How to determine whether a traffic management practice is reasonable

Scott Jordan

Department of Computer Science, University of California, Irvine

Arijit Ghosh

Department of Computer Science, University of California, Irvine

Abstract

As part of the wider debate over net neutrality, traffic management practices of Internet Service Providers have become an issue of public concern. An increasing number of vendors offer network equipment to implement a variety of traffic management practices using packet classification and packet filtering. The Federal Communications Commission has asked for public input on whether such traffic management practices are reasonable forms of network management. Little attention has been paid to this issue within the academic computer science community, and many communications policy researchers have recommended a case-by-case analysis. In contrast, in this paper we propose a framework for the classification of traffic management practices as reasonable or unreasonable. To build the framework, we focus both on the technical aspects of traffic management techniques and on the goals and practices of an ISP that uses these techniques. The framework classifies traffic management practices as reasonable or unreasonable on the basis of the *technique* used and on the basis of *how and when* the techniques are applied. We suggest that whether a traffic management practice is reasonable largely rests on the answers to four questions regarding the techniques and practices used. We consider examples of how these techniques are used by ISPs, and how the answers to these four questions collectively affect the degree to which a traffic management practice is reasonable. Based on these questions, we propose a framework that classifies techniques as unreasonable if they are unreasonably anti-competitive, cause undue harm to consumers, or unreasonably impair free speech.

Email addresses: `sjordan@uci.edu` (Scott Jordan), `Arijit.Gjosh@uci.edu` (Arijit Ghosh).

1 Introduction

The traffic management practices of Internet Service Providers (ISPs) have become an issue of public debate. In 2007, Comcast started using reset packets to terminate selected peer-to-peer connections (Comcast Corporation, 2008). This practice, when uncovered by a few users, generated a firestorm of debate, largely because it dovetailed into an existing debate over net neutrality (Weitzner, 2008).

Net neutrality represents the idea that Internet users are entitled to service that does not discriminate on the basis of source, destination, or ownership of Internet traffic. Proponents of net neutrality argue that without a prohibition on discrimination, ISPs may charge application providers discriminatory prices for access to dedicated bandwidth or for quality of service (QoS), or may outright block access to certain applications or websites, and that such activity will inhibit development of new Internet applications (Jordan, 2009). To proponents of net neutrality, Comcast's practices seemed like blocking of certain applications; to Comcast, however, its practices seemed like reasonable traffic management designed to limit network congestion.

The debate centers not only on Comcast's practices, but also on the wider use of deep packet inspection techniques which allow ISPs to identify and control traffic streams on the basis of transport and application layer information. An increasing number of vendors offer equipment that can be placed in the network to implement a variety of traffic management practices using packet classification and packet filtering.

In response to the early net neutrality debate in the United States, in 2005 the Federal Communications Commission (FCC) issued a set of principles (FCC, 2005b). The principles express the sentiment that consumers should be entitled to connect devices and to access content and applications of their choice. In a footnote, the FCC comments that these principles are subject to "reasonable network management", but does not define what this terms means. In response to the discovery of Comcast's traffic management practices, a few organizations petitioned the FCC to rule that an ISP is violating these principles (and thus not practicing reasonable network management) when it intentionally degrades a targeted Internet application such as peer-to-peer (Free Press, Public Knowledge et. al., 2007) and to adopt rules that would prevent such practices (Vuze Inc., 2007).

The FCC asked for public input on whether this practice and other traffic management practices are reasonable forms of network management (FCC, 2007). They asked whether ISPs use traffic management practices to prioritize latency-sensitive applications, to block unwanted traffic, to implement parental controls, to improve network performance, or to gain advantage over competitors. They also asked whether these practices are helpful or harmful to consumers and whether they are reasonable. In 2008, the FCC concluded that Comcast violated a principle concerning users' rights to access lawful Internet content and use applications of their choice, and that its practices do not constitute reasonable network management (FCC, 2008). However, the FCC has not adopted rules that would delineate reasonable network management.

These questions have largely gone unanswered by the academic networking community. Most networking technologists would have some concern about violations of layering such as that involved in deep packet inspection. However, there is no consensus about when layering violations are warranted or how to respond to them.

There have been only a few attempts in the networking literature to go beyond the technical aspects of traffic management and to consider the social and legal implications. Weitzner (2008) discusses the Comcast incident and the connections to net neutrality. Peha (2007) discusses the incentives that ISPs may have for using discriminatory practices, and the benefits and damages that may accrue from these practices. He also gives examples of what should be allowed and prohibited, but does not give a framework that allows one to classify practices. Frieden (2006) similarly gives examples of what he believes to be permissible and impermissible traffic management practices, and suggests a few best practices (including limitations on blocking and degradation) that ISPs should adopt. He similarly does not present a framework for classification, but instead proposes that the FCC should impose reporting requirements on ISPs and assess practices on a case by case basis. Lehr et al. (2007) discuss strategies that end-users may adopt in response to ISP discrimination, including technical counter-measures.

However, we have found no literature that proposes a method for classification of traffic management practices as reasonable or unreasonable. In this paper, we present such a framework for traffic management by Internet Service Providers within the United States. We restrict our attention to traffic management policies as a subset of a larger class of network management policies. To build the framework, we focus both on the technical aspects of traffic management techniques and on the goals and practices of an ISP that uses these techniques. The framework classifies traffic management practices as reasonable or unreasonable on the basis of the *technique* used and on the basis of *who* decides when the techniques are applied. The framework results in classifying practices as unreasonable when they are unreasonably anti-competitive, cause undue harm to consumers, or unreasonably impair free speech.

The paper proceeds as follows. In section 2, we suggest that whether a traffic management practice is reasonable largely rests on the answers to four questions regarding the techniques and practices used. Section 3 considers examples of how these techniques are used by ISPs, and how the answers to these four questions collectively affect the degree to which a traffic management practice is reasonable. Based on these questions, in section 4 we propose a framework that classifies techniques as unreasonable if they are unreasonably anti-competitive, cause undue harm to consumers, or unreasonably impair free speech.

2 Key questions about traffic management techniques and practices

Traffic management is applied to implement a variety of functions, at a variety of layers, by a variety of actors, in a variety of manners, for a variety of purposes. To delineate these components, define a *traffic management technique* as a specific function that is offered at a specific layer. The function should determine whether traffic is transmitted and/or the rate at which traffic is transmitted, or should enable such functions in other techniques. Define a *traffic management practice* as a collection of traffic management techniques, used by a specific type of actor, in a specific manner, for a specific purpose. In this section we suggest that whether a traffic management practice is reasonable largely rests on the answers to four questions regarding the techniques and practices used. In the next section, we will consider examples of how these techniques are used by ISPs to form traffic management practices.

The first two questions apply to traffic management *techniques*, because they are directed at the layer ("where") and functionality ("what"). The second two questions apply to traffic management *practices*, because they are directed at the actor ("who") and the manner and purpose ("when").

The first question is:

(1) WHERE: Where in the network, and at which layer, is the traffic management technique applied?

To answer the "where" question, we need to understand layers. The Internet is based on the concept of a layered architecture, where each layer provides certain functionalities. The reference model for layered architectures is the OSI model, developed by the International Standards Organization. The OSI model is composed of 7 layers, while the Internet model collapses this down to 4 layers, as pictured in figure 1. It is useful to think of the physical connection, e.g., wire, as being located below the bottom-most layer (layer 1) and the user, e.g., you, as being located above the top-most layer (OSI layer 7). Traffic management can be applied in any of the layer. Access control is often im-

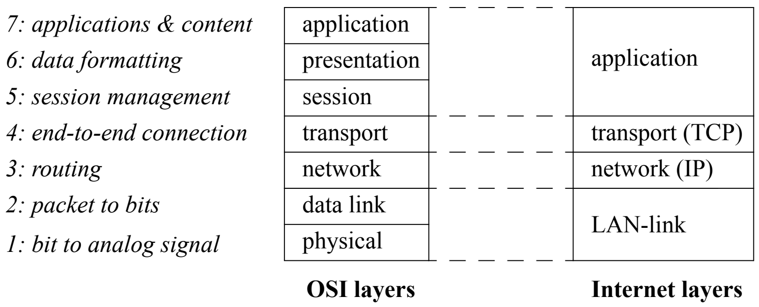

Fig. 1. OSI and Internet layered models

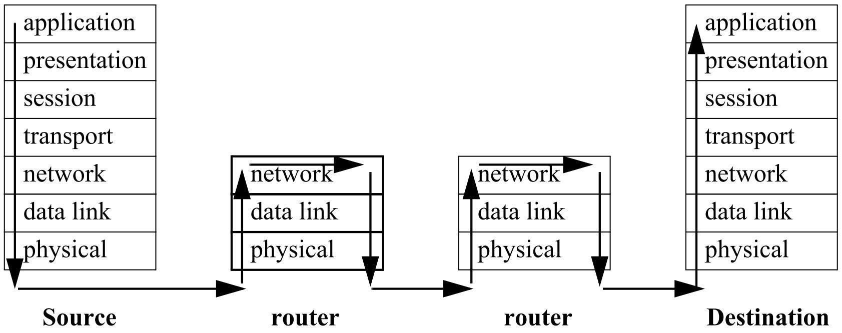

Fig. 2. Routing

plemented in the Internet Transport or Application layers, e.g. the Integrated Services (IntServ) protocol is at the transport layer. Traffic shaping can be implemented at any Internet layer, e.g. the TCP protocol at the Transport layer, the Differentiated Services (diffServ) protocol at the Network layer, and similar protocols which control flow on a local area network at the LAN-link layer. Techniques that use deep packet inspection (DPI) are usually operating at multiple layers, e.g. termination of selected peer-to-peer connections identified using DPI operates at both Transport and Application layers.

It is worthwhile to distinguish between techniques operating purely at endpoints versus techniques operating at transit points such as routers. Not every network device contains all 7 layers. Personal computers do contain all 7 layers - the network interface card (e.g., Ethernet card) implements OSI layers 1 and 2, the operating system (e.g., Windows) implements OSI layer 3 and part of layers 4 through 7, and user-installed software implements the remainder of layers 4 through 7. A network router, however, often contains only layers 1 through 3. As a result, communication from source to destination follows a complicated path as pictured in figure 2. Indeed, one of the Internet's key design paradigms, the end-to-end principle (Saltzer et al., 1984) suggests that network functionality should be implemented in OSI layers 1 through 3, and hence in each router, only if it cannot be implemented effectively in higher layers. This principle has been followed in much (but not all) of Internet design to date.

With respect to the “where” question, we thus propose that the pertinent distinction should be whether the technique is applied (a) at or above transport layer *and* in a transit node or (b) either below transport layer *or* at an endpoint. If a technique is applied at or above the transport layer, then good network design recommends that it be applied only at an endpoint. Therefore, techniques that are applied at or above the transport layer *and* in a transit node likely violate this design principle; although this violation is not sufficient to make a traffic management practice unreasonable, it should raise a red flag. None of the standardized traffic management techniques discussed in the previous section is intended to be used in a manner that violates this principle. However, a number of proprietary traffic management techniques do violate layering. Proprietary products such as Sandvine (Sandvine Incorporated, 2004) or PacketShaper (Packeteer Incorporated, 2007) are used in transit nodes and involve DPI; this violates layering since transit node devices should not inspect transport or application layer headers or the application layer payload. Firewalls also violate layering if they are placed in transit nodes.

If a technique is applied below transport layer, then the design principle allows implementation at transit nodes. For instance, guaranteed QoS can only be provided by offering QoS in every portion of the network that may experience congestion. Thus, the reasonableness of traffic management practices implemented below transport layer may also depend on whether such techniques are available at transit nodes.

Are these practices reasonable? The answer to the “where” question alone is not sufficient to make this determination; other questions must be considered.

The second question is:

(2) WHAT: What type of traffic management functionality is applied?

The purpose of a set of traffic management techniques is to management the overall performance of the network. The goal of this performance, however, must come from the objective of the network operator, which is considered below. However, one can usually define this goal either in terms of ensuring a maximum acceptable congestion level or in terms of maximizing revenue. Given a goal, the network operator has the choice of a wide variety of techniques to try to achieve this goal, and will usually apply several techniques that complement each other.

Different techniques work on different time scales. For congestion that lasts only a short time (e.g. less than a second), traffic shaping at transit points is generally considered effective. If the goal is to ensure a maximum overall acceptable congestion level, then the burst of packets is simply queued. If there are multiple types of flows, then priority service techniques such as diff-Serv and weighted fair queueing can ensure different minimum performance levels for each type or can attempt to maximize revenue. If the congestion lasts for more than a short time, however, such techniques are insufficient and queues overflow. For congestion that lasts for a moderate time (e.g. more than a second but less than a minute), traffic shaping at endpoints is generally considered effective. Applications that can tolerate large fluctuations in instantaneous throughput, such as file transfer, email, or web browsing, are called *elastic*. End-to-end flow control, usually TCP, can be applied to such elastic applications, and this is typically an effective method of reducing moderate time scale congestion. If the congestion lasts for more than a moderate time, however, such techniques result in unacceptable performance for inelastic applications. Therefore, for congestion that lasts a long time (e.g. more than a minute), access control is often required to reduce the number of flows competing for resources.

Access control is thus a more severe form of traffic management than end-to-end flow control, which is in turn more severe than traffic shaping at transit points. With respect to the "what" question, we thus propose that the pertinent distinction should be whether the functionality of the traffic management technique is (i) blocking or termination of a session versus (ii) enhancement or degradation of QoS. Blocking or termination is a severe form of traffic management and should raise a red flag. In contrast, enhancement or degradation of QoS is much less severe if applied in moderation.

The technique must be matched broadly to the type of application, since a user's perception of performance depends on the way the application responds to fluctuations in QoS. The choice is generally between suffering temporary but significant reductions in rate and having your connection blocked or terminated. Users of elastic applications would generally prefer that throughput fall in times of congestion rather than the flow be blocked or terminated; thus end-to-end flow control is appropriate. In contrast, Users of applications that are inelastic, e.g. voice over IP (VoIP) or video conferencing, are unwilling to tolerate a significant reduction in rate and are upset with call terminations and would prefer that the connection be blocked before it starts if network capacity is insufficient to maintain an acceptable quality connection.

The majority of the traffic management techniques discussed above use enhancement or degradation. A few, however, use blocking or termination. The IntServ architecture includes provisions to block new connections if adequate resources are unavailable. VoIP applications may block or terminate connections if sufficient QoS cannot be maintained. Sandvine's traffic management products can terminate selected TCP connections (Sandvine Incorporated, 2007). Firewalls are intended to block selected connections. The red flag raised by the use of blocking, therefore, is not sufficient to determine whether a practice is reasonable; other questions must be considered.

The third question is:

(3) WHO: Who decides whether the traffic management practice is applied?

We propose that the pertinent distinction should be whether the traffic management practice is applied (i) directly by a user or by an ISP only when a user desires this action versus (ii) by an ISP independent of a user's wishes. A user or a user's application can communicate its desires in several manners. If the technique is applied solely at the endpoints, then the source and destination devices can jointly determine whether and how to use the technique without the need for network operator participation. If the technique requires implementation at transit points, then a user can be presented with choices by the network and can respond according to the user's objectives; most likely this would be done in an automated manner by the user's application based on a profile selected by the user. For example, if the network used IntServ to reserve network capacity for some flows, then the user could decide which flows should receive enhanced performance. Similarly, if the network used diffServ to prioritize some packets, then the user could mark those packets that should receive priority. In this manner, the user decides which applications are important to him/her and the perceived value of good performance.

Actions taken by a user or under the user's direction are generally not deemed to be unreasonable. However, actions taken unilaterally by an ISP should raise a red flag, worthy of further investigation. Examples of each of these are given in the next section.

The final question is:

(4) WHEN: On what basis is it decided to apply the traffic management practice?

Traffic management can be used in various manners and for various purposes. Rather than relying on case-by-case analysis, we propose that the pertinent distinction should be whether the traffic management practice is applied to certain traffic on the basis of (i) the application, (ii) the source and/or destination, (iii) service provider, and/or (iv) payment. Practices applied to certain applications may be reasonable if they are done in a nondiscriminatory manner. Practices applied to traffic based on source and/or destination, however, are likely to raise a red flag out of anti-competitive concerns. Similarly, practices applied only to traffic carried by certain service providers are likely to raise a red flag for the same reason. In contrast, the reasonableness of practices applied on the basis of payment is likely to rest on the reasonableness of the payment amount. Examples of each of these are given in the next section.

3 Traffic management practices

In the previous section, we proposed four questions that affect the degree to which a traffic management practice is reasonable or unreasonable. In this section, we consider five examples of traffic management practices and discuss how the answers to these four questions affect their reasonableness. In the next section, we will use the lessons learned here to construct a framework for determination of whether a traffic management practice is reasonable.

First, consider the use of session management techniques that started this debate over traffic management practices. Sandvine's traffic management products are capable of identifying and terminating file-sharing connections (Sandvine Incorporated, 2004, 2007). In 2008, Comcast used products such as Sandvine to terminate TCP connections carrying BitTorrent packets used for uploading files from a Comcast subscriber to a destination outside the Comcast network, when the Comcast subscriber was not simultaneously downloading files (Comcast Corporation, 2008). For this traffic management practice, the answers to the four questions are:

Where: at or above the transport layer, in a transit node (red flag).
What: termination (red flag).
Who: by an ISP independent of a user's wishes (red flag).
When: on the basis of the application and the destination (possible red flag).

This practice raises at least three red flags: (1) it violates layering, because a transit node operates at or above the transport layer; (2) it involves termination of a connection; and (3) it is done independent of a user's wishes. With so many red flags, we easily find this practice to be unreasonable. The principal reason is that causes undue harm to consumers, since there are more direct and transparent manners to limit traffic from a user. Indeed, the FCC concluded that the practice is unreasonable, by relying on the following aspects of the practice: blocking, anti-competitive harm, lack of disclosure, and lack of tailoring of the practice to combat network congestion (FCC, 2008).

Next, consider another class of practices that involves blocking or termination of connections – firewalls. The answers to the four questions for firewalls are:

Where: at or above the transport layer (ok), at the endpoint or in transit nodes (red flag).
What: blocking (red flag).
Who: directly by a user or by an ISP only when a user desires this action (ok), or by an ISP independent of a user's wishes (red flag).
When: on the basis of the application and/or the source and/or destination (possible red flag).

The use of firewalls as a traffic management practice can thus also raise several red flags. First, firewalls can be implemented in endpoints (e.g. Windows Firewall) or in transit nodes (e.g. in wireless routers or network gateways). When implemented in transit nodes, this is a layering violation which raises a red flag. In addition, firewalls such as parental control software can be used to block traffic from certain sources, which raises another red flag. However these uses of firewalls are universally accepted as reasonable forms of traffic management. Why? The answer is that such firewalls are under the control of the end user. In contrast, firewalls have sometimes been used by ISPs independent of a user's wishes. In 2005, Madison River Communications blocked ports used by VoIP applications, which the FCC concluded is unreasonable traffic management (FCC, 2005a). Currently, many ISPs block connections to or from specific ports to combat spam (e.g. blocking outgoing SMTP traffic to port 25) or to prohibit residential servers (e.g. blocking incoming traffic to selected server ports). While combating spam is a worthy goal, users may desire to send email via other ISP's servers; hence, we conclude that the use of firewalls in this manner is a traffic management practice that should be used only with the consent of the user. Similarly, while ISP contracts may prohibit operation of a residential server, there are more direct and transparent manners to limit traffic to and from a user, and this practice should be considered unreasonable. We conclude that if the "who" question is resolved in favor of user choice, then the other red flags do not matter.

What about traffic management practices that involve limited degradation of traffic without blocking or termination? Many products offer proprietary traffic shaping techniques, and a number of ISPs use these techniques to limit file-sharing traffic. The answers to the four questions for this practice is:

Where: at or above the transport layer, in a transit node (red flag).
What: degradation (possible red flag).
Who: by an ISP independent of a user's wishes (red flag).
When: on the basis of the application (possible red flag).

Many educational institutions implement this practice by configuring products such as PacketShaper to limit the network bandwidth used by file-sharing applications (Packeteer Incorporated, 2008). The practice delays the transmission of file-sharing packets and hence slows down the rate at which these streams are forwarded through the device. This type of traffic shaping could be implemented at the network layer if low priority packets were labeled by the user. However, without the user's involvement to identify low priority packets, products such as PacketShaper use DPI to determine which packets belong to file-sharing applications. Use of DPI classifies this practice as an application layer practice; because an application layer practice is applied at a transit node, it violates layering, which raises one red flag. A second red flag is raised because the practice is typically applied without the consent of the user. This type of practice is less severe than blocking of termination; opinions differ as to whether these two red flags are sufficient to classify the practice as unreasonable. Since there are more direct and transparent manners to limit traffic from a user, we are reticent to classify such techniques as acceptable. However, because these alternative practices involve different business models that may require some time to be accepted by the public, we recommend classifying traffic shaping for file-sharing traffic as a borderline traffic management practice that could be used for a limited period of time if properly disclosed in the user contract.

Next, consider another class of practices that involves limited degradation – tiering. The answers to the four questions for tiering are:

Where: at or below the network layer, in a transit node (ok).
What: degradation (possible red flag).
Who: by an ISP on the basis of a user's wishes (ok).
When: on the basis of consumer payment (ok).

Tiering is typically accomplished in transit nodes (the user modem and/or ISP routers) at the data link and network layers by limiting the user download and upload rates to the maximum rates dictated in the user contract. This is a form of degradation, since the equipment is capable of transmitting at higher rates. However, since this practice is applied on the basis of user choice (and clearly displayed in user contracts), this is universally considered to be reasonable traffic management.

Finally, consider an example of a traffic management practice that involves enhancement of QoS. Currently, this is commonly used to support an ISP's own offering of VoIP or video-over-IP:

Where: at or below the network layer, in transit nodes (ok).
What: enhancement (possible red flag).
Who: by an ISP on the basis of a user's wishes (ok).
When: on the basis of the application and the service provider (possible red flag).

Enhanced QoS for real time applications such as voice and video typically requires the use of traffic management techniques that offer QoS in the data link and/or network layers in every portion of the network where congestion may occur (see e.g. Cox Communications (2004)). When an ISP uses enhanced QoS for its own VoIP and/or video-over-IP offerings, it uses these practices within its own network. In the case of VoIP, the traffic is then transited onto the public switched telephone network which offers similar QoS. In the case of video-over-IP, the video source usually resides on the ISP's network, so the entire network path (up to the subscriber premises) is under the control of the ISP. Although the practice is applied without the ability for a user to decline this enhancement, presumably no user would desire their voice or video service to have a lower QoS. This practice does however raise one red flag because it is applied only to voice and/or video service offered directly by the ISP. We do not object to the use of QoS, nor to charging for QoS; however, we do believe it is an acceptable traffic management practice *only if* the ISP offers the same QoS service for services offered by other providers at a rate that is not unreasonably discriminatory (Jordan, 2009).

4 A framework for determination of whether a traffic management practice is reasonable

In the previous two sections, we proposed four questions that affect the degree to which a traffic management practice is reasonable or unreasonable, and investigated the reasonableness of five examples of traffic management practices on the basis of the answers to these questions. In this section, we propose a framework for determination of whether a traffic management practice is reasonable.

The order in which the questions are considered is important. Start with one part of the "where" question, the location in the network where the traffic management technique is applied. If the technique is *applied at an endpoint*, we propose that it be classified as a reasonable traffic management practice regardless of the answers to the other questions. One endpoint is the user; practices applied directly by the user are not in question. The other endpoint is the entity with which the user is communicating. When this entity is an ISP, the ISP is acting in the role of an application provider. Common examples of this situation are ISPs that offer email and/or web hosting services. However, a user can (or should be able to) receive such application services from a large number of potential providers. Since this market is competitive, practices applied at an endpoint that negatively impact the user's experience may drive users to change application providers, buy they need not change their ISP. Therefore, any traffic management practice applied at an endpoint should be classified as reasonable. In contrast, if the traffic management practice is applied *at a transit node*, we must consider the remaining questions.

Next consider the "who" question, namely who decides whether the traffic management practice is applied. If the traffic management practice is applied *directly by a user or by an ISP only when a user desires this action*, we propose that it should be classified as a reasonable traffic management practice because the user has control over whether the practice is applied. Such practices are common, and include many firewalls, parental control software, and tiering. If an ISP were to provide enhanced QoS for voice or video purely on the basis of consumer payment, then this payment for QoS would not be discriminatory and we propose that it be classified as a reasonable traffic management practice. In contrast, if the traffic management practice is an action *taken unilaterally by an ISP*, then it is worthy of further investigation. If a practice is used without user consent, then we believe it should be disclosed in sufficient detail in the user contract. If so disclosed, then we must consider the remaining questions to determine if it is a reasonable practice.

Before progressing to these remaining questions, however, we should limit the scope of the traffic management practices considered here. We only consider techniques that are applied to networks such as the the Internet that use a public right-of-way; private networks are free of such regulation. We only consider techniques that affect Internet applications; if an ISP offers a voice service under Title II of the Communications Act (which regulates common carriers) or offers a video service under Title VI of the Communications Act (which regulates cable communications), then these restrictions need not apply. We only consider lawful uses; ISP rights to detect and interfere with illegal uses are addressed elsewhere in law. we only consider non-harmful uses of the network; security measures may require special considerations. We do not consider issues of privacy, which intersect with many of the techniques discussed here but which require considerations beyond those detailed here. Finally, prohibition of unreasonable practices should implemented only where sufficient competition does not exist; Title I of the Communications Act includes a provision which instructs the FCC to forbear from applying regulations unless they are in the public interest and required to ensure just and reasonable practices. Toward this end, regulation of reasonable traffic management should only apply to access networks, specifically to the portions of an ISP's network which must be transversed to form routes from the Internet to its subscribers.

The next aspect to be considered is the "what" question, in particular whether the practice involves blocking or termination of a session versus enhancement or degradation of QoS. If the practice involves *blocking or termination*, we propose to classify it as unreasonable. Blocking or termination practices that are applied at a transit node without user choice are unreasonably anti-competitive, cause undue harm to consumers, or unreasonably impair free speech. When blocking is applied at a transit node without user choice *on the basis of the source or destination or on the basis of the speech within the packet*, the practice unreasonably impairs free speech; this type of blocking includes blocking of specific web pages or blocking on the basis of the content of the speech. When blocking is applied at a transit node without user choice *on the basis of the application*, the practice is unreasonably anti-competitive and/or causes undue harm to consumers; this type of blocking includes blocking of specific applications (e.g. blocking or terminating VoIP or file-sharing connections) and blocking of specific ports (e.g. SMTP or server ports). There is no reasonable justification for the use of these techniques. In some cases, the ISP's goal may be to limit congestion, reduce spam, or implement security;

however, such goals can be implemented either through less severe methods that do not involve blocking or with the consent of the user. If a traffic management practice is implemented in a transit node, without user choice, but does not block or terminate connections, we must consider the remaining questions.

Practices that *enhance or degrade QoS in a transit node without user choice* are the concern of the remainder of this section of the paper. To address such practices, consider the "when" question, which asks on what basis is it decided to apply the traffic management practice. This question considers the manner and purpose of the practice. We propose that the pertinent distinction should be whether the traffic management practice is applied to certain traffic on the basis of (i) the application, (ii) the source and/or destination, (iii) service provider, and/or (iv) payment.

First, consider using *source and/or destination and/or service provider* as the basis. A common example of this practice is an ISP that provides enhanced QoS for its own VoIP service, but does not provide this same QoS to competitors VoIP packets. Another example of an exclusive arrangement would occur if an ISP were to provide access to enhanced or degraded QoS to some third party application providers but not others. Use of source and/or destination and/or service provider without user choice involves the use of exclusivity. Such exclusive arrangements are unreasonable, since they tilt the playing field between application providers through use of Internet infrastructure. Thus, we propose that these traffic management practices be classified as unreasonable, because they are unreasonably anti-competitive.

Next, consider using *payment* as the basis for the decision of when an ISP uses enhanced or degraded QoS. For instance, an ISP could charge a consumer for enhanced QoS for all packets to or from that subscriber. Alternatively, an ISP could charge an application provider for enhanced QoS for all packets to or from that application provider. Consumer payment for QoS places the use of the practice under the control of the user, and hence this framework would already have classified such practices as reasonable. We thus only need consider charging of application providers. We considered this case in detail in Jordan (2009). If the price is not unreasonably discriminatory (e.g. if an ISP sells QoS to all application providers at the same price as it passes on to its own applications that require QoS), then we argued in Jordan (2009) that the practice is reasonable. However, if prices for QoS are unreasonably discriminatory, then a traffic management practice that uses such prices as the basis is unreasonable since the practice is unreasonably anti-competitive.

Finally, consider cases in which the practice is applied on the basis of the *application*. In these cases, if the practice is applied entirely *at or below the network layer*, then we propose that the practice be classified as reasonable. Enhancement or degradation of QoS is thus applied to specific packets iden-

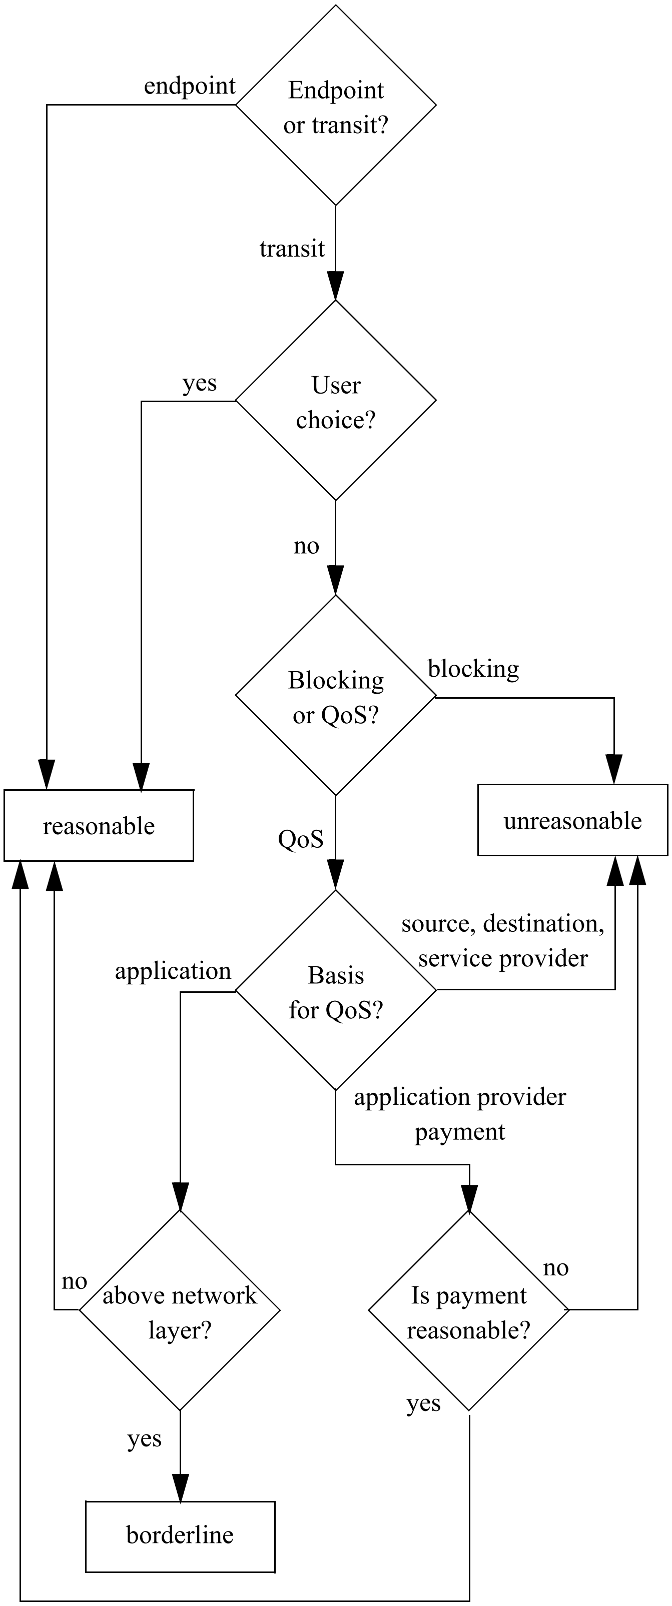

Fig. 3. The framework

tified by the user, for instance if an ISP chose to give enhanced QoS to all packets identified using diffServ codepoints by the user as VoIP.

The last remaining case consists of practices that are applied *at or above the transport layer at transit nodes without user consent and enhance or degrade QoS on the basis of the application.* Practices of this sort use DPI to identify which packets should receive high or low priority or dedicated bandwidth. A common example of this practice is traffic shaping for file-sharing. Because DPI is used (rather than user identification of these packets), this practice violates layering. The question is whether this violation of layering is severe enough to cause this practice to be classified as unreasonable. There are more direct techniques that can be used that rely on user identification of packet priorities and that do not violate layering, as discussed in the next section of this paper. However, because these alternative practices involve different business models that may require some time to be accepted by the public, we recommend classifying any such practice that uses DPI to apply QoS as a borderline traffic management practice that could be used for a limited period of time if properly disclosed in the user contract.

The resulting framework is summarized by the flowchart in Figure 3.

5 Conclusion

We have proposed a framework that can be used to determine whether a traffic management practice is reasonable. The next step is to illustrate how unreasonable practices can be modified to become reasonable and to achieve similar purposes.

References

Comcast Corporation, Feb 2008. Comments of Comcast Corporation before the Federal Communications Commission in the matter of broadband industry practices.

Cox Communications, May 2004. Whitepaper: Voice over Internet protocol: Ready for prime time.

FCC, 2005a. DA 05-543, Madison River Communications Consent Decree. Available at http://hraunfoss.fcc.gov/edocs_public/attachmatch/DA-05-543A1.pdf.

FCC, 2005b. FCC 05-151, Internet Policy Statement. Available at http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-05-151A1.pdf.

FCC, 2007. FCC 07-31, Broadband Market Practices Notice of Inquiry. Available at http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-07-31A1.pdf.

FCC, 2008. FCC 08-183, Comcast Order. Available at http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-08-183A1.pdf.

Free Press, Public Knowledge et. al., November 2007. A petition before the Federal Communications Commission for declaratory ruling that degrading an Internet application violates the FCCs Internet Policy Statement and does not meet an exception for reasonable network management.

Frieden, R., 2006. Network neutrality or bias? – Handicapping the odds for a tiered and branded Internet. In: Proceedings of the 34th Research Conference on Communication, Information and Internet Policy (TPRC). TPRC.

Jordan, S., May 2009. Implications of Internet architecture upon net neutrality. ACM Transactions on Internet Technology 9 (2), 5:1–5:28.

Lehr, W. A., Gillett, S. E., Sirbu, M. A., Peha, J. M., 2007. Scenarios for the network neutrality arms race. International Journal of Communication 1, 607–643.

Packeteer Incorporated, July 2007. Best practices: Monitoring and controlling peer-to-peer (p2p) applications.

Packeteer Incorporated, 2008. Packeteer education customers. Retrived on July 10, 2008.
URL `http://www.packeteer.com/customers/education.cfm`

Peha, J. M., 2007. The benefits and risks of mandating network neutrality, and the quest for a balanced policy. International Journal of Communication 1, 644–668.

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

Sandvine Incorporated, December 2004. Session management: BitTorrent protocol, managing the impact on subscriber experience.

Sandvine Incorporated, July 2007. Sandvine DPI-based policy solutions.

Vuze Inc., November 2007. A petition before the Federal Communications Commission to establish rules governing network management practices by broadband network operators.

Weitzner, D., May-June 2008. Net neutrality... seriously this time. Internet Computing, IEEE 12 (3), 86–89.