2

INFORMATION TECHNOLOGY

5

Doing the DPI Dance

Assessing the Privacy Impact of Deep Packet Inspection

Alissa Cooper
Center for Democracy and Technology
and Oxford Internet Institute

Massive growth in data processing power has spurred the development of deep packet inspection (DPI) equipment that potentially allows providers of Internet service and other intermediaries to collect and analyze the Internet communications of millions of users simultaneously. DPI has come to permeate numerous Internet policy discussions, including those related to net neutrality, behavioral advertising, content filtering, and many others. Although the policy concerns that DPI raises differ in each case, one theme that recurs is the potential for DPI to eliminate online privacy as it exists today, absent the pervasive use of encrypted communications. As a technology that can provide Internet service providers (ISPs) and their partners with broad and deep insight into all that their subscribers do online, its potential to facilitate privacy invasion has been described in the most dire of terms: as "wiretapping" the Internet (Barras 2009, 1), "unprecedented and invasive ISP surveillance" (Ohm 2009, 1417), and even "the end of the Internet as we know it" (Riley and Scott 2009, 1).

Residential ISPs' use of DPI has drawn scathing privacy criticism—and attention from policy makers in both the United States and the EU—despite the fact that numerous other entities are capable of conducting content inspection. Content delivery networks and caching services could have similar capabilities, as can individual Internet users employing firewalls, home gateways, or packet sniffers. Likewise, web- and software-based service providers have been providing many of the services that DPI can facilitate for ISPs—security protections, behavioral advertising, and content filtering, for example—for years.

There are several characteristics inherent to residential ISPs and their use of DPI that significantly increase the privacy stakes as compared to these

other entities, however. ISPs are uniquely situated in three respects: they serve as gateways to all Internet content, changing ISPs can be difficult for Internet users, and the use of a tool as powerful and versatile as DPI makes it prone to invisible mission creep. All of these characteristics are difficult or impossible to mitigate, and together they form the fundamental basis for the heightened privacy alarm that has characterized DPI discussions. To some, these characteristics are enough to reject DPI altogether and call for its prohibition (NoDPI 2008).

Few legal prohibitions of such technologies exist today, however. The application of existing communications privacy laws to new uses of DPI is unclear at best and inadequate at worst, providing space for ISPs to experiment with many different uses of DPI. In the United States, for example, the Wiretap Act as amended by the Electronic Communications Privacy Act (PL 99-508, 1986) prohibits the interception and transfer of electronic communications. But both of these prohibitions have exceptions for business-related uses that are not well defined, and the application of the law depends on the arrangement of flows of intercepted data between network operators and other parties. Furthermore, the case law that exists to interpret the prohibitions comes largely from outside the Internet context, leaving uncertainty about the meaning of the law in the DPI context.

ISPs, meanwhile, see promising opportunities of many kinds in the growth of DPI. The technology can provide them with a powerful tool to address constantly evolving challenges in managing network congestion and security threats. It can provide insight into how their networks are being used, allowing them to make more informed decisions about network upgrades and architecture. And perhaps most importantly, DPI is among a set of tools that can provide ISPs with new revenue streams, whether by funneling data about users to advertisers, selling expedited delivery to content providers, or levying extra fees on heavy network users. As the services that telephone and cable companies have traditionally relied on for the bulk of their revenue—multichannel video and voice—are increasingly forced to compete with similar web- and IP-based offerings from unaffiliated applications providers, developing and monetizing intelligence in the network's core is becoming increasingly tempting.

ISPs around the world are taking up DPI for all of these different reasons (Finnie 2009); one DPI vendor claims that its ISP clients serve 20 percent of the world's fixed broadband subscribers (Verhoeve 2009). Given the attractiveness of deep packet inspection for ISPs and the nebulous legal landscape, ISPs are likely to continue to deploy and experiment with DPI on their networks for the foreseeable future. Under these circumstances, the potential that DPI creates for privacy invasion should be examined, the

tools available to mitigate its associated privacy risks should be explored, and both of these concerns should be assessed against deep packet inspection's benefits as perceived by Internet users. This chapter takes on those tasks. The first section articulates a clear, comprehensive definition of the term *deep packet inspection* for use in conducting privacy analysis. Using that definition, the following section explores the inherent privacy risks of ISPs' use of DPI. The second half of the chapter delineates DPI's most prevalent uses and discusses how a set of techniques to mitigate privacy invasion could be applied in each case. The chapter concludes with a summary of the entire analysis.

UNDERSTANDING DEEP PACKET INSPECTION

Taken together, the meanings of "deep," "packet," and "inspection" describe both a technology and a practice that warrants analysis from a privacy perspective. This section explores the meaning that each component contributes to the term as a whole by discussing the different parts of a packet, defining "deep" in terms of which of these parts might be examined and taking an expansive view of the types of analyses that might be performed of packet data.

"Packet"

All Internet communications are comprised of packets: small pieces of data that get transmitted from one end of a communication medium to the other. Internet packets follow a layered structure, with each successive layer of information inside a packet corresponding to a particular function that is necessary to have the packet successfully delivered to and used by the recipient. One generic network model that is often cited as a basis for this layered structure is the Open Systems Interconnection (OSI) Reference Model (International Organization for Standardization 1996). The OSI model consists of seven layers: (1) physical, (2) data link, (3) network, (4) transport, (5) session, (6) presentation, and (7) application. These layers form a complete representation of the network, from its physical capabilities (electrical signals sent on a cable or telephone wire, for example) at layer 1 to applications and services (email or web browsing, for example) at the highest layers.

For the purpose of defining DPI and understanding its privacy impact, a mapping between the structure of Internet packets and the OSI layers is useful (see figure 5.1). Internet packets are addressed according to the Internet Protocol (IP) (Postel 1981b). Every Internet packet contains an IP header—a

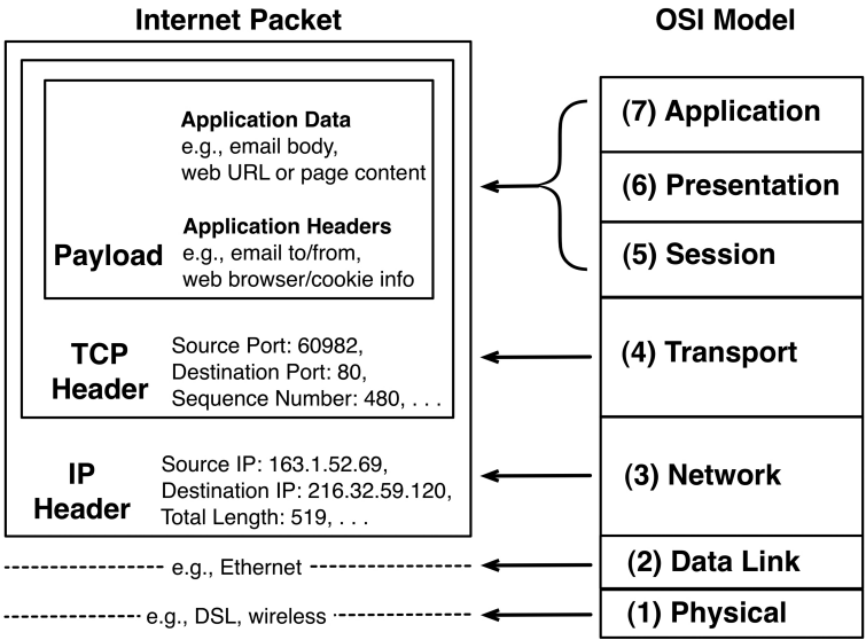

Figure 5.1. Mapping between Internet packets and OSI layers (layers 3 to 7 are of interest for defining DPI).

set of metadata fields that describe characteristics of the packet that allow it to successfully reach its destination. IP header fields include the source and destination addresses for the packet, the length of the packet in bytes, and several other items. The IP header corresponds to OSI layer 3, the network layer. Historically, ISPs' inspection of packet data has been limited to data in the IP header, with the remainder of the packet treated as the data to be delivered. This distinction between headers and data exists at each successive layer of Internet packets.

Beyond IP headers, Internet packets also contain transport-layer headers. Although several transport-layer protocols are in wide use, the most common one is the Transmission Control Protocol (TCP) (Postel 1981a). TCP serves several layer-4 functions, including identifying a device-level address (known as a port) where packets should be delivered and providing error detection for when packets get dropped. The port and sequence number fields shown in figure 5.1 are used for each of these purposes, respectively. As with IP, a division exists between TCP headers and data—the network nodes that implement TCP can find all the information they need in the headers and treat the remainder of the packet as data to be delivered.

For simplicity, this chapter will use the term *payload* to refer to the data portion of a packet carried by TCP, although in practice a packet's TCP headers and data together can also be considered as a "payload" carried by IP.

The payload, which comprises OSI layers 5, 6, and 7, contains all of the application- and content-related information in the packet. All of the bits that comprise an email message, Voice over IP (VoIP) call, or web surfing session reside in this section of the packet. Although the OSI model separates the payload into three separate layers of functionality, for the purpose of analyzing the privacy impact of DPI, the entire payload can be considered as a single unit.

Just as with IP and TCP, many Internet applications use the header/data model, with the headers facilitating some functionality necessary for the application to function and the data portion containing the data communicated by the application. Web requests, for example, usually contain headers that serve various functions in executing the request, including providing information about the user's browser or cookies. The data portion is the URL being requested (which in many cases also includes search terms) or the content of the web page. Similarly, the "to," "from," and "subject" fields of an email message have in some circumstances been considered to be application headers, whereas the body of an email message has been treated as application data, although in some cases all email fields are considered as application data.

"Deep"

There has been some controversy about just how "deep" the inspection of packets needs to be for it to qualify as DPI and potentially jeopardize users' privacy. The strictest conception of "deep" draws a line between IP addresses and all other headers and data in the packet, claiming that the use of any data other than the destination IP address constitutes DPI (Bowman 2009). A slightly more expansive conception makes a distinction between IP headers and the rest of the packet: inspection of any packet data other than IP header fields is considered deep (Reed 2008). Parsons (2008) provides an even more nuanced definition by classifying three levels of depth: "shallow," which uses OSI layers 1–3; "medium," which uses OSI layers 1–5; and "deep," which uses data at all OSI layers.

Taking these definitions together with an understanding of the layered structure of Internet packets, a conception of how "deep" DPI needs to be before it raises its own privacy implications begins to emerge. It seems quite clear that Internet service providers' use of IP headers, at OSI layer

3, is unobjectionable. IP headers have always been used by ISPs to route packets to their destinations, and thus it would be difficult to argue that their continued use creates some new privacy risk that has not existed since the Internet Protocol was first developed.

TCP headers provide a minimal amount of additional information about an Internet user's online activities, primarily in the form of port numbers. Many applications adhere to specific registered port numbers; for example, HTTP uses port 80, and email frequently uses port 25. Thus, by inspecting TCP headers, ISPs may glean some limited information about their subscribers' activities that goes beyond what IP headers reveal. However, in part because web browsing is a popular form of Internet usage, many nonweb applications have migrated to port 80 in order to take advantage of web optimizations or to avoid restrictions placed on other ports. Some applications also use nonstandard ports, or they change the ports they use over time. Thus, while TCP headers provide some application-level information, using them in isolation raises limited privacy concerns.

Determining the user's privacy interest in packet payloads is a thornier task. Payloads often contain application headers, and many of these headers—such as the HTTP version type and content encodings cited earlier—are fairly innocuous from a privacy perspective. However, other kinds of headers can reveal much more sensitive information about a person's Internet activities, such as URLs, email recipient addresses, user names, addresses, and many other kinds of data.

Furthermore, for most application headers, it is next to impossible for an ISP to pluck out the header information without also inspecting at least a small number of other data within a payload. While IP and TCP headers have standardized formats, there is no standard format for all payloads, nor a mechanism for ISPs to know definitively from only IP or TCP header information whether a particular packet will contain an HTTP request, email to/from headers, or any other specific application data. ISPs can make good guesses about what a packet contains by observing the characteristics of the traffic flow—the sequences, sizes, and timing of streams of packets—together with TCP port information (Allot Communications 2007), but ultimately, correctly identifying the application requires inspecting the packet payload itself.

The same analysis holds true for application data, including the content that Internet users access online and generate themselves. Internet users may not have a particularly strong privacy interest in some of these data—the content of an online weather report, for example. But the breadth of activities that Internet users engage in is increasingly large and can incorporate infor-

mation at all levels of sensitivity. Email, instant messaging, VoIP, file sharing, and the innumerable list of web-based activities—from reading the news to visiting social networks to searching for health information—each carry with them a particular set of privacy expectations. To attempt to ferret out particular bits that an ISP could inspect without sweeping in privacy-sensitive content is not likely to be feasible.

The thread tying all of these pieces together is the inspection of application-level data. When ISPs go beyond their traditional use of IP headers to route packets, they begin to implicate data that their subscribers have an interest in protecting. The privacy interest may be minimal, as with the mere inspection of TCP headers. But beyond OSI layer 4, drawing any firm conclusions about which parts of a packet are or are not privacy sensitive becomes exceedingly difficult. Thus, the data of concern from a privacy perspective are application-level data—any data above OSI layer 3 that relate to an application—with the understanding that the inspection of OSI layer 4 data alone may incur limited privacy risk.

"Inspection"

Although common understandings of "inspection" include merely scrutiny or examination, in the DPI context it has taken on a much broader meaning that may encompass interception, collection, observation, analysis, and storage of application-level data. This expansive conception of "inspection" suits DPI because the technology that comprises a single DPI system—and that DPI vendors sell as a single product—can be put to so many different uses. Some of these uses require only real-time analysis of individual packets crossing the network, while others involve intercepting, storing, and later analyzing users' packet content. This broad meaning of the word is likewise a suitable foundation for privacy analysis because it is not merely the inspection of data that may create privacy risks but also what happens before (collection) and after (analysis and storage) data inspection. Not all DPI systems will necessarily perform every one of these functions, but all of the functions that are performed should be included in the privacy analysis of any DPI system.

Combining this notion of inspection with the discussions of "deep" and "packet" yields the following definition of DPI for use in analyzing its privacy impact: *Deep packet inspection is the collection, observation, analysis, and/or storage of data related to an application that is found in Internet packets above OSI layer 3.* Figure 5.2 shows the subset of packet data and OSI layers that are covered by this definition.

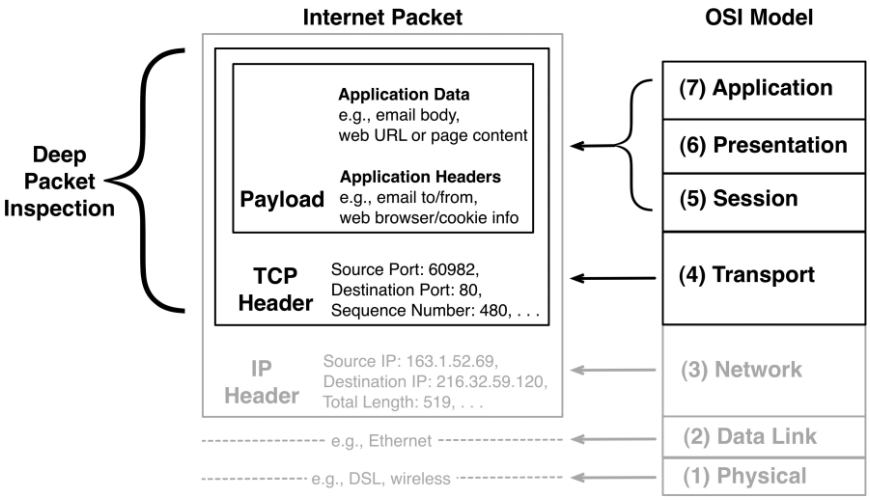

Figure 5.2. Subset of packet data and OSI layers involved in DPI.

SINGULAR CHALLENGES OF ISP USE OF DEEP PACKET INSPECTION

The definition of DPI given above is highly generic and not limited to any particular service or functionality that an ISP might want to implement. But even at its most generic, abstract level, an ISP capable of observing, analyzing, and storing application-level data presents unique risks to privacy that do not apply to other kinds of service providers with the same capability. As compared to their web- and software-based kin, ISPs using DPI are uniquely situated in three respects: they serve as the gateways to the Internet, the costs of switching between ISPs are high, and their use of DPI is prone to invisible mission creep. Each of these characteristics is both inherent to ISPs' use of DPI and difficult or impossible to mitigate, raising the privacy stakes above and beyond those for other service providers in many cases.

ISPs as Internet Gateways

The Internet is often thought of as a dramatically free medium for speech, where little stands between an Internet user and the expression of her ideas to friends, colleagues, and the world at large. It is also an intensely personal medium used to maintain familial and social ties, to find information related to personal activities and pursuits, and to transact personal business. The Internet provides a single communications platform that supports services tra-

ditionally offered by disparate infrastructure providers, including broadcasters, telephone companies, banks, and many others. Millions of Internet users worldwide trust the medium enough to engage in a wide range of personal and commercial communications and transactions online. While "the medium" is composed of many services and applications providers at different levels, the foundation for this trust is the connectivity itself as provided by ISPs. Ohm (2009, 1446) has aptly described this service provider trust as the "sense of repose" that Internet users have as they use the network to conduct their lives.

DPI has the potential to disrupt this sense of repose by inserting a middleman—and potentially a gatekeeper—between Internet users and those with whom they communicate. To the extent that Internet users find themselves at ease conversing and transacting online, ISPs' increased use of DPI presents the potential to chip away at that sense of security by introducing surveillance where it did not exist previously. ISPs are an important element of the trust that Internet users place in the network, and increased use of DPI calls that trust into question.

The effects of this loss of trust could be wide-ranging. As with other technologies of surveillance, increased use of DPI creates the potential for self-censorship and inhibition online (Lyon 2007). It may also serve to deter online commerce if consumers and businesses question the confidentiality of their transactions. These risks are plausible whether specific uses of DPI are known to Internet users or indeed breach confidentiality, because even a general awareness that surveillance may be occurring can prompt people to alter their behavior (Foucault 1977). Introducing DPI on the network thus has the potential to turn what was a trusted conduit into a suspicious eavesdropper, even if Internet users are only vaguely aware that DPI is in use.

Many other trusted service providers exist on the Internet, and many of them would be similarly capable of damaging users' trust should they begin to examine their users' communications in an unexpected way. In fact, there are clearly intermediaries in existence today that are capable of collecting more application-level data about many more Internet users than any single ISP could—Google is the obvious example. But neither Google nor any other service provider is as capable as an ISP of comprehensively monitoring the entirety of each individual subscriber's online activities. Every one of a subscriber's packets, both sent and received, must pass through the ISP's facilities. What separates ISPs from other service providers is the potential for their gaze over their subscribers to be omniscient.

ISPs may be far from realizing that potential, and encryption tools exist to help protect Internet users from the prying eyes of their ISPs. But as long as the majority of Internet users pursue their online activities without encrypting

their communications, the mere existence of DPI on the network jeopardizes the bond between them and their ISPs.

High ISP Switching Costs

The potential for ISPs to abuse their gate-keeping power is further exacerbated by the fact that switching to a new ISP is comparatively more difficult than switching between other services like search engines or web browsers. While the latter may involve a simple mouse click or software download, changing ISPs can be a much more elaborate process, involving a time investment to explore new alternatives and bundled services, installing new equipment, setting up new bill payments, and time at home waiting for an engineer to hook up new service (Krafft and Salies 2008). Because of these barriers to switching, subscribers may be unwilling or unable to switch ISPs even if their current ISPs introduce DPI-based practices to which they object. Internet users may perceive their choice of ISP to be much more binding than their choice of other online services, which reduces their ability and inclination to avoid ISPs' privacy-invasive practices. Moreover, users in many U.S. markets may have few competing ISPs to choose from (Horrigan 2009).

Notably, even where consumers have many ISP choices, switching costs may still impede consumers from changing ISPs because of DPI concerns. For example, in the competitive UK market, many ISPs indicate in their website disclosures that they are using DPI of some form to manage congestion, and the majority of Canadian ISPs that responded to a recent regulatory inquiry indicated that they are using DPI for some network management purpose (Parsons 2009). Whereas competition for privacy is appearing in other online sectors with low switching costs—the major search engines, for example, continue to improve upon each others' data retention policies (Center for Democracy & Technology 2007)—higher ISP switching costs may reduce ISPs' incentives to compete on privacy.

While there may be limited steps that ISPs can take to make changing easier—lowering or eliminating contract termination fees, establishing flexible schedules for hooking up new service, and so forth—the burdens of changing to a new ISP are in some ways inherent to the provision of Internet service. Because these burdens are largely unavoidable, relying on competition to discipline ISPs' privacy behavior is not likely to be sufficient.

Invisibility of Mission Creep

Another distinguishing feature of ISPs' use of DPI is the potential for mission creep: having DPI equipment that was installed for one purpose used

for multiple new purposes over time (Werbach 2005). The potential uses of DPI are nearly as wide as computing itself. Many of the capabilities of DPI equipment are generic computing capabilities: intercepting packets, pattern-matching their content, and storing the raw data, aggregations of the data, or conclusions drawn from the data. Because the wide variety of DPI uses employs some or all of these generic capabilities, DPI vendors are finding it more efficient and less costly to build their equipment to suit multiple uses. Several vendors tout the fact that a single one of their products can be used for congestion management, usage monitoring, and prioritized or tiered service offerings, for example (Arbor Networks 2010; ipoque 2008). The trend is toward more functionality built into individual DPI products, not less.

When mission creep occurs, it may be invisible to users. Because ISPs' use of DPI occurs in the middle of the network, there need not be any indication to subscribers that inspection takes place. There is also no technical reason why DPI equipment should leave any trace on users' computers; although, for some uses, DPI can work in conjunction with files or software stored on users' computers. This lack of visibility is in contrast to other kinds of technologies that can perform similar functions to those of DPI. For example, while many web-based behavioral advertising networks deposit cookies on users' computers for tracking purposes, an ISP could employ a DPI-based behavioral advertising system without storing anything on users' machines. Furthermore, one of the core design goals of DPI vendors is to build equipment that has the least possible impact on the network's performance and users' experiences (Allot Communications 2007). The combination of these technological elements creates the potential for DPI to be deployed—and subsequently put to new uses—mostly invisibly on the network.

Perhaps because of the fact that DPI technology does not need to reveal itself, several early DPI systems were deployed without any indication to users (European Commission 2009; Federal Communications Commission 2008). Furthermore, despite the limited public scrutiny that ISPs' DPI practices have been subjected to thus far, one large ISP has already admitted that "even though DPI equipment was originally intended to introduce usage data collection functionality . . . it was subsequently determined that DPI should be used for traffic shaping" (Engelhart 2009, 3). This sort of mission creep is precisely what raises concerns about the misuse of the technology and its ability to erode consumers' trust in the network. Features of the technology that are not easy to overcome drive concerns over mission creep: the cost effectiveness of producing general-purpose DPI equipment and its lack of transparency on the network.

PRIVACY ANALYSIS OF DPI USES

Internet service providers' use of DPI thus creates unique challenges to privacy—because ISPs serve as the trusted on-ramp to a medium that is intensely expressive and personal, Internet users may not easily extricate themselves from an ISP whose DPI practices they disagree with, and DPI presents particularly promising territory for invisible mission creep. For many stakeholders, these characteristics together are so serious that the stakeholders insist on its prohibition (NoDPI, 2008). ISPs, however, will continue to experiment with DPI as long as the legal landscape remains unsettled. A comprehensive privacy analysis of any DPI system cannot end with these generic considerations because the specific purposes for which DPI is used can create additional privacy challenges—and opportunities. The following sections discuss DPI's prevalent uses and a set of potential mitigation techniques that create particular challenges and opportunities in the DPI context.

DPI Uses

Various actors have already used DPI for a wide variety of purposes, and new uses will no doubt emerge. The following categories of prevalent DPI uses are discussed below: congestion management, prioritized service offerings, behavioral advertising, proactive security measures, troubleshooting, usage monitoring, and content filtering. While all of these functions can be performed using other means, the discussion below emphasizes how ISPs are using DPI, or are envisioning its use, in each case.

Congestion management—DPI can be used in a number of different ways to help ISPs manage congestion on their networks. Because deep packet inspection captures application-level data, ISPs can use DPI to identify applications with a particular set of network usage characteristics. For example, real-time applications such as VoIP transfer small amounts of data and require minimal network delay (known as latency) to ensure good call quality, whereas bulk transfer applications such as peer-to-peer file sharing may transfer high amounts of data and be more tolerant of higher delays. ISPs can use DPI (usually in combination with other techniques) to determine the application in use by each packet stream. They can then apply a particular bandwidth management rule to particular applications or classes of applications, such as prioritizing low-latency traffic or deprioritizing delay-tolerant traffic during times of congestion (Werbach 2005). ISPs in the United States (Casserly et al. 2008; Wilson et al. 2010) and Canada (Engelhart 2009; Henry and Bibic 2009; MacDonald 2009) have experimented with or deployed this

technique on their networks, drawing significant attention from regulators and privacy advocates in the process.

ISPs may also use DPI to as a tool to implement pricing practices aimed at managing traffic (Henry and Bibic 2009). For example, ISPs may charge subscribers based on which applications they use (Anderson 2007). Or an ISP may allocate to each subscriber a particular monthly volume of application-specific traffic—gaming or VoIP, for example—and levy fees on only those users who exceed their monthly thresholds. In either case, ISPs need the ability to determine which applications each subscriber is using, and DPI provides that capability. If ISPs institute bandwidth usage caps such that subscribers who exceed the caps are charged additional fees, they may want to exempt certain content (such as software upgrades) from counting against subscribers' usage allotments. DPI provides the tools necessary to make such distinctions.

Prioritized service offerings—The same kinds of capabilities that DPI provides for congestion management could also be used to provide affirmative prioritization to particular applications and services. For example, an ISP could use DPI to identify its own video content on the network and expedite the delivery of that content, it could charge other video content providers for such expedited delivery, or it could offer expedited delivery of specific video content to subscribers for an additional fee. These kinds of prioritized service offerings have been at the center of net neutrality controversies (Felten 2006), but the fact that DPI could facilitate prioritized services has been somewhat overlooked.

Behavioral advertising—Internet users' online activities present a wealth of information to advertisers seeking to better target their online advertisements. Web-based advertising companies have for many years used web-based technologies (such as cookies) to track the sites that users visit, allowing the companies to compile profiles of users' behavior for advertising purposes. DPI creates that same possibility for ISPs by allowing them to identify the websites that their subscribers are visiting, the content of those sites, and the other kinds of applications and data that subscribers are using. ISPs or their advertising partners can extract this information from individual packets and compile it into profiles that can later be used to show targeted ads to subscribers as they surf the web. ISPs in the United States (Johnson 2008; Martin 2008; Post 2008) and the UK (Bohm 2008; Clayton 2008) have experimented with the use of such systems.

Proactive security measures—ISPs provide many kinds of proactive security protections to their networks: filtering spam, blocking malware and viruses, and monitoring for intrusions and attacks. DPI, in combination with other tactics, can be useful for all of these purposes because it allows ISPs to

screen the traffic crossing their networks and identify potentially malicious content within packet payloads. DPI vendors claim that ISPs have been using DPI for security purposes for many years (Bowman 2009; Mochalski and Schulze 2009), and many ISPs reserve the right to do so (or to do something vaguely DPI-like) in their terms of service.

Troubleshooting—DPI can also help ISPs address network problems, including security issues, as they arise. An ISP may receive signals from below OSI layer 4 that indicate the presence of a problem, such as a denial-of-service attack or a failing router; the ISP can switch on a DPI system to investigate further. ISPs may also switch on their DPI tools in response to customers' service calls (MacDonald 2009). While some kinds of customers' problems, such as network outages, may not require application-level data to be resolved, others may be caused or exacerbated by specific applications or content, in which case DPI can provide better or more relevant information than information from lower layers.

Usage monitoring—Information about network usage can be helpful to ISPs as they seek to understand how usage differs across subscribers in different regions or with different service plans and where network upgrades may be necessary. While ISPs have long collected information about the volume and frequency of transmissions on their networks in the aggregate, DPI provides an extra level of insight into which applications and services are generating more or less traffic. Application-level insight in turn allows ISPs to manage their networks according to the performance demands of different applications (as discussed in the congestion management section) by adding bandwidth, changing congestion management policies, or rerouting traffic depending on its application or application type. ISPs all over the world are using DPI in this manner (Cho et al. 2006; MacDonald 2009).

Metrics describing web usage—namely, which websites are most popular among an ISP's subscriber base—may be of particular interest, both internally for the ISP and externally for website owners and advertising companies. Website owners have a keen desire to understand the volume of traffic to their sites, as do advertisers and advertising companies seeking to draw the largest or most targeted audiences possible to their ads. DPI can facilitate the process of compiling data about the volume of traffic to individual websites that can then be shared or sold to website owners and marketers. Some ISPs have partnerships with data vendors that purchase these data from ISPs and sell them to website owners and marketers (Hindman 2008).

Content filtering—DPI can serve as the basis for content filtering schemes of many flavors, including those aimed at child pornography, unlicensed uses of copyrighted works, or content that an ISP deems to be objectionable. Although government-mandated filtering—which in many cases makes use of

DPI—has garnered significant attention throughout the world, there have also been many instances of ISPs using DPI or contemplating its use for content filtering in the absence of any regulatory mandate (Hunter 2004; Mateus and Peha 2008). DPI is also a tool for ISPs to use in offering content filtering services, such as parental controls, that individual subscribers may elect to apply to their Internet connections, as opposed to ISPs applying a blanket filtering mechanism to all subscribers' traffic.

Lawful intercept—ISPs and DPI vendors frequently cite the interception of communications at the behest of law enforcement or government intelligence officials (Bowman 2009; Henry and Bibic 2009; Werbach 2005). This function is often known as lawful intercept, although as Bendrath (2009) notes, DPI may be used whether a particular government request is lawful or not. Because this use case arises purely as a result of government requirement or coercion, it introduces a host of privacy concerns that the other uses, which result largely from ISPs' own initiative, do not share. Thus, government-requested interception will not be considered in the privacy analysis below.

Although the list above represents the most prevalent use cases, the list is by no means exhaustive. As Reed (2008) has pointed out, DPI constitutes a general-purpose capability whose evolution is directly tied to the evolution of processing power and functionality built into Internet routing and switching devices. As such, its potential uses are as diverse as computing itself, and new uses are likely to continue to be invented at a rapid pace.

The extent to which Internet users perceive the benefits of DPI will vary from use to use and will also depend on whether users already have software or use web-based services that serve the same functions. Proactive security measures and troubleshooting capabilities are likely to be the most easily appreciated uses, especially if Internet users experience reduced levels of spam, malware, or loss of connectivity on the network as a result. For congestion management, prioritized service offerings, and content filtering, users' judgments will likely depend on how their own online activities are affected by the services. If their user experiences are dramatically improved once DPI technology is deployed, they may be more likely to accept such offerings. Conversely, if their experiences worsen as a result of DPI—because their content is targeted for throttling or filtering—they may be less likely to accept them. Congestion management, content filtering, and prioritization are three uses that are likely to be better appreciated as they are deployed as user empowerment tools that users can turn on and off as they wish. Since usage monitoring is primarily conducted for the benefit of ISPs, users may have difficulty perceiving its value. Finally, because many Internet users appear hesitant to embrace behavioral advertising (Turow et al. 2009), that use case may prove to be the least compelling from the user perspective.

Mitigations of Privacy Risk

Although the uses of DPI vary widely, many of the steps that ISPs can take to mitigate the privacy risks of their DPI systems are common to multiple uses of the technology. Among the many privacy protections that data collectors of all sorts have at their disposal, several create particular challenges and opportunities for ISPs using DPI: limiting the depth of inspection, limiting the breadth of inspection, limiting data retention, disclosing the presence of DPI, and offering users choices. This selection by no means exhausts the scope of protections that ISPs could apply, but each tactic provides a useful illustration of how ISPs can reduce some privacy concerns and how mitigating privacy risk may be similar across multiple DPI uses.

Limiting the depth of inspection—As the discussion of the term *deep* alluded to, the depth of DPI—the extent to which DPI tools delve into individual packets—can vary greatly, from the mere inspection of port numbers to the capture and analysis of entire packet payloads. At its core, DPI concerns pattern matching: inspecting packets to determine if the bits they carry matches some predetermined sequence of bits that is of interest to the ISP. The depth of DPI measures how much of each packet is subjected to this pattern matching. ISPs seeking to mitigate privacy risk can limit their inspection to only the depth necessary for the purpose at hand.

Only some DPI uses provide the ISP with a choice of how deep inspection can be. For example, some peer-to-peer protocols declare their names in their application headers, so that when a peer-to-peer file transfer is first initiated, the name of the protocol in use is always visible at the same location inside the packet. When DPI is used to identify these protocols as part of congestion management or a prioritized service offering, the ISP could seek to inspect only the portion of the packet payload where it is possible for peer-to-peer protocol names to appear as opposed to inspecting entire payloads. Likewise, an ISP looking to generate approximate usage data about the amount of email traffic on its network could inspect only port numbers to identify traffic using common email ports, or alternatively it could inspect entire payloads to look for well-known email protocol signatures or "to," "from," and "subject" fields. The ISP's choice will likely depend on a number of factors, including the perceived accuracy of each alternative. If shallower inspection is sufficient for the ISP's purpose and it chooses to do deeper inspection, however, it would be introducing additional privacy risk unnecessarily.

Other uses, such as blocking viruses (a proactive security measure), filtering objectionable content, and behavioral advertising likely necessitate the capture and inspection of entire payloads in order to determine if the traffic on the network matches a virus signature or a content fingerprint that is known to the ISP or contains information that indicates a particular user's

interest. Troubleshooting, because of its investigatory nature, is also likely to be designed to capture entire payloads.

Limiting the breadth of inspection—Breadth measures the volume of packets that are subject to DPI. The possible levels of breadth of DPI are wide-ranging, from just a sample of packets from a small number of an ISP's subscribers to every packet of millions of subscribers. As with depth, ISPs can limit the breadth of inspection to what is necessary for the DPI system's purpose. Cost and performance likely also factor in to decisions about breadth. Although today's DPI equipment is capable of processing vast amounts of data at high speeds, deploying it on a large scale may be costly or may introduce performance losses that ISPs must also take into account.

One way that breadth can be calibrated is by changing the number of an individual user's packets that pass through DPI equipment. When an application running on a user's computer initiates a communication with another computer on the Internet, the sequence of packets exchanged between the two computers is known as a "flow." Each individual user and many applications can sustain multiple flows simultaneously. For many DPI uses, conducting a pattern match on only the first few packets in a flow may provide sufficient information for the ISP to take whatever action is appropriate (Mochalski and Schulze 2009). In the peer-to-peer protocol-matching example from above, the ISP may know that the protocol name appears only in the first packet of a flow between two computers. If the first packet of a flow contains the peer-to-peer protocol name, the ISP can apply its congestion management or prioritization rules to the entire flow without further inspecting its packets. Otherwise, it can ignore the flow altogether. The same logic applies to filtering content of particular types, e.g., images.

Similarly, some networks experience congestion much more frequently in one direction (upstream or downstream), and some content slated for prioritization may travel in only one direction. ISPs can limit their DPI use to the relevant direction only, thereby limiting the number of packets that get inspected. Behavioral advertising presents a similar choice. The behavioral advertising firm Phorm inspects traffic flowing in both directions between a user and a website, collecting data about both the user's website request (the URL and, for searches, the search terms) and the website content sent to the user (Clayton 2008). In contrast, the system used in a trial by NebuAd, another DPI-based firm, captures only URLs and search terms, not the content of web pages (McCullagh 2008), thus requiring inspection in only one direction. Both companies have limited their inspection to web traffic, and, upon identifying a flow destined for certain sensitive sites (web-based email sites, for example), they have exempted the remainder of the flow from inspection (Clayton 2008; Dykes 2008). Thus, the breadth of packets involved

in Phorm's system is far greater than that of the NebuAd system, but both companies have sought to limit breadth in particular ways.

Another breadth-related design decision concerns the total number of subscribers whose packets pass through a DPI engine. ISPs can deploy DPI throughout the network such that all subscribers' traffic is subjected to inspection, or they may selectively deploy it on certain groups of subscribers or at certain network nodes. For usage monitoring and troubleshooting, it may be sufficient to inspect traffic at only a sample of network nodes that give an indication of how subscribers are using the network or to diagnose a specific network problem, rather than monitoring every single subscriber. For other uses, such as behavioral advertising or content filtering, the efficacy of the DPI solution likely increases with each subscriber whose traffic is subjected to it, giving ISPs an incentive to include as many subscribers as possible.

Finally, DPI's breadth depends on whether a DPI system proactively monitors traffic or only in response to other network events or signals. The difference between the proactive security and troubleshooting use cases highlights this distinction. For proactive spam management, for example, an ISP may scan every email message to determine whether each message matches any predefined spam signatures. Conversely, ISPs employing the troubleshooting approach might wait for subscribers or other ISPs to report spam problems, and only then would they use DPI to investigate. For some use cases, such as behavioral advertising and prioritized service offerings, the only reasonable choice is to deploy DPI proactively because the continual monitoring of the network is what allows the ISP to offer the DPI-based service.

Limiting data retention—Only a handful of DPI use cases demand that some individualized (per-user) data be retained beyond a short time, and even in those cases the data need not be raw packet payloads. For behavioral advertising to work, behavioral profiles need to be retained, but DPI systems, as Phorm and NebuAd have done, can draw their conclusions about users' behavior immediately and retain those conclusions without keeping the packet payloads. Proactive security measures may require retaining information about specific traffic patterns coming to or going from certain IP addresses, but again these logs need not necessarily contain all collected packet data.

ISPs may have reasons to retain more detailed individualized records of what their DPI systems are doing, however. These records can help ISPs evaluate the effectiveness of their congestion management, prioritized service offerings, troubleshooting arrangements, or content filters. ISPs need to balance the benefit of storing additional records in individualized form against the privacy risks associated with storing data, which include accidental or malicious disclosure, compelled disclosure to litigants and governments, and internal misuse (Brown and Laurie 2000; Cooper 2008a).

There are also performance trade-offs involved between doing analysis in real time with little data storage and storing packet data for later analysis. Behavioral advertising and usage monitoring, for example, do not necessarily require real-time analysis because behavioral profiles and usage data can be calculated based on stored data. Performing sophisticated analysis in real time on large quantities of traffic can affect the network's performance (Messier 2009). On the other hand, given the vast streams of data that may potentially pass through a single DPI unit, the storage costs associated with retaining the data for any significant length of time may be prohibitive for some ISPs. Both of these factors come into play in cases where ISPs have a choice about when to analyze packet data.

Disclosing the presence of DPI—Disclosing the presence of DPI is one of the most basic steps that ISPs can take to mitigate the privacy impact of DPI. For most uses, the level of detail included in the disclosure can be quite specific, providing details of the circumstances under which packets are inspected, the purpose for conducting the inspection, the data retention policy, the choices that users have with respect to the DPI system, and the mechanism they can use to access the data collected about them. With regard to proactive security, however, ISPs may have cause to refrain from disclosing much so as to prevent attackers from circumventing the system.

ISPs' lack of proximity to their subscribers complicates the task of disclosure (Ohm 2009). While interaction with online service such as search engines or social networks is usually obvious, it is not immediately obvious to Internet users when they are interacting with their ISPs' service infrastructures. Frequent visits to search engines or social networks can create for Internet users a strong sense of their personal relationship or association with those sites. Conversely, their ISPs contact them only infrequently, usually only to send a bill. Despite the fact that an ISP transmits all of a subscriber's online communications, the subscriber is unlikely to feel as though he is directly communicating with the ISP on a frequent basis, whereas it may be much more obvious to Internet users that they directly communicate with the websites they visit.

ISPs have several conventional mechanisms available to them to try to bridge this gap, including their own websites (Bell Canada n.d. provides one example) and written rules and procedures, their printed communications sent to subscribers or potential subscribers, and media campaigns and press reports. ISPs and vendors can also seek innovative ways to raise awareness about DPI more generally. German DPI vendor ipoque, for example, has released an open source version of its DPI engine for bandwidth management (Anderson 2009), showing exactly how the pattern matching works for a wide variety of protocols. Third parties have developed tools that allow

Internet users to detect some of the potential effects of DPI, such as packet modification (Dischinger et al. 2008; Electronic Frontier Foundation 2008). ISPs could incorporate and expand such tools to show their users how DPI is functioning on the network.

Offering users choice about DPI—For some uses of DPI, offering users a choice about whether to have their packet data inspected contravenes the purpose of deploying DPI in the first place. For congestion management or the filtering of unlawful content, for example, offering subscribers a choice about their participation may defeat the purpose of conducting DPI; those who create substantial congestion or transmit illegal content could simply choose not to have their communications subject to inspection. For the offering of prioritized services, ISPs may build terms into their contracts with content providers that require them to prioritize content delivered to all subscribers, thus eliminating their ability to offer users a choice.

For the other DPI uses, however, the idea of offering subscribers a choice is reasonable, and the challenge for ISPs is to facilitate that choice. Conventional mechanisms available to ISPs for offering choice each have their own drawbacks and benefits. ISPs may reach all their subscribers by sending printed materials or including the choice option in the initial contract for Internet service, but it may be difficult to draw subscribers' attention to printed disclosures. Sending a notice and obtaining subscribers' consent via email may also be possible, but many ISPs do not collect their subscribers' primary email addresses. Phone calls may be a good way to engage subscribers directly, but calling millions of subscribers may be prohibitively expensive.

ISPs also can show subscribers an online pop-up or overlay notice to inform them about DPI and what their choices are. There are at least two ways such a notice could work: as a captive screen or walled garden that forces subscribers to make a decision before they can use their Internet connections for other purposes (similar to how some hotels and WiFi hotspots operate), or as a recurring notice that shows up from time to time if subscribers click away from it without making a decision. Opperman (2009) discusses this type of pop-up for a purpose unrelated to notification of DPI. In either case, ISPs will need to account for the fact that the novelty of receiving a pop-up from the ISP may take many subscribers by surprise. Furthermore, the task of deciding when to insert the notice in the midst of subscribers' online activities may itself require some form of DPI. If the ISP wants to display a web-based pop-up, for example, it will need to intercept a user's web request in order to deliver the notice as part of the web response. For web-based notices, at least, it is not immediately obvious how to bypass the conundrum of using DPI to inform subscribers about the use of DPI.

ISPs face additional challenges because many Internet connections are shared by multiple members of a household. To give individuals control over whether their communications will be monitored, ISPs would need to find ways to offer choice separately to different household members and to treat different individuals' traffic differently depending on their individual preferences. For most residential Internet connections, ISPs do not seek to differentiate between users within a household, and doing so may not be technically feasible or could create privacy risks of its own by revealing more to the ISP about which person is associated with each online activity. Thus, for any particular use of DPI, ISPs face the need to weigh the ability of a single individual to choose on behalf of the entire household against the potential that individuals within the household may have differing preferences.

For the choice about a DPI-based service to be meaningful, ISPs would need to differentiate between users who do and do not choose to participate, and the traffic of those who decline should not be subject to inspection above OSI layer 3. Early DPI-based behavioral advertising technologies were configured such that opting out merely discontinued the creation of an advertising profile for the user but did not discontinue the use of DPI on that user's traffic (Clayton 2008; Cooper 2008b). Moreover, given high ISP switching costs, offering a meaningful choice would also require ISPs to provide mechanisms for users to change their preferences over time so they are not locked in to their initial decisions.

Summary of Analysis

There is obviously a great deal of variability in determining exactly which privacy protections can be applied to each kind of DPI deployment, but the analysis above provides some insight into the kinds of steps available to ISPs for each particular use case. The chart in figure 5.3 summarizes this analysis, using a plus sign to indicate a protection that is feasible to apply, a minus sign to indicate a protection that is not likely feasible, and a bullet to indicate a protection whose application is unclear or highly specific to each individual DPI deployment. The three overarching ISP challenges discussed earlier are overlaid across all uses and mitigations.

The summary aptly depicts why individual uses of DPI deserve their own purpose-specific privacy analyses. Proactive security, for example, likely one of the most acceptable uses of DPI to users given its potential to help safeguard their network connections, is also the least amenable to the privacy risk mitigations, with little opportunity to limit retention, disclose the presence of DPI, or offer users choice about its use for security. Behavioral advertising, meanwhile, might raise greater privacy concern for users, but it

		Privacy Risk Mitigation				
		Limiting Depth	Limiting Breadth	Limiting Retention	Disclosing DPI	Offering Choice
Use Case	Usage Monitoring	+	+	+	+	+
	Congestion Management	+	+	+	+	–
	Prioritized Services	+	•	+	+	-
	Troubleshooting	–	+	•	+	+
	Behavioral Advertising	–	•	•	+	+
	Content Filtering	-	•	•	+	-
	Proactive Security	–	-	•	•	•

Figure 5.3. A summary of the application of privacy mitigations to the DPI use cases and the overarching challenges for ISPs using DPI.

also better accommodates mitigations, with clear means available to offer disclosure and choice and some possibility for limiting breadth and retention. The differences among all of the DPI uses demonstrate that a one-size-fits-all assessment of DPI is inadequate to the task of understanding DPI's practical privacy impact.

CONCLUSION

Given the above analysis, the cause for concern over DPI is warranted. As a technology with the capacity to provide targeted insight into Internet communications on a mass scale, it has the potential to dramatically alter the way that people approach the Internet and the trust that they place in the network. Because of ISPs' unique position as network on-ramps, their use of DPI creates singular privacy challenges that may be difficult or impossible to overcome.

The proliferation of DPI is likely to continue, however, as long as the legal landscape remains somewhat flexible, evolving network usage profiles continue to create new challenges for congestion and security, and ISPs continue to feel the pressure to monetize the content crossing their networks. The ability to address the privacy concerns has not yet been fully explored or exploited. As DPI-based systems are designed and developed, ISPs have many tools at their disposal to help mitigate privacy risks. A comprehensive privacy analysis of DPI must explore the details of individual uses and deployments

because the applicability of each mitigation tactic and the benefits of DPI systems that Internet users perceive are highly dependent on the specific use(s) to which DPI is put.

The extent to which public attention will be devoted to DPI-related privacy issues will likely continue to depend on the specific policy issue under discussion. For example, while privacy has been (and will likely continue to be) a core focus for policy makers in analyzing behavioral advertising, the same has not been true for net neutrality, and thus the focus on DPI-related privacy has been stronger in the former context than in the latter. But as long as ISPs continue to use DPI-based solutions, the privacy of Internet users will hang in the balance among ISPs' unique power as Internet gateways, the steps that ISPs take to mitigate privacy risk, and the benefits that users perceive of DPI-based systems.

REFERENCES

Allot Communications. 2007. Digging deeper into deep packet inspection. http://www.allot.com/Common/FilesBinaryWrite.aspx?id=3053.

Anderson, Nate. 2007. Deep packet inspection meets 'Net neutrality, CALEA. ars technica. July 26. http://arstechnica.com/hardware/news/2007/07/Deep-packet-inspection-meets-net-neutrality.ars.

———. 2009. Deep packet inspection engine goes open source. ars technica. September 9. http://arstechnica.com/open-source/news/2009/09/deep-packet-inspection-engine-goes-open-source.ars.

Arbor Networks. 2010. Arbor e100 datasheet. http://www.arbornetworks.com/de/docman/arbor-e100-data-sheet-english/download.html.

Barras, Colin. 2009. Tim Berners-Lee: Internet at risk from "wiretapping." *Computer Weekly*, March 16. http://www.computerweekly.com/Articles/2009/03/16/235279/Tim-Berners-Lee-Internet-at-risk-from-39wiretapping39.htm.

Bell Canada. n.d. Bell: Network management. http://service.sympatico.ca/index.cfm?language=en&method=content.view&content_id=12119.

Bendrath, Ralf. 2009. Global technology trends and national regulation: Explaining variation in the governance of deep packet inspection. International Studies Annual Convention, February 15–18, New York City. http://userpage.fu-berlin.de/~bendrath/Paper_Ralf-Bendrath_DPI_v1-5.pdf.

Bohm, Nicholas. 2008. The Phorm "Webwise" system—A legal analysis. Foundation for Information Policy Research, April 23. http://www.fipr.org/080423phormlegal.pdf.

Bowman, Don. 2009. Sandvine presentation to the Canadian Radio-television and Telecommunications Commission. CRTC Public Notice 2008-19, July 6. http://www.crtc.gc.ca/public/partvii/2008/8646/c12_200815400/1241688.DOC.

Brown, Ian, and Ben Laurie. 2000. Security against compelled disclosure. In *Proceedings of the 16th Annual Computer Security Applications Conference*. New Orleans, LA, USA. http://www.apache-ssl.org/disclosure.pdf.

Casserly, James L., Ryan G. Wallach, Daniel K. Alvarez, Joseph W. Waz, Kathryn A. Zachem, Mary McManus, Thomas R. Nathan, and Gerard J. Lewis. 2008. Comments of Comcast Corporation in the matter of broadband industry practices, WC Docket No. 07-52. February 12. http://fjallfoss.fcc.gov/ecfs/document/view?id=6519840991.

Center for Democracy & Technology. 2007, August. *Search privacy practices: A work in progress*. http://www.cdt.org/privacy/20070808searchprivacy.pdf.

Cho, Kenjiro, Kensuke Fukuda, Hiroshi Esaki, and Akira Kato. 2006. The impact and implications of the growth in residential user-to-user traffic. In *Proceedings of the 2006 conference on applications, technologies, architectures, and protocols for computer communications*, 207–18. Pisa, Italy: ACM. doi:10.1145/1159913.1159938. http://portal.acm.org/citation.cfm?id=1159938&dl=GUIDE&coll=GUIDE&CFID=69766487&CFTOKEN=47276760#.

Clayton, Richard. 2008. The Phorm "Webwise" system. May 18. http://www.cl.cam.ac.uk/~rnc1/080518-phorm.pdf.

Cooper, Alissa. 2008a. A survey of query log privacy-enhancing techniques from a policy perspective. *ACM Transactions on the Web* 2, no. 4: 1–27. doi:10.1145/1409220.1409222. http://portal.acm.org/citation.cfm?id=1409220.1409222.

———. 2008b, July 17. Statement of Alissa Cooper before the House Committee on Energy and Commerce, Subcommittee on Telecommunications and the Internet: What your broadband provider knows about your web use: Deep packet inspection and communications laws and policies. http://energycommerce.house.gov/images/stories/Documents/Hearings/PDF/Testimony/TI/110-ti-hrg.071708.Cooper-testimony.pdf.

Dischinger, Marcel, Alan Mislove, Andreas Haeberlen, and Krishna P. Gummadi. 2008. Detecting bittorrent blocking. In *Proceedings of the 8th ACM SIGCOMM conference on Internet measurement*, 3–8. Vouliagmeni, Greece: ACM. doi:10.1145/1452520.1452523. http://portal.acm.org/citation.cfm?id=1452520.1452523&coll=Portal&dl=GUIDE&CFID=75047289&CFTOKEN=91399611.

Dykes, Bob. 2008, July 17. Summary of testimony of Bob Dykes, CEO NebuAd, Inc. before the House Subcommittee on Telecommunications and the Internet, What your broadband provider knows about your web use: Deep packet inspection and communications laws and policies. http://archives.energycommerce.house.gov/cmte_mtgs/110-ti-hrg.071708.Dykes-testimony.pdf.

Electronic Communications Privacy Act. 1986, October 21. PL 99-508.

Electronic Frontier Foundation. 2008, August 1. EFF releases "Switzerland" ISP testing tool. http://www.eff.org/press/archives/2008/07/31.

Engelhart, Kenneth G. 2009. Response to Interrogatory: Rogers(CRTC)4Dec08-1. CRTC Public Notice 2008-19, January 13. elecom Public Notice CRTC 2008-19. http://www.crtc.gc.ca/public/partvii/2008/8646/c12_200815400/1005723.zip.

European Commission. 2009. Commission launches case against UK over privacy and personal data protection. IP/09/570, April 14. http://europa.eu/rapid/pressReleasesAction.do?reference=IP/09/570&format=HTML&aged=0&language=EN&guiLanguage=en.

Federal Communications Commission. 2008. *Memorandum opinion and order in the matters of Free Press and Public Knowledge against Comcast Corporation for secretly degrading peer-to-peer applications; Broadband industry practices; Petition of Free Press et al. for declaratory ruling that degrading an Internet application violates the FCC's Internet policy statement and does not meet an exception for "reasonable network management."*

Felten, Edward W. 2006, August. Nuts and bolts of network neutrality. *AEI-Brookings Joint Center for Regulatory Studies*, no. 6. http://www.reg-markets.org/admin/authorpdfs/redirect-safely.php?fname=../pdffiles/php9e.pdf.

Finnie, Graham. 2009, January. *ISP traffic management technologies: The state of the art*. Heavy Reading. http://www.crtc.gc.ca/PartVII/eng/2008/8646/isp-fsi.htm#_toc219621630.

Foucault, Michel. 1977. *Discipline and punish: The birth of the prison*, trans. Alan Sheridan. New York: Pantheon Books.

Henry, Denis E., and Mirko Bibic. 2009. Response to Interrogatory: The Companies(CRTC)4Dec08-1 PN 2008-19 Abridged. CRTC Public Notice 2008-19, January 13. Telecom Public Notice CRTC 2008-19. http://www.crtc.gc.ca/public/partvii/2008/8646/c12_200815400/1006810.zip.

Hindman, Matthew. 2008. *The myth of digital democracy*. Princeton, NJ: Princeton University Press.

Horrigan, John. 2009, June 17. *Home broadband adoption 2009*. Pew Internet & American Life Project. http://pewinternet.org/Reports/2009/10-Home-Broadband-Adoption-2009.aspx.

Hunter, Philip. 2004, September. BT's bold pioneering child porn block wins plaudits amid Internet censorship concerns. *Computer Fraud & Security* 2004, no. 9: 4–5. doi:10.1016/S1361-3723(04)00109-5. http://www.sciencedirect.com/science/article/B6VNT-4DCVHTP-7/2/b091ae807e76df39019a04ddeadcfaa8.

International Organization for Standardization. 1996. Information technology—Open Systems Interconnection—Basic reference model: The basic model. ISO/IEC, June 15. http://webstore.iec.ch/preview/info_isoiec10731%7Bed1.0%7Den.pdf.

ipoque. 2008. Datasheet PRX-10G. http://www.ipoque.com/userfiles/file/datasheet-prx10g.pdf.

Johnson, Roger L. 2008, August 8. Knology Letter RE: Internet advertising inquiry. http://markey.house.gov/docs/telecomm/knology.pdf.

Krafft, Jackie, and Evens Salies. 2008, May. The diffusion of ADSL and costs of switching Internet providers in the broadband industry: Evidence from the French case. *Research Policy* 37, no. 4: 706–19. doi:10.1016/j.respol.2008.01.007. http://www.sciencedirect.com/science/article/B6V77-4S2VFTM-1/2/114350a2884bda8e3a889cb706498606.

Lyon, David. 2007. *Surveillance studies: An overview*. Oxford, UK: Polity.

MacDonald, Natalie. 2009. Response to Interrogatory: Bragg (CRTC) 4 December 08-1 PN 2008-19. CRTC Public Notice 2008-19, January 13. Telecom Public Notice CRTC 2008-19. http://www.crtc.gc.ca/public/partvii/2008/8646/c12_200815400/1005749.zip.

Martin, D. Craig. 2008, August 13. WOW! Letter to Hon. John D. Dingell, Hon. Joe Barton, Hon. Edward J. Markey, and Hon. Cliff Stearns. http://markey.house.gov/docs/telecomm/20080808_wow_to_ejm.pdf.

Mateus, Alexandre M., and Jon M. Peha. 2008. Dimensions of P2P and digital piracy in a university campus. In *Proceedings of 2008 Telecommunications Policy Research Conference (TPRC)*. Alexandria, VA. http://digitalcitizen.illinoisstate.edu/press_presentations/documents/mateus-peha-TPRC-paper.pdf.

McCullagh, Declan. 2008. Q&A with Charter VP: Your web activity, logged and loaded. *CNET News*, May 15. http://news.cnet.com/8301-13578_3-9945309-38.html.

Messier, Michel. 2009. CRTC File No: 8646-C12-200815400—Telecom Public Notice CRTC 2008-19, Review of the Internet traffic management practices of Internet service providers—Cogeco reply comments. CRTC Public Notice 2008-19, April 30. http://www.crtc.gc.ca/public/partvii/2008/8646/c12_200815400/1110488.pdf.

Mochalski, Klaus, and Hendrik Schulze. 2009. *Deep packet inspection: Technology, applications & net neutrality*. http://www.ipoque.com/userfiles/file/DPI-Whitepaper.pdf.

NoDPI. 2008. No deep packet inspection FAQ. https://nodpi.org/faq/.

Ohm, Paul. 2009, September. The rise and fall of invasive ISP surveillance. *University of Illinois Law Review* 2009, no. 5: 1417–96. http://lawreview.law.uiuc.edu/publications/2000s/2009/2009_5/Ohm.pdf.

Opperman, Jay. 2009. Security scene: Introducing constant guard. Comcast Voices. October 8. http://blog.comcast.com/2009/10/security-scene-introducing-constant-guard.html.

Parsons, Christopher. 2008, January. Deep packet inspection in perspective: Tracing its lineage and surveillance potentials. *The new transparency: Surveillance and social sorting*. https://qspace.library.queensu.ca/bitstream/1974/1939/1/WP_Deep_Packet_Inspection_Parsons_Jan_2008.pdf.

———. 2009. *Summary of January 13, 2009 CRTC filings by major ISPs in response to Interrogatory PN 2008-19 with February 9, 2009 Updates*. February 13. http://www.christopher-parsons.com/PublicUpload/Summary_of_January_13_2009_ISP_filings_with_February_9_2009_Updates_version_1.0%28for_web%29.pdf.

Post, Glen F. 2008, August 7. CenturyTel letter to Chairman Dingell, ranking member Barton, Chairman Markey, and ranking member Stearns. http://markey.house.gov/docs/telecomm/centurytel_080708.pdf.

Postel, Jon, ed. 1981a. RFC 793, Transmission Control Protocol. Internet Engineering Task Force, September. http://www.rfc-editor.org/rfc/rfc793.txt.

———, ed. 1981b. RFC 791, Internet Protocol. Internet Engineering Task Force, September. http://www.ietf.org/rfc/rfc791.txt.

Reed, David P. 2008. Statement of Dr. David P. Reed to Subcommittee on Telecommunications and the Internet Committee on Energy and Commerce U.S. House of Representatives. July 17. http://energycommerce.house.gov/images/stories/Documents/Hearings/PDF/Testimony/TI/110-ti-hrg.071708.Reed%20-testimony.pdf.

Riley, M. Chris, and Ben Scott. 2009, March. *Deep packet inspection: The end of the Internet as we know it?* Free Press. http://www.freepress.net/files/Deep_Packet_Inspection_The_End_of_the_Internet_As_We_Know_It.pdf.

Turow, Joseph, Jennifer King, Chris Jay Hoofnagle, Amy Bleakley, and Michael Hennessy. 2009. Americans reject tailored advertising and three activities that enable it. *SSRN eLibrary* (September 29). http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1478214.

Verhoeve, Michael. 2009, July 28. Final reply: Telecom Public Notice CRTC 2008-19. http://www.crtc.gc.ca/PartVII/eng/2008/8646/c12_200815400.htm#reldoc.

Werbach, Kevin. 2005. Breaking the ice: Rethinking telecommunications law for the digital age. *Journal on Telecommunications & High Technology Law* 4, no. 1: 59. http://heinonline.org/HOL/Page?handle=hein.journals/jtelhtel4&id=65&div=&collection=journals.

Wilson, Alexandra M., Lauren M. Van Wazer, Grace Koh, John P. Spalding, and Alysia M. Long. 2010, January 14. Comments of Cox Communications, Inc. In the matter of preserving the open Internet; Broadband industry practices. http://fjallfoss.fcc.gov/ecfs/document/view?id=7020378714.