when an AS arranges to interconnect with another AS, they use BGP to exchange routing information (the path to get to some IP address).

The BGP session between two ASes is important because it implements their interconnection agreement. When an AS arranges to interconnect with another AS, they will have to decide what traffic they will carry from each other. These interconnection agreements come in many flavors.⁴¹

To implement the agreement each AS must configure their externally facing BGP router. They will specify which routes it will accept or share in the BGP session with their neighbor. It also specifies which IP addresses are within their AS. In this way the BGP session between pairs of neighboring ASes technically defines their interconnection.

Network researchers have two ways to get information about the various BGP sessions ongoing between ASes on the public internet. The first way is to set up their own AS and interconnect with networks on the internet which provide a view of the global internet routing system from that perspective. This kind of perspective of the internet's routing system is revealed by some networks using “Looking Glass” software which lets the public issue (usually limited) sets of commands to display the internet's routing system from that network's perspective.⁴² The University of Oregon's “Route Views” project aggregates BGP data from a number of different locations (perspectives) around the world to provide a more complete view.⁴³

Another source for BGP session data is the Internet Routing Registry (IRR).⁴⁴ This is a place where an AS can voluntarily register their BGP configuration.⁴⁵ If the AS

registers their BGP configuration in the IRR, they can then set up their BGP routers to be automatically configured. Some backbone providers require this registration of their customers.⁴⁶ But others do not; and providers are free to input misleading data. This can make the IRR data inaccurate.⁴⁷

We have now seen three ways to infer the interconnection relationships between networks, and each gives a view along a different dimension: (1) using traceroute to get IP address paths and then mapping these to ASNs (a “data” view); (2) the “Route Views” method of setting up an AS and recording routing tables (a “control” view); and (3) the voluntary IRR where networks register themselves (a “management” view).⁴⁸ None of these methods provides enough reliable or useful data.⁴⁹ Unless and until there is an overhaul of the inter-network routing architecture such that the routing is exposed to the end-user⁵⁰, our disclosure proposal aims to improve the quality of data produced by these existing methods so that they become more useful.

IV. Network Transparency: A Proposal

We would require disclosure of internal traffic management and external interconnection relationships. We begin with an overview of this proposal, and later describe each part in more rigorous detail. In describing these disclosures we must beware the dangers of specificity: it risks fixing the state-of-the-art and thereby ossifying network engineering practices, forcing it into a rigid mold.

We propose a disclosure aimed at the more scattered audience that we term the “internet vanguard.” The internet vanguard are those with the technical competence and the inspiration—whatever its source—to use the disclosure. This is a varied group with shifting and often conflicting interests including hobbyists, academics, application developers and infrastructure owners. We seek to expose network practices so the knowledge will serve as another arrow in the quiver of the internet vanguard. The hope is to blunt any negative impacts of non-neutral networks without outright prohibiting them.⁵¹

The first set of disclosures reveals the network’s internal traffic management, its “QOS policy.” This shows all that happens to a user’s traffic *inside* the network. This requires disclosing the combinations of techniques and algorithms that are used to classify traffic and to actually implement a “QOS policy.”⁵²

The second disclosure reveals the ISPs interconnections. After (or before) surviving the ISP’s internal traffic treatment, this disclosure looks *outside* to show where traffic travels. There are at least two ways to require this. For reasons we discuss below, accurate and timely participation in the IRR seems the most cost-effective way to reveal the interconnections.

a. Internal QOS

The fear of ossifying network engineering is especially felt in describing the QOS disclosure requirement. The form of disclosure required must be open ended to allow flexibility in implementing these mechanisms. The hard part is balancing disclosure uniformity, on the one hand, against design flexibility on the other. We want to craft the internal disclosure requirement with enough reference to a standard (like the widely used “DiffServ architecture” we discuss below) but without stifling design flexibility. Complete flexibility, however, undermines the usefulness of the disclosure; the less uniform the disclosure, the more complex and difficult it is to develop useful tools.

QOS architecture is important because it outlines the form of this disclosure requirement. Before introducing the components of QOS architecture, it is important to carefully explain what “architecture” means in the network engineering context. When we talk about architecture, the term is not merely a creative way to say design and construction. Instead it means a “kind of abstract thinking about resources and relationships” that results in “a set of high-level design principles that guide the technological design” of the implementation.⁵³ In this way, architecture does not mechanically instruct the technical design, rather it provides a way “to ensure that the resulting technical design will be consistent and coherent—the pieces will fit together smoothly—and that the design will satisfy the requirements” of the architecture.⁵⁴ Our disclosure requirement within such a flexible framework strikes a balance between a usefully uniform disclosure, and lack of restrictions when designing the network.

There are two primary QOS architectures: integrated services (“Intserv”) and differentiated services (“Diffserv”). Intserv lets each application reserve a portion along each node of the network.⁵⁵ In contrast, Diffserv classifies the traffic and then treats it according to the class's previously-selected policy.⁵⁶

We focus on Diffserv; it is more widely deployed (in private and public networks) since it overcomes the scalability issues of Intserv.⁵⁷ The Diffserv QOS architecture describes a system where each router has a set of classes each corresponding to a treatment. A set of Diffserv routers that together implement a QOS policy are called a Diffserv “domain.” At each router, a packet is classified and then treated accordingly. Our disclosure requirement compels disclosure of the composite treatment across the domain.

i. Classification or “Trust”

Before treating the traffic, a Diffserv domain must classify each packet into one of a limited number of classes. The selected class is important because it dictates how the traffic will be treated. The classification is done as near as possible to the edge of the network (near the source of the traffic), to simplify the network design.⁵⁸ Subsequent hops along the way need only respond to the class as marked on the packet, rather than keep track of each connection.

There are generally two ways to classify and mark traffic. “Source marking” is when the end-system (the computer that generates the traffic) marks its own traffic as belonging to some class.⁵⁹ “Ingress marking” is when a node (a router) close to the source device uses a classification technique to mark the traffic.⁶⁰

The choice between marking at the source or upon ingress brings us another vocabulary gem: “trust.” When the Diffserv domain will rely on the classifying marks made by a device, the device is “trusted.”⁶¹ This trust may be limited: some kinds of classification mechanisms (sometimes called “color-aware”) perform their own complex classification and only incorporate the incoming marking as one input to a new classification.

This “trust” has not been as misused as some other network engineering terms. But it is easy to see how it might be: “My ISP does not trust me to say how important my traffic is!” While this trust is fairly close to its literal meaning, it is important to understand exactly what it means: it describes the reliability of the pre-existing traffic classifications. If a network disclosure indicates that the network does not “trust” its users, this should be understood in its precise technical meaning.

ii. Routing or “Fairness”

After classifying the traffic (or trusting the pre-existing marking) the router then schedules their departure. This scheduling is where the QOS policy will be translated into algorithms to actually control the traffic. When a router's scheduling algorithms faithfully implement a QOS policy, it is considered “fair.”

This “fairness” is particularly vulnerable to misunderstanding; when interpreting a network’s disclosure, we must keep in mind its precise meaning. It is generally used simply as a metric when evaluating how well a scheduler achieves a desired bandwidth allocation.⁶² This sense of “fairness” is closer to “in conformity with the rules” rather than “even-handed.” While there has been some effort to alter the focus of the technical research community, the focus has remained on flow-rate fairness: each flow *of a given class* should be treated equally.⁶³

This flow-fairness is surprising in two ways. First, “fairness” so-defined takes the classification policy as given. To an outsider, saying something “conforms with the rules” just begs the question: what rules? So this sense of fairness is very much *not* a proxy for any broader equitable sense of the term, since the classifying criteria and their associated treatment—the rules—are not evaluated for their normative desirability. The second, related surprise is this “fairness” does not look at whose computer is generating the flow. If one user's application generates 10,000 flows of a certain class, while a second user generates only 1, each of the 10,001 flows are seen as formally equal in the eyes of this “fairness.” This highlights a difference of perspective.

From the ISP perspective, it is a problem when “fairness” ignores the identity of source—they must concern themselves with allocating network access among their consumers. But this “fair” allocation among all consumers is not the concern of each individual consumer. Indeed this is the fundamental joint-resource problem faced by the ISP.⁶⁴

This stilted notion of fairness causes added confusion when raised in the wider public debates over “fair” use of bandwidth.⁶⁵ It conflates flow-fairness with the notion of “fairness” in the context of the economic free-rider problem that an ISP faces. Free-riding refers to when a person uses an “unfair” amount of a common good, whether by overuse or underpayment. Furthermore, the “fairness” reported by network engineering risks being misunderstood as the ever-controversial “fair share” that I am obligated to do if I benefit from others cooperative action.⁶⁶

The technical idea captured by the “fairness” is important—it describes how effectively the QOS policy is implemented. But since it does not really touch on the human actors and does no evaluation of the classifying criteria or QOS policy itself, it should not be understood to make the moral claim that “fair” might normally suggest.

To effectively implement a QOS policy, network engineers can select from a variety of algorithms to piece together the desired treatment. These algorithms are sometimes grouped under labels: queuing, scheduling, shaping, dropping, etc. These labels are not equivalent (e.g. a dropped packet is never sent, whereas a "shaped" packet might be sent eventually), but they are all ways of describing the treatment that packets receive after they are classified. The fairness of a QOS implementation reflects the effectiveness of this arrangement of algorithms to implement the desired policy. But beware when analyzing a network’s disclosed QOS policy; this sense of fairness is descriptive and should not be misunderstood as giving normative guidance.

b. External Interconnection

In the preceding paragraphs we discussed how to interpret a disclosed QOS policy; the QOS policy is effected within the network . We move now to the border of a network, to discuss the disclosure requirement at the boundary of administrative control: the interconnections between separately administered networks.

As we discussed earlier, AS interconnections form the topology of the public internet. The three methods we discussed for acquiring this information (traceroute, “Route Views”, and the IRR database) reinforce each other, so if our disclosure can improve one then the effect is amplified across all three. We compare two ways to compel disclosure of this interconnection information to take advantage of these methods.

The first is to require network’s accurate and timely participation in the distributed IRR database. Some networks already maintain IRR entries. But, as discussed previously, this data is often stale and often inaccurate. The second way to compel the disclosure is to require networks to enable real-time (programmatic) access to

their “Looking Glass” (LG) portal and to expose it to the public. As discussed previously, many networks already have an LG portal setup but they often limit the access and prohibit automated access.⁶⁷

There are at least three reasons to prefer the first (accurate and timely participation in IRR) over the second (LG portal). First, the IRR provides all of the routing data in a single place making it easier to use. A person trying to use the data does not need to locate all the different network LG portals. Second, the IRR database is already set up to be distributed and mirrored, making it more reliable. This ensures that the data can be accessed, without relying on each network to keep their LG portal responsive. Third, maintaining an IRR entry is likely to be less costly for the network than providing an LG portal. A network’s LG portal could be overwhelmed by queries whereas the IRR is duplicated and mirrored by third parties to ensure capacity—a heavy user could even create their own copy. Also the IRR need only be updated when the network adds a new interconnection. The network will already have to reconfigure their routers to add the new interconnection, so requiring entry in the IRR adds only minimal extra work since the IRR is set up to enable this reconfiguration to be automated by registry updates.

Any interconnection disclosure will be amplified since all three methods of inferring interconnection reinforce and check each other. Compelling networks to maintain accurate entries in the IRR seems the most cost-effective way to reveal the technical details of their interconnections.

Finally, and perhaps most controversially, we would require the filing of all interconnection agreements of networks providing either internet access or transmission to the public. The FCC could play a role in collecting and making public these documents. While such publication would undoubtedly hurt the bargaining of certain access providers, it is not clear that there would be any welfare loss. Further, such disclosure would provide a powerful curb against strategic behavior.

V. Conclusion

Efforts to regulate the internet, like the BitTorrent order, have proceeded tentatively in mandating any internet policy. Our proposed disclosure advances this work with a careful response to potential anticompetitive abuses on the net. It harnesses the generative internet to provide a new type of regulatory disclosure—one that is aimed neither directly at consumers nor at regulators—but instead seeks to provide the raw instrumentalities to those who can aggregate, transform, and interpret the data needed to guide any regulatory policy or consumer choice. Before we can even say what network neutrality is, we first must work towards network transparency.

¹ Open network architecture constituted the FCC's effort to make the Bell Operating Company's telephone networks more compatible with new, non-BOC data-based enhanced services. See Filing and Review of Open Network Architecture Plans, 4 F.C.C.R. 1 (1988) (setting forth the requirements that the Bell Operating Companies submit ONA plans).

² 47 C.F.R. § 68 (2008) provides the technical and procedural standards under which customer premises equipment (CPE), *i.e.*, telephones, can be connected directly to the United States public switched telephone network (PSTN). An industry group subsequently has assumed the role of setting standards. See 2000 Biennial Regulatory Review of Part 68 of the Commission's Rules and Regulations, Report and Order, 15 F.C.C.R. 24944 (2000).

³ See 47 U.S.C. § 251 et seq. (setting forth unbundling requirements for incumbent telephone companies to encourage entrance into the local telephone market).

⁴ See Kevin D. Werbach, *Only Connect* (2007), available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=%20964991.

⁵ See, *e.g.*, Phil Weiser, *The Next Frontier for Network Neutrality*, 60 ADMIN. L. REV. 273 (2008). Interestingly, the FCC in the recent BitTorrent states that it will decline from instituting rule-making because it feared "unduly tying our hands should the known facts change." Formal Complaint of Free Press and Public Knowledge Against Comcast Corporation for Secretly Degrading Peer-to-Peer Applications, Memorandum Opinion and Order, 23 F.C.C.R. 13028 ¶ 30 (2008) ("BitTorrent Order"). At the same time, however, the FCC declined to impose any systemic disclosure requirements to cure its acknowledged ignorance. The FCC curses the darkness but has no interest in lighting a candle.

⁶ See A. Michael Froomkin, *Habermas@discourse.net: Toward a Critical Theory of Cyberspace*, 116 HARV. L. REV. 749 (2003).

⁷ Note to reader: we find this locution accurate but somewhat ridiculous and ask for any suggestions for different terms.

⁸ For an attempt to fix a value for this uncertainty, see Mark Gaynor & Scott Bradner, *Statistical Framework to Value Network Neutrality*, 17 MEDIA L. & POL'Y 24 (2007) ("[W]hen market uncertainty is large it decreased the total value of a non-neutral network by 50 percent because the transport provider has picked such an average bundle of services that half of the potential users don't value it as high as its price.").

⁹ Some have made wildly speculative attempts to flesh out the utility function sketched above. For example, Hermalin and Katz put forth a model to try and capture "a household's utility from connecting to the [ISP] platform and consuming the available applications." Benjamin E. Hermalin & Michael L. Katz, *The Economics of Product-Line Restrictions With an Application to the Network Neutrality Debate*, at 5 (Inst. of Bus. & Econ. Research Competition Policy Center, Paper CPC06-059, 2006), at <http://repositories.cdlib.org/iber/cpc/CPC06-059>.

Such economic models by necessity make simplifying assumptions in hopes of revealing some underlying phenomenon. We do not weigh in on whether a model should be judged by the reality of its assumptions or instead on its predictive power. But we can still examine assumptions of their model to see what exactly the model aims to predict and what it deems safe to suppress.

Their model assumes that "there is diminishing marginal utility from consuming the content of any specific provider." This is suspect in fact and troubling in its predictive usefulness. Social networks are clear examples of communities where people seem to reap *increasing* "utility" as they "consume" more of the "product." And if we assume that they have diminishing returns, the model will suppress the value of human relationships by suggesting they become ever *less* important as they are lived.

There are a number of other troubling explicit and implicit assumptions in their model. But for now we focus on the most sweeping: the model assumes "[t]here is a unit mass of households with identical preferences." Without any qualifications, this is clearly not accurate (people like different things online); but, more damning, it suppresses the very defining characteristic of the long-tail of the internet. If we want an economic model of a user's utility from accessing internet, it will have to capture this world of niches.

Here the demanding challenge for economists is preserving their disciplinary assumption of exogenous uniform preferences, while adequately explaining the felt reality of heterogeneity that defines the internet. See generally, George J. Stigler & Gary S. Becker, *De Gustibus Non Est Disputandum*, 67 AM. ECON. REV. 76, 89-90 (1977) ("Our hypothesis is trivial, for it merely asserts that we should apply

standard economic logic as extensively as possible. But the self-same hypothesis is also a demanding challenge, for it urges us not to abandon opaque and complicated problems with the easy suggestion that the further explanation will perhaps someday be produced by one of our sister behavioral sciences.”) We do not aim to discourage modeling efforts by economists; rather, we hope they can find a way to creatively explain this seemingly “opaque and complicated problem.”

¹⁰ *Unexpected Holidays: Tourist Films Tibet Riots*, <http://www.youtube.com/watch?v=NflbbNHpP5A> (last accessed July 28, 2009).

¹¹ *Skateboarding dog*, <http://video.google.com/videoplay?docid=-8762816342722081405> (last accessed July 28, 2009).

¹² See Testimony of Barbara van Schewick, Official Testimony at Second FCC En Banc Hearing, at 3 (April 17, 2008) (“[D]isclosed information must provide enough detail to enable customers to make an informed decision and to enable them to adjust their behavior. Comcast’s current acceptable use policy falls short of these goals.”).

¹³ See Indus. Analysis & Tech. Div., Wireline Competition Bureau, *High-Speed Services For Internet Access: Status As Of June 30, 2006*, at 3 (2007), at http://hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-270128A1.pdf.

¹⁴ Or a consumer may forgo a pharmacology degree and simply search the internet. See e.g., WebMD—Better Information, Better Health – Drug Index, <http://www.webmd.com/drugs/index-drugs.aspx> (last accessed Sept. 7, 2008) (“Search for prescription and over-the-counter medications by brand or generic name, or by condition.”).

Medical advice online is controversial in part because of gravity of the subject and rent-protecting behavior of medical professionals. But this sourcing of experts to digest the disclosures into a widely useful form mirrors part of the mechanism we envision in our disclosure proposal, albeit in a much different field.

¹⁵ The technical terms of network management policy will be incomprehensible to consumers who sign up on the phone for internet access. C.f. Douglas A. Hass, *The Never-Was-Neutral Net and Why Informed End Users Can End the Net Neutrality Debates*, 22 BERKELEY TECH. L.J. 1565, 1630-31 (2007) (describing how disclosing traffic control policies directly to consumers, similar to the Fair Credit Act, “would help consumers more easily compare different service offerings”).

And the ISP’s peering relationships are not meaningful without reference to the topography of the rest of the internet; an ISP cannot be expected to know this topography let alone to disclose it to the consumer over the phone. Instead, there must be a bridge to close the gap between users who are positioned to signal their preferences and the meaning of the technical information that an ISP discloses.

¹⁶ RICHARD POSNER, *ECONOMIC ANALYSIS OF LAW* § 4.6 at 112 (6th Ed. 2003).

¹⁷ *Id.* (“A monopolist (or cartel) may have a greater incentive than a firm in a competitive industry to misrepresent the qualities of its product”).

¹⁸ *Id.* at 113.

¹⁹ “BitTorrent Order,” *supra* note 5, at ¶ 44.

²⁰ *Id.* ¶ 45.

²¹ *Id.* ¶¶ 44-8.

²² JERRY KANG, *COMMUNICATIONS LAW & POLICY* 225 (3d ed. 2009).

²³ *Buckley v. Valeo*, 424 U.S. 1, 20-2 (1976).

²⁴ The term “hacker” in this context refers to the tinkering sort. See *Internet Users’ Glossary*, (Gary Scott Malkin & Tracy LaQuey Parker eds.) (Network Working Group, Request for Comments No. 1392) at 21 (Jan., 1993), at <http://www.apps.ietf.org/rfc/rfc1392.html#page-21> (defining “hacker” as “[a] person who delights in having an intimate understanding of the internal workings of a system, computers and computer networks in particular”).

²⁵ See generally *DoxPara Research*, <http://www.doxpara.com/> (last accessed July 29, 2009) (listing previous *Black Ops* presentations along the right-hand side).

²⁶ Dan Kaminsky, *Black Ops of TCP/IP 2007* at 32:06 available at <http://video.google.com/videoplay?docid=3470502418262982787> (explaining that “[i]t’s not about defending network neutrality; network neutrality is the fricken status quo” instead “[w]hat’s actually suggested is ‘provider hostility’”).

He describes one kind of provider hostility, the “Times Square Effect”: like the altered ads in a movie depicting Times Square in New York City, the ISP alters a webpage to display its own

advertisements. An ISP, he explains, may alter search results to insert a link atop the results, or to replace web contents with its own. *Id.*

²⁷ In his 2006 talk, he introduced a method for detecting when an ISP was treating some traffic differently than others. His method simulates traffic patterns and checks whether one receives better (web surfing) or worse (P2P file sharing) treatment from the network. See Dan Kaminsky, *Black Ops of TCP/IP 2006*, available at <http://video.google.com/videoplay?docid=2235822817954898588>.

This method is highly error-prone. The user must guess at the classes of traffic being mistreated, then must fool the ISP by simulating it realistically. Also, measuring the relative treatment that the ISP is giving each class of traffic introduces large margins of error because of conflating factors (like unintentional network performance fluctuations). Finally, as Kaminsky mentioned in his talk the following year, the tool has the effect of flooding the network. Dan Kaminsky, *Black Ops of TCP/IP 2007* available at <http://video.google.com/videoplay?docid=3470502418262982787> (“Right, like I want to be the guy who tests the network by breaking it.”). This shows the delicate position of a principled internet exceptionalist: he wants to win the war without destroying the countryside.

²⁸ The 2007 technique is a kind of benign exploit that makes use of an unrelated security bug. Roughly speaking, his technique involves use of a trusted third-party to verify the contents of ongoing connections, to detect when an ISP tries to make this sort of alteration. Dan Kaminsky, *Black Ops of TCP/IP 2007*, at <http://video.google.com/videoplay?docid=3470502418262982787>. He explains that this technique is very messy (since it relies on bugs in another protocol) and that it is not a long term solution. He predicts that complete network encryption, while costly and inefficient, may be inevitable; the ad-supported model of many websites will demand it.

²⁹ *Switzerland Network Testing Tool*, <http://www.eff.org/testyourisp/switzerland> (last accessed Aug. 23, 2007).

³⁰ Peter Eckersley, *Switzerland Design* (May 2008), available at <http://switzerland.svn.sourceforge.net/viewvc/switzerland/trunk/doc/design.pdf>.

³¹ This forging of packets is what Comcast was caught doing, so switzerland can now detect that technique.

³² See Dmitri Krioukov, Fan Chung, kc claffy, Marina Fomenkov, Alessandro Vespignani & Walter Willinger, *The Workshop on Internet Topology (WIT) Report*, ACM SIGCOMM COMPUTER COMM. REV., Jan., 2007 at 69.

³³ Each IP packet has a “time to live” value (TTL) which indicates the number of hops the packet should “survive”—after that many hops, the packet should be dropped by the network. This works by having each hop decrement the TTL by 1. When the packet is dropped, the last hop it reached sends a packet to the original sender indicating that the packet did not reach its destination. This notice (from the intermediate host along the path between source and destination) includes the address of the intermediate host. So the traceroute tool crafts each packet to go just one-hop further than the last (i.e. has a TTL of 1 more than the last), recording the address of the intermediate hosts along the way.

³⁴ See Priya Mahadevan et al., *The Internet AS-Level Topology: Three Data Sources and One Definitive Metric*, ACM SIGCOMM COMPUTER COMM. REV., Jan., 2006 at 17, 18 (describing how CAIDA’s “skitter” collects traceroute data).

³⁵ David D. Clark, *The Design Philosophy of the DARPA Internet Protocols*, ACM SIGCOMM COMPUTER COMM. REV., Aug., 1988 at 106 (describing “some of the early reasoning which shaped the internet protocols” to explain “why the protocol is as it is”).

³⁶ Two prominent examples of IGPs are IS-IS and OSPF. See generally Richard P. Colella, Ross Callon, Ella P. Gardner & Yakov Rekhter, *Guidelines for OSI NSAP Allocation in the Internet* (Network Working Group, Request for Comments No. 1629) 8-12 (May 1994), available at <http://www.apps.ietf.org/rfc/rfc1629.html#page-8> (describing IS-IS); and John Moy, *OSPF Version 2* (Network Working Group, Request for Comments No. 2328) 6-7 (April 1998), available at <http://www.apps.ietf.org/rfc/rfc2328.html#page-6> (describing OSPF as an IGP which “means that it distributes routing information between routers belonging to a single Autonomous System”).

³⁷ Note that while the internal routes are not important here, the internal traffic management *is* important but is not part of this external interconnection discussion.

³⁸ We refer here to the general category of EGPs, not to the old “EGP” (itself an instance of an EGP) that BGP (another, more modern EGP) builds upon. See Yakov Rekhter, Tony Li, Susan Hares, *A Border Gateway Protocol 4 (BGP-4)* (Network Working Group, Request for Comments No. 4271) at 7 (Jan. 2006), available at <http://www.apps.ietf.org/rfc/rfc4271.html> (“[BGP] is built on experience gained with EGP.”).

³⁹ The number space for ASNs is scarce. Currently, the ASN numbering system is transitioning from the original 2-byte sized ASN (with a maximum of around 65,000 ASNs) to a 4-byte sized ASN (raising the count to around 4-billion). Still, these numbers must be assigned.

The Internet Corporation for Assigned Names and Numbers (ICANN) develops the broad policy for issuing AS numbers. The Internet Assigned Numbers Authority (IANA) then implements this policy by allocating blocks of ASNs to Regional Internet Registries (RIRs). A qualifying network can then register for an ASN from their RIR. See *Global Policy for Allocation of ASN Blocks to Regional Internet Registries* (July 31, 2008), available at <http://www.icann.org/en/general/global-policy-asn-blocks-31jul08-en.htm> (describing how the IANA must allocate ASN blocks to RIRs who in turn assign these ASNs to organizations); see also, *ARIN Number Resource Policy Manual* § 5 (Aug. 5, 2008), available at <http://www.arin.net/policy/nrpm.html#five> (describing one RIR's policy for assigning ASNs to organizations); and *ASO - RIR Links*, available at <http://aso.icann.org/rirs/index.html> (listing by ICANN of all RIRs and their associated regions); and John Hawkinson & Tony Bates, *Guidelines for creation, selection, and registration of an Autonomous System (AS)* (Network Working Group, Request for Comments No. 1930) (Mar. 1996), available at <http://www.apps.ietf.org/rfc/rfc1930.html> (describing scenarios when a network should, or should not try to acquire an ASN).

⁴⁰ See Yakov Rekhter, Tony Li, Susan Hares, *A Border Gateway Protocol 4 (BGP-4)* (Network Working Group, Request for Comments No. 4271) (Jan. 2006), available at <http://www.apps.ietf.org/rfc/rfc4271.html> (specifying the BGP standard).

⁴¹ One way to categorize these agreements is as either “peering” or “transit.” While the distinction is not always clear, peering generally refers to a more balanced exchange (carrying each others traffic as “peers”) whereas transit involves one AS paying another to carry traffic (carrying the traffic as “provider”).

It is sometimes conceptually helpful to make a strong distinction between these two sorts of arrangements, but from a technical perspective the amount of money changing hands is not important. What is important is the way traffic is actually routed. This implementation is reflected in the reachable addresses announced by an AS to his BGP neighbors.

⁴² See generally *Looking Glass Sites*, <http://www.nanog.org/> (listing various “looking glass” resources).

⁴³ University of Oregon Route Views Project, <http://www.routeviews.org/>

⁴⁴ *IRR Overview*, <http://www.irr.net/docs/overview.html> (“The Internet Routing Registry (IRR) is a distributed routing database development effort.”).

⁴⁵ Within the IRR, these routing policies are described using the Routing Policy Specification Language (RPSL). See David Meyer et al., *Using RPSL in Practice* (Network Working Group, Request for Comments No. 2650) (Aug. 1999), available at <http://www.apps.ietf.org/rfc/rfc2650.html> (explaining in practice how to specify, register, and analyze routing policies using RPSL).

⁴⁶ *Id.* at 1 (“CW, ANS and CA*Net customers are generally required to register their policies in their provider's registry. Others may register policies either at the RIPE or RADB registry, as preferred.”).

⁴⁷ See Priya Mahadevan et al., *supra* note **Error! Bookmark not defined.**, at 19 (discussing both intentional and unintentional inaccuracies in the IRR).

⁴⁸ *Id.* at 25 (“It remains an open question which data source most closely matches actual Internet AS topology, given that each graph approximates a different view of the Internet looking at the *data* (skitter), *control* (BGP), and *management* (WHOIS) planes.”) (emphasis added). Note that the “WHOIS” tool is used to access the IRR.

⁴⁹ “In its current state, Internet topology research is not an informed discipline since available data is not only scarce, but also severely limited by technical, legal, and social constraints on its collection and distribution.” See, Dmitri Krioukov, Fan Chung, kc claffy, Marina Fomenkov, Alessandro Vespignani, Walter Willinger, *The Workshop on Internet Topology (WIT) Report*, 37 ACM SIGCOMM Computer Communications Review 69 (2007), available at <http://www.caida.org/publications/papers/2006/wit/wit.pdf>.

⁵⁰ One particularly interesting proposal—“A New Inter-Domain Routing Architecture” (NIRA)—to overhaul the interconnection architecture was proposed in 2004. Xiaowei Yang et al., *NIRA: A New Inter-Domain Routing Architecture*, 15 IEEE/ACM TRANSACTIONS ON NETWORKING 775 (2007) (describing a new architecture “that gives a user the ability to choose the sequence of providers his packets take”).

NIRA lets end users control the path their packets travel. This puts the end-user in the position to transact more directly with the backbone providers. This not only aligns the industry structure along user's preferences, but it also makes end-to-end QOS more feasible. If something like NIRA is adopted it would

remove the conceptual distinction we make between *internal* traffic management and *external* interconnection—the traffic management becomes part of the end-user’s interconnection selection.

For more on NIRA, Xiaowei Yang’s dissertation goes into even greater detail. See Xiaowei Yang et al., *NIRA: A New Internet Routing Architecture* (CSAIL Technical Report, 2004) available at <http://dspace.mit.edu/handle/1721.1/30495>; see also *User Selected Routes*, http://nds.ics.uci.edu/dwiki/user_selected_routes (discussing ongoing research).

⁵¹ See generally Jonathan L. Zittrain, *The Future of the Internet and How to Stop It* (2008).

The P4P effort attacks the problem from a similar approach (exposing infrastructure to application). See <http://openp4p.net>. But P4P targets the narrow realm of P2P file sharing and focuses mostly on performance increases (P2P applications perform better when they know what kind of network they are running on: they can pick nearby peers). This only indirectly addresses NN concerns since it aims only at reducing network load from P2P traffic (the decreased load would theoretically diminish the incentive to discriminate). It does not address the broader structural concerns about human communications or about content development and control.

Also the P4P solution (ISPs sharing their topology to P2P applications) rests precariously on the accuracy of the ISP’s disclosure. Without some teeth to force ISPs to accurately disclose, this leaves the applications vulnerable to ISPs incentives to misrepresent in order to minimize their costs (to the detriment of application performance). Still, the P4P project is—in theory—a creative and interesting attempt to cooperate; hopefully it can resolve these issues.

⁵² We explore this terminology below.

⁵³ Robert Braden, David Clark, Scott Shenker & John Wroclawski, *Developing a Next-Generation Internet Architecture* 3-4, (ISI White Paper, 2000), available at: <http://www.isi.edu/newarch/DOCUMENTS/WhitePaper.pdf>.

⁵⁴ *Id.* In network engineering, the architecture is to the implementation much the way the principle is to the rule in law. The implementation flows from the architecture, and changes to the implementation must still fit within the general principles of the architecture.

⁵⁵ Bob Braden, David Clark, & Scott Shenker, *Integrated Services in the Internet Architecture: an Overview*, (Network Working Group, Request for Comments No. 1633) (June 1994), available at <http://www.apps.ietf.org/rfc/rfc1633.html> (discussing the IntServ architecture).

⁵⁶ Steven Blake et al., *An Architecture for Differentiated Services*, (Network Working Group, Request for Comments No. 2475) (Dec. 1998), available at <http://www.apps.ietf.org/rfc/rfc2475.html> (discussing the DiffServ architecture).

⁵⁷ Intserv is unworkable (and thus unused) on the public internet because it is expensive at large scale and it requires end-to-end coordination. JOHN EVANS & CLARENCE FILSFILS, *DEPLOYING IP AND MPLS QOS FOR MULTISERVICE NETWORKS: THEORY AND PRACTICE* [hereinafter “*Deploying QOS*”] 149-50 (2007).

The Intserv QOS architecture describes a system for each application to reserve network resources. An application running on a user’s computer wants a certain QOS from the network (that is, the application wants its traffic to be given a certain treatment by the network). The request (including the class of quality required) is forwarded on to each subsequent node in the network along the path to the destination, with each node confirming the request if it can. In this way, the path is guaranteed to provide the quality required.

The problem is a router has to keep track of all ongoing traffic to avoid confirming too many requests. And this has to be done in each router on the network. In this way, Intserv is both expensive to deploy in a large network and further complicated by the diversity of owners of the routers on the public internet. A Diffserv domain scales better than Intserv because it does all the classifying of traffic at the edges of the network, and only keeps track of each class, rather than tracking each flow of traffic. *Id.*

⁵⁸ *Deploying QOS*, p. 99 (“as close to the traffic source as possible in order to simplify the network design”).

⁵⁹ *Deploying QOS*, p. 99 (“Packet marking may be applied at the source end-system itself; if the end-system is considered to be *trusted* then this marking may be relied upon throughout . . .”) (emphasis added).

⁶⁰ *Deploying QOS*, p. 100 (describing how a trusted device close to the end-systems will perform implicit or complex classification when the end-systems are “not capable of” or “cannot be trusted to” mark their own traffic).

⁶¹ *Deploying QOS*, p.99.

⁶² *Deploying QOS*, p. 118 (“how closely the scheduler achieves the intended bandwidth allocation”); repeated at p. 123.

⁶³ Bob Briscoe, *Flow Rate Fairness: Dismantling a Religion*, ACM SIGCOMM COMPUTER COMM. REV., April, 2007, at 65 (arguing that current research on “fairness” continues to fail because it focuses on the wrong thing—flows; instead it should focus on the cost of user’s actions on each other).

⁶⁴ From a strategic standpoint, the rhetoric of “fairness” puts ISPs in a tricky position. They will want to support it if it prompts development of a social norm for customers to be well-behaved. Then their customers have solved the ISP’s joint-resource problem. But social norms are not so easily controlled as contractual terms. The prevailing “fairness” norm today may treat high-bandwidth P2P applications as “unfair” (perhaps because of presumed copyright infringement), but tomorrow some new high-bandwidth application might not have the same social stigma. This would force ISPs to instead fight against “fairness.”

⁶⁵ E.g., Ars Technica, <http://arstechnica.com/news.ars/post/20071205-growth-of-p2p-leads-ietf-to-debate-fair-bandwidth-use.html>.

⁶⁶ JOHN RAWLS, A THEORY OF JUSTICE: ORIGINAL EDITION at 112 (Reprinted 1st Ed. 2005) (“We are not to gain from the cooperative labors of others without doing our fair share.”) (citing H. L. A. Hart, *Are There Any Natural Rights?* 64 PHILOSOPHICAL REVIEW 185 (1955); see also *The Free Rider Problem*, STANFORD ENCYCLOPEDIA OF PHILOSOPHY, <http://plato.stanford.edu/entries/free-rider/>). When a journalist describes how P2P applications do not behave “fairly” it implies that they are taking advantage of some unstated generosity in violation of a broader “fair” use of the network. See e.g., Ars Technica, *supra* note 65. “Fairness” in this context conceptually glosses over the distinction between a product purchase and a group asset.

A customer should not blame their neighbor for defects in a service they purchase. Their neighbor’s use might be a but-for cause of their unsatisfactory internet access, but it is the ISPs responsibility (not their neighbor’s) to ensure internet access. This mindset treats the purchased product as if it were part of a collective group (like corporate network users), but each consumer purchases internet access from the ISP separately. A customer has no obligations to the other customers to preserve network resources. Consumers are not being “unfair” when they use what they have purchased.

⁶⁷ E.g., *Sprint Looking Glass*, <https://www.sprint.net/lg/> (“Sprint’s Looking Glass is NOT to be used with ANY automated scripts/applications unless expressly authorized by Sprint.”) (last accessed Aug. 29, 2008).