Telecommunications Policy, vol. 38, no. 1, pp. 66-85, Jan. 2014.

Policy Implications of Technology for Detecting P2P and Copyright Violations

Jon M. Peha[1] and Alexandre M. Mateus
Carnegie Mellon University

Abstract

The effectiveness of many proposed policies regarding both online copyright protection and network neutrality depend on the extent to which it is technically possible to detect peer-to-peer file sharing (P2P), the transfer of copyrighted files, or both. There are many detection approaches, some performed primarily by network operators and some by application-layer agents. This paper describes capabilities, limitations, privacy issues, and policy implications of detection technologies and their countermeasures, in part through quantitative analysis of empirical data. Different approaches are better for different purposes. Network operators are well-positioned to estimate how widespread copyright violations are, but application-layer detection from outside entities has important advantages when the purpose is punishment. Detection is also imperfect, so policies should require more transparency regarding how it is done than we see today. It is shown that, although network operators may not detect every transfer, and they typically miss more video than audio, they can identify most individuals who share copyrighted files via P2P after several weeks of monitoring provided that traffic is unencrypted, which is useful for some purposes. However, it is also shown that encryption is already in use, and it effectively prevents network operators from detecting transfers of copyrighted content. Thus, if network operators are held responsible for monitoring illegal file sharing, there is a tension between using detection to identify violators of copyright law for punishment, which may motivate even greater use of encryption, and using detection for other purposes such as creating fair compensation schemes for copyright-holders, warning users that they may be violating copyright law, or allocating network resources. Alternatively, there are forms of detection that are not evaded through encryption, and application-layer agents rather than network operators are primarily responsible for these. These copyright policy issues are intertwined with network neutrality policy in subtle ways. Network neutrality rules do not protect illegal transfers of copyrighted content, but if network operators are responsible for enforcement (as in "graduated response") then regulators must determine when it is reasonable to terminate or degrade service based on allegations of copyright violation given the limitations of detection technology to prove those allegations. Allegations of copyright violation should be considered invalid unless they are accompanied with information about how detection was performed and an opportunity for rebuttal. Such transparency has been routinely lacking in both laws and industry agreements.

Keywords: peer to peer file sharing (P2P); copyright; network neutrality; open Internet; graduated response; deep packet inspection

[1] peha@cmu.edu, www.ece.cmu.edu/~peha/bio.html

1 Introduction

Peer-to-peer (P2P) networks lie at the center of two of the most controversial debates on Internet policy. As a result, there have been policy proposals in recent years that would have the effect of allowing, mandating, or prohibiting the use of technologies that identify P2P traffic on networks and respond in some way. There are two sets of issues involved: copyright protection and network neutrality. P2P traffic has made up almost one third of the world's Internet traffic for the last few years (Cisco, 2012). When demand on limited network resources from all network traffic including P2P becomes too high, quality of service (QOS) is degraded for some or all traffic. This has led some Internet Service Providers (ISPs) and enterprise networks to use technical means to identify P2P traffic, and then reduce P2P's consumption of network resources, thereby improving QOS for applications other than P2P. One particularly high-impact example occurred when a U.S. ISP started identifying and terminating P2P transfers (Peha, 2008; Federal Communications Commission, 2008; Mueller & Asghari, 2012). This case led the U.S. Federal Communications Commission (FCC) to hold network neutrality hearings, and ultimately to produce "Open Internet" rules that prohibit some forms of discrimination based on application (e.g. P2P) or content (FCC, 2010). The other reason that P2P is so controversial is that a large majority of the content transferred in P2P networks globally (Mateus & Peha, 2012b) and in some local environments such as university networks (Mateus & Peha, 2012a) consists of copyrighted material distributed without the authorization of copyright holders. Indeed, around the world, far more copies of copyrighted music and movies are distributed illegally through P2P than legally through sale of physical media such as CDs and DVDs and through legal download sites such as iTunes (Mateus & Peha, 2012b). Calculating the economic impact of these P2P transfers is notoriously difficult and estimates differ greatly. Many prominent representatives of copyright holders have asserted that P2P transfers greatly decrease their profits (Recording Industry Association of America, 2011), and thus their incentive to create new content. Some believe this is the cause of the rapid decline in revenues from CD sales. This has led to a series of controversial calls for changes in public policy and law enforcement. To further complicate matters, some of those proposals must be carefully reconciled with Open Internet rules.

Many of the policies suggested by advocates and researchers to address both network neutrality and online copyright protection are based on assumptions about the extent to which it is technically possible to detect P2P file sharing in general, or to detect the transfer of specific copyrighted files, that are not always valid. For example, some assume that a network operator either knows with certainty whether copyright infringement is occurring or has no knowledge. In the U.S., the Digital Millennium Copyright Act (DMCA) provides an ISP with safe harbor protection from liability for copyright infringement that occurs on its network, but this protection is lost if the provider has knowledge of the infringement (US Congress, 1998). On the other hand, the FCC determined in its Open Internet proceeding that a fixed broadband Internet access service provider is not allowed to block lawful network traffic on the grounds that it was

produced by an application such as P2P that is generating a large amount of traffic, and the provider is not allowed to block content that is transferred lawfully, but the provider is free to block a transfer that violates copyright law (FCC, 2010). These provisions become difficult to apply in a world where technology can shed some light on whether copyright infringement is occurring, but *not always with certainty*. Given the levels of certainty that technology can realistically provide, when does an ISP know enough to block P2P on the basis of copyright law without running afoul of net neutrality regulations, and when does an ISP have sufficient knowledge to be held liable for copyright violations? (The matter is further complicated in cases where the provider has financial incentive to block content, perhaps because that content competes with the provider's own offerings.) As another example, the DMCA requires ISPs to terminate service for repeat infringers if the ISPs are to avoid liability (Frieden, 2008), but according to Bridy (2010), "wrongful terminations might themselves create the potential for provider liability to customers for breach of contract," which creates a dilemma for ISPs when knowledge is less than 100% certain. For these and many other current and proposed policies in the realms of copyright protection and network neutrality, policymakers and courts should consider just how effective detection technology can be.

For copyright enforcement, policymakers should also consider who is best positioned to use effective tools. There has been a loud debate over whether ISPs should be required to detect violations on their network, or whether this task should fall to copyright-holders, or to law enforcement. Some even argue that ISPs should be required to prevent copyright infringements by blocking transfers rather than just detecting them, or be held accountable for not doing so (Mueller, Kuehn and Santoso, 2012; Horten 2011). The viability of such policies depends on what current technology can realistically do.

Not only should policy choices depend on the effectiveness of technology, but the effectiveness of technology depends on policy choices. For example, a policy that holds ISPs liable for illegal activities of their customers only if the ISP has sufficient knowledge may deter ISPs from adopting technology that provides any knowledge (Frieden, 2008). Similarly, whether or not consumers adopt encryption to evade detection may depend on whether there are policies that punish observed transfers of copyrighted material or network practices that degrade quality of service for P2P (Lehr, Sirbu, Gillett, & Peha, 2007). Thus, to understand what can be known about use of P2P and the exchange of copyrighted content, we must consider the tools available for monitoring, the tools available to evade monitoring, and a wide range of policies that affect incentive to use them both

The goal of this paper is to describe and categorize the range of detection technologies available, qualitatively and quantitatively assess what they can realistically do, and explore the policy implications of the capabilities and limitations of these technologies. The technologies considered can (i) detect P2P transfers and individuals who perform them, (ii) identify the specific content being shared via P2P and

whether it is copyrighted, and (iii) use P2P file-sharing despite the technologies above without being detected or without revealing what content is being shared.

These technologies are quantitatively assessed in a number of scenarios using empirical data that was collected in the Digital Citizen Project at Illinois State University (ISU) (Illinois State University, 2006) during three collection periods in 2007-8 which total 111 days. ISU created this project to advance innovative and comprehensive ways of dealing with P2P downloads of copyrighted content. (See appendix and (Mateus & Peha, 2012a) for more detail about data collection and analysis.)

A P2P detection technology cannot be evaluated without considering the intended purpose; a technology that is effective for one purpose might be inadequate for another, and vice versa. In the realm of copyright, one use of detection is as prelude to punishment, perhaps through civil proceedings initiated by copyright-holders, as occurred in the U.S. under the Digital Millennium Copyright Act (DMCA) (U.S. Congress, 1998), Alternatively, ISPs may play a more direct role in punishment by terminating or substantially degrading Internet service of customers who are thought to be guilty of too many violations (or "strikes"). This is known as *graduated response,* whereby Internet users that are accused of copyright infringement online receive notices from their ISP, and the consequences increase with the number of notices. Typically, the first notice is just a warning, and the last may come with a termination of Internet service or a serious degradation of service. France was a prominent example with their HADOPI law (*Haute Autorité pour la diffusion des œuvres et la protection des droits sur internet*), and other nations followed (e.g. South Korea, United Kingdom, New Zealand. However, after legal challenges and intense debate, there has been a shift towards graduated response policies established through stakeholder agreements rather than laws (Mueller, Kuehn, & Santoso, 2012; Horten 2011). The U.S. is on its way to a private-sector version; a recent voluntary agreement (Center for Copyright Information , 2011) between a number of major copyright-holder representatives and ISPs would lead participating ISPs to degrade the service of customers who are found to be repeat offenders. A non-governmental organization created through this agreement rather than government would be responsible for ensuring justice.

As either an alternative or a complement to punishment, network operators may try to prevent infringement by blocking transfers that would violate copyright law. This use of detection may be less demanding in that it is not necessary to correctly identify who is responsible for the transfer, but a high degree of certainty is still required that the transfer would violate copyright law or this approach will suppress the legal exchange of ideas. As will be discussed, this is difficult to achieve.

Another related use of detection is to warn users that they *may* be violating copyright rather than to punish them. For example, some U.S. universities have adopted this approach in the hope that a warning alone (without punishment or threat of punishment) will deter copyright violations. This could be particularly useful for students from other countries, who may be unaware of the legal issues and risks.

Indeed, the Higher Education Opportunity Act (HEOA) requires U.S. universities to warn their students (U.S. Congress, 2008; Mueller, Kuehn, & Santoso, 2012; Storch & Wachs, 2011), although not necessarily in conjunction with monitoring. (HEOA requires universities that operate campus enterprise networks to take other actions as well, and penalizes them if they do not.) Uses for both punishment and warnings are similar in that they require the technical ability to identify specific individuals, but the occasional false accusation would be a serious problem for the former and not the latter. For a similar reason, punishment requires reliable identification of content to show that it is copyrighted, whereas warnings could be issued to P2P users even when there is uncertainty over whether the content is copyrighted.

A fourth and very different use of detection is to determine which content is being transferred to compensate copyright-holders. There have been a number of proposals to create compulsory licenses for large classes of content (Lincoff, 1995). Copyright holders could no longer prohibit access to their content, but they would be compensated from a common pool of funds in proportion to the number of times their content is downloaded. Such a fund could come from a variety of sources, including general taxes (Fisher, 2004), broadband fees (Page & Touve. 2010), voluntary contributions (Electronic Frontier Foundation, 2008), or fines for copyright violations. In this case, the detection mechanism must identify content with sufficient accuracy to make compensation fair, but the mechanism need not identify the individuals involved. Information on which content is most popular clearly has other uses as well, such as for the marketing of digital content.

Finally, detection is useful simply to find out whether copyright violations are common in a given context, and therefore to determine whether a response is needed. For example, national policymakers benefit from knowing how prevalent this is within their borders, and university administrators benefit from knowing how prevalent it is on their campuses. For this purpose, it is important to know the number of files transferred, and perhaps to differentiate content type (e.g. movies versus software), but little more.

Detection for resource allocation is simpler in some ways than detection for the punishment of copyright violations or for compensation of copyright-holders, except that it must be done in real time which can be challenging. The goal is typically to protect quality of service of some favored traffic during periods of congestion, so the mechanisms only have to identify which traffic comes from the presumably less favored P2P; it does not matter which specific content is being transferred, or whether that content is copyrighted.

Although this paper focuses on the effectiveness of detection technology and its policy implications, there are other factors that must be considered before deciding on the suitability of any of these technologies. These include the privacy issues that come with any form of observation, total system cost, and the magnitude of the problem that detection is supposed to solve. If detection leads to punishment, as some

propose for copyright enforcement, then there is also need for a just system of due process for contested cases. Moreover, justice must be achieved with little overhead, because if the process itself is too time-consuming or costly, then this alone may deter people from sharing content legally out of fear of being falsely accused, or it may deter people from creating content out of fear that the cost of defending their copyright will make that copyright worthless. These issues are important, but some aspects are outside the scope of this paper.

Moreover, this paper focuses on P2P even though not all online copyright violations occur through P2P. There are sites that illegally stream copyrighted videos, and "cyberlockers" where users can upload and download files, many of which may be copyrighted. However, P2P is particularly challenging from both legal and technical perspectives, and has therefore been the focus of a great deal of attention. In a P2P network, there is no need for a centralized entity that is in charge, as there is with a cyberlocker or streaming site. That means there is no central IT staff that can check whether content is copyrighted when that content is uploaded, or be legally held accountable for not doing so. It also means that detecting which content is being shared and by whom is a more challenging technical problem.

Technology monitors the transfer of files. The question of when these transfers violate copyright law is addressed in Section 2. Section 3 describes and categorizes technologies that can detect P2P and/or transfers of copyrighted content. Section 4 focuses on use of these technologies for resource allocation, quantitatively demonstrates some benefits, and qualitatively describes some dangers and network neutrality implications. Section 5 quantitatively demonstrates strengths and limitations of deep packet inspection. Section 6 shows growing use of encryption, and discusses its implications for detection technologies. Section 7 discusses legal implications of the distinctions between detection technologies for privacy and safe harbor protections. Conclusions and policy implications are presented in Section 8.

2 When does a transfer constitute a copyright violation?

As the next section will discuss, technology can often determine when copyrighted works are being transferred. Many such transfers do violate copyright law, but not all. This section will review enough copyright law to determine which relevant factors cannot be shown through technology, which can, and under what circumstances.

Consider the case where technology has determined that a clip from a copyrighted song has been transferred over the Internet. Further assume that the song's copyright is still valid, and that the rightful owner wants that copyright enforced wherever possible, The transfer is still legal if the transferring party is the copyright-owner or has obtained rights from the copyright-holder. For example, a consumer may purchase a song legally and then transfer it to a remote server for storage (although this would be an unusual use of today's file-sharing software). The transfer is also legal if use of this song or at least the clip that has been observed falls within the bounds of fair use. Fair use provisions vary from country to

country, which makes it even more challenging for software to incorporate fair use concerns, as observed in a recent European Union ruling (European Union, 2011). We focus here on U.S. law as it applies to digital media (Samuelson, 1994). The Copyright Act of 1976 establishes the right to use copyrighted material without permission "for purposes such as criticism, comment, news reporting, teaching (including multiple copies for classroom use), scholarship, or research" under certain conditions, and identifies four factors to consider. The first factor is the nature and character of the use, including whether it is for commercial purposes or nonprofit educational purposes, and whether it is transformative (and therefore a creative work that benefits the public) or merely derivative. For example, the critique of a copyrighted work that includes portions of the original may be considered transformative and therefore a new creative work, unless it contains so much of the original that it supersedes the original. (Folsom v. Marsh, 1841). The second factor is the nature of the copyrighted work, including whether limiting access to that content serves the public interest. A third factor is the amount of the copyrighted work that was used, such as whether a small insubstantial excerpt was copied or the entire work. A fourth factor is the economic impact on the copyrighted work.

Some factors pose greater challenges for detection technology than others. The easiest factors are those used to determine whether a given work should be protected based on the characteristics of the work itself; the technology simply need not report the detection of material that does not warrant copyright protection. There are other factors that can be accommodated relatively easily if detection technology can examine the entire work being transferred, but can be difficult when only a portion is available. For example, if a movie is transferred in its entirety, there can be no fair use defense that this is a transformative work that exhibits some creativity of its own, nor can it be a new work that contains only an insubstantial excerpt from the original. However, such arguments are problematic if technology can only identify one scene from the copyrighted material, and that is often the case when detecting P2P because files are typically transferred in many pieces, and sometimes over multiple paths. Other factors are based not on what is being transferred but on who is doing the transferring and why. For example, is this transfer for commercial purposes or for nonprofit educational purposes? Is the responsible individual someone who has obtained rights from the copyright holder? As the next section will describe, the importance of this kind of context is more problematic for some monitoring technologies than others. Finally, some factors are simply beyond the capability of technology alone. For example, technology cannot determine whether the amount of material reproduced from a copyrighted work is "substantial." According to law professor Rob Frieden, "copyright law, although carefully worded, simply cannot be expressed in the kind of algorithmic language that is required of computer programs to automate functionality like printing or copying. This is especially true of 'fair use.'" (Frieden, 2008)

3 Detection technologies and their capabilities

This section discusses technologies that can detect P2P regardless of content, and technologies that can detect transfers of specific copyrighted material. It will describe how the technologies work, and discuss their effectiveness and their susceptibility to accidental error and deliberate countermeasure.

Others have also looked at detection technologies in a policy context. Peha (2007) characterized different types of detection technologies available to commercial ISPs, and then discussed their implications for network neutrality policy. Finnie (2009) described traffic management approaches available to ISPs for the Canadian Radio Television and Telecommunications Commission, including a section on detection. Parsons (2008) presented a different way of categorizing various detection technologies to conclude that one technology (Deep Packet Inspection) should be viewed as a surveillance technology with broad policy implications. Fuchs (2012) provides a useful survey of Deep Packet Inspection products using the claims of their developers, before discussing how these products have been or could be applied and the policy implications. Unlike the above, this paper will consider a broader range of detection technologies, including those not used by network operators, it will explicitly address effectiveness in the face of countermeasures, and it will include quantitative results.

The ability to detect copyrighted content is best quantified with each *Detected Attempt to Transfer Copyrighted Media* (DATCoM), which the authors have previously defined (Mateus & Peha, 2012a) as a communication event that corresponds to a transfer or attempted transfer of content identified as being protected by copyright. Note that not every DATCoM violates copyright law. As discussed in Section 2, the person transferring the file could have a legal right to do so, or the transfer could comply with fair use provisions. Moreover, some technologies cannot distinguish between the successful transfer of an entire copyrighted file and an incomplete and therefore unsuccessful transfer, so neither does the DATCoM measure, but this distinction may matter from a legal perspective.

Fundamentally, detection can take place at either the application layer or at the network layer, and these two approaches demand different technologies and different public policies. Network-layer detection requires greater involvement of the network operator, so its use may lead to new policies for ISPs, corporate enterprise networks, or university enterprise networks. Application-layer detection can be performed from anywhere on the Internet, with little involvement from network operators. Thus, a debate over technology is in part a debate over who is responsible for detection.

3.1 Application-layer detection

For years, copyright-holders in the music industry have identified individuals engaging in copyright infringement using the application-layer technique of *swarm infiltration*. Agents use a modified P2P client to connect to a P2P network as a regular user and collect information about other users sharing particular titles, including their IP addresses. Thus, network operators are not involved in the costly and arduous

task of monitoring. Network operators still have a lesser role. These agents can periodically enlist the help of network operators to map the IP addresses of apparent infringers to the offline identities of these individuals, or even to contact these individuals directly. There is some cost; network operators must keep records of IP address assignments so they can later determine who had a given IP address at the precise time of a DATCoM.

Swarm infiltration is often effective, but has limitations. Agents can easily infiltrate swarms that are openly advertised (e.g. on public trackers), but some swarms are missed. In general, false accusations of DATCoMs are avoided as long as there are no errors in each of two steps: determining whether a transferred file is copyrighted, and determining which individual is associated with a DATCoM. Through swarm infiltration, an agent can learn the content that is offered by other computers and the IP addresses listed for those computers. However, the listed IP addresses can be wrong. Indeed, it is relatively easy for an attacker to deliberately insert the IP address of an innocent person (Piatek, Kohno, & Krishnamurthy, 2008). Moreover, the computer might not actually be offering the copyrighted content that it is advertising. These problems are not insurmountable; it is possible to verify both IP address and content by actually downloading the content from the listed IP address rather than depending on the available metadata, thereby bringing greater reliability to the system. However, this verification takes time, so there is incentive for agents to skip this step, and there is evidence that some have done so in the past, which leads to false accusations (Freedman, 2009; Piatek, Kohno, & Krishnamurthy,2008). The other technical challenge is in ensuring that the IP address is correctly mapped from IP address to individual, particularly in systems where IP addresses are frequently reassigned to different individuals. While some end users have static IP addresses which they keep indefinitely, other users have dynamic IP addresses which are assigned every time the user reengages the network. Again, the technical problem is not insurmountable, but there are opportunities for error. Thus, if detection is combined with punishment, then detection procedures must be sufficiently transparent for the reliability of the outcome to be assessed. Otherwise, it is impossible to know the risk of false accusations.

Users can deliberately try to evade detection by undermining one or both of these steps. One approach is to obfuscate identity, perhaps using onion routing (Dingledine, Mathewson, & Syverson, 2004) which can conceal IP address, or a virtual private network (VPN) service that prevents mapping of IP address to identity (Ernesto, 2008; Ernesto 2009; Cheng, 2009). (For a modest fee, some VPN providers aggregate traffic from multiple users in a pool of IP addresses, and choose not to record data that would allow them to map users to IP addresses if asked.) Alternatively, users can try to prevent infiltration of their swarms by only granting access to trusted parties. Obviously, this limits their ability to share illegally on a large scale.

One advantage of an application-layer method like swarm infiltration is that the agent doing the monitoring is a party to the transfer, which means the agent knows something about the context of the

transfer. As described in Section 2, context can sometimes determine legality. Many (but not all) fair use claims can be assessed. The agent knows that this transfer is not a teacher sending content only to a handful of students enrolled in his class for nonprofit educational purposes, or the legitimate owner of a copyrighted movie backing up her files.

Other application-layer detection techniques may be used to characterize the range of content that is available and/or transferred using P2P. Some studies have scanned the large public trackers that provide information about the content available via P2P and where to find it (Felton, 2010; Layton & Watters, 2010; Envisional, 2011; Mateus & Peha, 2012b). This can show what content is available, although it does not directly measure the extent to which content is shared. A new technique (Mateus & Peha, 2012b) has been introduced that can estimate the number of times per day that each piece of content known to a public tracker (such as a song or a movie or an executable piece of software) is actually transferred. Estimates are more accurate for larger files than smaller files (so this method was not used for files smaller than 1 KB). Unless they are supplemented with swarm infiltration, these techniques are not useful for identifying specific individuals who are sharing content, but they could be useful for characterizing which content is being shared and how much, perhaps in support of some kind of compensation fund for copyright-holders.

3.2 Network-layer detection

The alternative to application-layer detection is network-layer detection, wherein network traffic is monitored and analyzed to determine whether P2P traffic exists, and possibly whether that P2P traffic carries copyrighted content. There is growing interest in network-layer detection, fueled by new products especially for enterprise networks, and by a variety of policy proposals that would require commercial ISPs and university networks to more actively seek and respond to DATCoMs or face legal consequences. Network-layer techniques are based on observations of streams of packets as they flow through a network. This means collection devices must be placed in sensitive points within the network, either to monitor network traffic in real time, or to divert traffic elsewhere where it can be analyzed, perhaps at a later time. Thus, unlike application-layer detection, network operators must be heavily involved in the deployment and operation of this technology, although these network operators may or may not be responsible for the data analysis as well. Similarly, network operators must be heavily involved when placing wiretaps for law enforcement, even though it is law enforcement agents who analyze content. (In some circumstances, others can use network-layer techniques as well. Anyone who uses wi-fi in their home may be granting neighbors that opportunity, but this is obviously no way to monitor millions of households.)

From network–layer information that Internet Protocol (IP) networks were initially designed to use, the most a network operator can ever tell about the likelihood of copyright violation is whether a user is communicating with a suspicious address, such as that of a public tracker known to be associated with

illegal transfers. It is impossible to tell whether anyone is sharing files via P2P, or whether transmitted content is copyrighted, from just the control information at the head of each IP packet. Thus, some propose to look deeper into the actual payload of a packet, which is known as *deep packet inspection* (DPI). DPI has strong supporters and strong detractors, but like many technologies, its specific uses can be either beneficial or detrimental. For example, using DPI it is possible to determine whether a particular communication session is carrying malware, and thereby help prevent infection. It is also possible to determine that a user is uploading to a controversial blog, and then censor that user (Peha, 2007).

There are products on the market that use DPI to identify P2P traffic as a means of limiting resources allocated to P2P, and products that use DPI to address copyright violations. The latter examine the specific content being transferred, compare that content to items in a database of copyrighted content, and interpret matches as DATCoMs. (Sections 5 and 6 will present empirical results derived from use of some of these products.)

As with all detection techniques, false accusations of DATCoMs are avoided as long as content is correctly identified as copyrighted, and the right individual is identified. The issues of mapping IP address to identity are the same as with swarm infiltration discussed above. However, identifying content as copyrighted is generally more challenging with network-layer approaches than with application-layer approaches. For example, when someone uploads a movie to a website like YouTube, that website can look at the entire file before deciding whether it is copyrighted, and this can be done with a high degree of accuracy. In contrast, with the network-layer DPI approach, determinations must typically be made with only a fraction of the file. Matching techniques to identify content can be tuned to be more or less conservative, i.e. more prone to false positives or false negatives, but even when special care is taken to avoid false positives, some challenges persist. As discussed in Section 2, what if that captured piece of video does include part of a copyrighted movie, but the file being transferred is a parody that uses only one scene? This could be counted as a DATCoM, although under U.S. law, such a parody is not covered under the original copyright. (Posner, 1992) Also, network-layer detection lacks the context available in application-layer detection, making it harder to make other fair use distinctions.

For those who would deliberately conceal DATCoMs from DPI, the approach is vulnerable to a simple countermeasure: encryption. If copyrighted content is encrypted before it is transferred, a network-layer approach cannot possibly identify the content (without breaking the encryption), whereas an application-layer approach is not affected by encryption. Most P2P users have clients that already support encryption, although many have not turned encryption on.

While only DPI can identify specific content and determine whether that content is copyrighted, there are other network-layer detection techniques to identify application, even when encryption is used. There are many indicators of application to consider for individual communications sessions, including packet size,

time between successive packets, and session duration. For example, a 50 kb/s bi-directional stream of small packets generated at a fairly constant rate and exchanged with a single IP address over a period of minutes is probably a VOIP conversation (Peha, 2007). Moreover, it is possible to look at many communications sessions from the same device. Previous work has yielded a variety of behavioral patterns that can be used to detect unencrypted BitTorrent (Collins & Reiter, 2006), In the course of this research the authors have found similar tests that are promising for detecting encrypted BitTorrent. For example, if in a single hour a host communicates with five or more other hosts at data rates under 15 kb/s, and this host communicates with four or more hosts with which connections failed at least 28% of the time, then this host is likely to be engaging in P2P. All of this information is apparent without DPI even for encrypted traffic.

The accuracy of these behavioral techniques is uncertain, and will probably change over time. This is particularly true in commercial ISPs as opposed to enterprise networks, because operators of commercial ISPs are less able to know which applications are generating traffic on their network. For example, a given technique might be highly accurate for years at identifying P2P traffic, until a new application emerges that coincidentally generates traffic that the algorithm does not distinguish from P2P. Thus, if detection were used for resource allocation, the emergence of this new application may cause the system to begin allocating resources poorly. It is also possible that application designers would try to generate traffic that such an algorithm would mischaracterize if there were sufficient incentive (e.g. if QoS would improve). The extent to which encrypted P2P could be detected in the face of active countermeasures is unknown.

Finally, TCP (Transport Control Protocol) and UDP (User Datagram Protocol) packets include a field called “port number” which is indicative of application. Many of the older clients for BitTorrent, which is the dominant P2P protocol, use port number 6881. Consequently, network-layer detection of this traffic is easy. Later versions of the client tend to choose port number randomly, to make detection more difficult.

4 Dangers and benefits of discrimination by application

Some enterprise networks currently use network-layer detection to identify P2P traffic, along with mechanisms that degrade quality of service (QOS) for P2P traffic. They do this to maintain adequate QOS for other applications. For example, the Illinois State University (ISU) network used this approach, and Fig. 1 demonstrates why. As shown in Fig. 1.a, twice as many people used P2P in the hours when other demands on network resources were high as compared to the early morning. Presumably because network administrators believe P2P is more likely to be used for recreation rather than the core mission of a university, and because users of P2P were likely to be less disturbed by poor QOS than users of applications like web browsing and voice over Internet protocol (VOIP), ISU limited the capacity available for P2P. For example, in Spring 2007 P2P traffic was allowed to use up to 5% of total capacity at a time, whereas traffic from web browsing (HTTP), email and other applications considered crucial for the

functioning of the university was allowed up to 45%. The result, as shown in Fig. 1.b, is that total P2P data rates did not increase dramatically even when the number of P2P users doubled, which protected the QOS of other traffic. Moreover, the data rate per P2P user was greatest in off-peak hours, which may have motivated some students to shift their P2P to late-night.

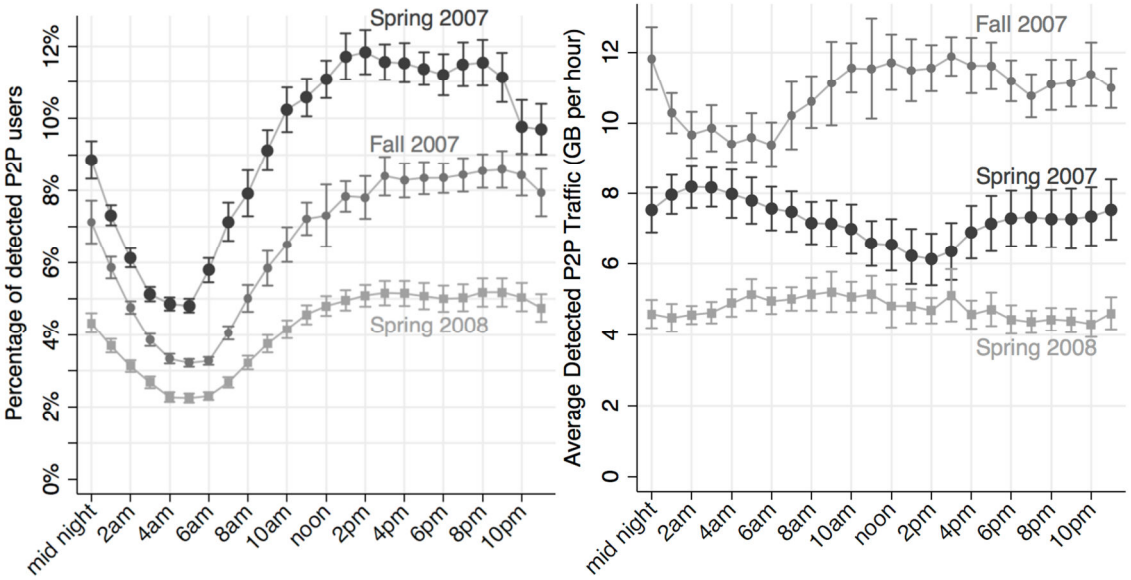

a. Percentage Detected P2P users b. Detected P2P traffic (GB)

Fig. 1. Breakdown by hour of day of P2P activity detected in each period. a) Hourly percentage of P2P users detected out of users living on campus. b) Hourly average GB of detected P2P traffic.

Some commercial ISPs have used similar network-layer detection to identify applications and adjust QOS. This is more controversial, and deservedly so. Section 1 mentioned the case of a U.S. ISP that blocked P2P. In 2007, it was discovered that Comcast was observing network traffic, identifying computers that it believed to be running P2P applications, and secretly terminating those connections in a way that users would probably attribute to their own hardware or software (Peha, 2008). Also noteworthy for our discussion of P2P detection technology is that Comcast allegedly blocked traffic from Lotus Notes, probably by misidentifying this as P2P traffic. This episode led to a class action suit against Comcast by its customers, which was settled for $16 million. It also led the FCC to rule against Comcast in an adjudicatory proceeding (FCC, 2008; Mueller and Asghari, 2012). After questions were raised about the FCC's authority and its adherence to administrative procedures in this case, the FCC decision was overturned by a federal appeals court (U.S. Court of Appeals, 2010; Mueller & Asghari, 2012). Part of the problem in this particular case was that subscribers were not informed. Consider the consumer who chose this ISP over its competitor and paid for premium service specifically to speed up her P2P transfers, not knowing that the ISP was surreptitiously terminating those transfers. That particular problem can be solved relatively easily with good transparency, but some important policy issues remain even when ISPs disclose everything. Another U.S. ISP, Cox Communications, ran a trial in 2009 in which it adjusted QOS based on its understanding of its users' QOS preferences (e.g. VOIP gets better QOS

than P2P), and this was explicitly disclosed to customers. Nevertheless, the trial drew some public criticism, and the company eventually chose to abandon the approach (Baumgartner, 2009). One of the issues raised in the public discussion over these trails (and in this paper) was whether technology would allow ISP to reliably determine which traffic deserves prioritization. Others raised issues about an ISP's incentives when using such technology.

Commercial networks and their customers benefit from this form of resource allocation in the same way enterprise networks do, but there are important differences. First, whether they serve universities, corporations, or non-profit organizations, enterprise networks serve the mission of a specific enterprise, and centralized technical staff can consider the extent to which a given application advances that mission. In contrast, with a commercial ISP, it is the countless missions of individual customers that matter, and the ISP's technical staff cannot always know what customers prefer. Moreover, as discussed in Section 3, the applications that customers consider important changes over time.

In a highly competitive market, ISPs might compete to comply with consumer preferences, perhaps solving or at least greatly limiting the problem above. The bigger problem occurs where ISPs have significant market power (e.g. when an ISP is the sole provider of 4 Mb/s or better Internet service in a region). Instead of prioritizing to maximize benefit to customers, such an ISP may choose to maximize oligopoly rents (Peha, 2007). For example, an ISP that also offers pay-per-view video service could have financial incentive to degrade QOS for all competing sources of video content, including P2P file sharing and streaming. It is less well understood that even if this ISP does not offer its own video, the ISP could degrade QOS for all video unless the customer or content provider pays extra, thereby boosting the price of video to its monopoly level even though the market is competitive (Peha, 2007). Thus, policies that give ISPs and their customers the benefits that are available to enterprise networks may unfortunately come with a risk that significant market power will lead to harm to consumers.

It is therefore reasonable that U.S. Open Internet rules do not apply to closed enterprise networks, and do apply to ISPs that offer Internet access service to the public (FCC, 2010). Thus, commercial ISPs may someday be subject to Open Internet rules that limit their ability to discriminate based on application or content, and copyright protection rules that require them to discriminate based on application or content. . As will be discussed further in Section 8 (on policy implications), these rules must be reconciled with each other, and with what technology is capable of doing with respect to detection and discrimination.

5 Effectiveness of DPI

Prior work has quantified how effective some of the leading DPI products are at identifying P2P through controlled laboratory experiments (Rossenhovel, 2008). Results were found to depend greatly on many factors, including which P2P protocol is in use, extent of IP fragmentation, and use of encryption. At fairly

high data rates but what otherwise might be considered optimistic circumstances, i.e. no encryption or fragmentation, 90% of P2P sessions were detected for the two most popular P2P protocols.

This section goes further by quantifying the effectiveness of DPI at identifying copyrighted content that is shared via P2P and the individuals involved. This is achieved using extensive (and carefully anonymized) data gathered from a college network during the Digital Citizen Project (Illinois State University, 2006). An advantage of examining an actual network as opposed to a laboratory experiment it that it is possible to consider user behavior in some measures of effectiveness; a disadvantage is that measured results cannot be compared directly with correct results, since the correct results are not known.

Effectiveness was assessed in part by comparing results with two different approaches. In one, DPI in a commercial product is used to obtain the content being transferred, so that it can be compared against a database of copyrighted songs, movies, and TV shows constructed for the purpose of copyright protection. Each match is a DATCoM, and the specific title of the associated song, movie, or TV show is known. (This particular database does not include adult content, software, books, or other types of content, although this limitation is not inherent in the technical approach.) In the other approach, self-describing metadata about the file being transferred is used, such as filename, which is routinely carried in P2P sessions. In most cases, the transferred content can be classified using this metadata, primarily by comparing metadata to the music and video titles advertised on amazon.com.[2] The metadata approach is not reliable enough for the resulting matches to be considered true DATCoMs, because there is no guarantee that the metadata accurately describes the actual contents. However, since today's most popular P2P networks include some type of content rating system, mislabeled files are likely to be unpopular, and metadata is a reasonable proxy for content for statistical purposes. In this analysis, it is only assumed that the accuracy of metadata is roughly the same for songs, movies, and TV shows. Over the three monitoring periods combined, over 36 thousand distinct media titles were detected in DATCoM, and over 100 thousand distinct filenames in metadata,

As discussed in Section 3, reliably identifying specific content as the packets flow through a monitoring point is a complex task. Section 5.1 shows DPI's ability to identify the content being transferred, regardless of the individual behind that transfer. Section 5.2 shows DPI's ability to identify users who engage in P2P and who share copyrighted content using P2P, which are referred to as P2P users and DATCoM users, respectively.

[2] The classification process was done automatically for most titles identified via DPI using information collected from Amazon.com's catalog. For metadata filenames, the filename extension was used to infer the content type, and further classification was performed using the same Amazon.com source, as well as other sources (e.g. catalogs of adult content studios). The automatic process classified most of the nearly 140 thousand titles and filenames, but a few thousand titles and filenames that could not be automatically classified were handled manually using the authors' best judgment.

5.1 Detecting audio versus detecting video

This section examines the ability of a state-or-the-art detection product to identify two types of content, audio and video, and finds that effectiveness differs considerably from one content type to another. This is problematic if the goal is to characterize the extent to which different titles or content types are transferred, as ease of detection can distort that picture of relative popularity. The effectiveness of the DPI approach that yields DATCoMs is assessed through comparison to the metadata approach described above. The metadata need not be perfectly accurate for this analysis, as long as metadata accuracy is roughly the same for songs, movies, and TV shows.

Table 1 shows the breakdown of content types observed for both approaches on the portion of the ISU campus network that serves residential dormitories. Nearly all DATCoMs are songs. The absence of software and other content types is expected, since these were not in the database, but the fraction of movies and TV shows is surprisingly small. Similarly, most P2P users were observed with DATCoMs for songs, while relatively few were observed with movies or TV shows, as shown in Fig. 2. However, this is not the case with metadata, which appears to show far more video transfers.

	Titles in DATCoM n = 36,313	**Filenames in metadata** n = 101,879
Unclassified		7.7%
Song	99.5%	66.0%
Album		2.3%
Movie	0.4%	3.9%
TV Show	0.2%	3.0%
Adult / Software / Books / Pictures		17.0%

Table 1. Percentage of copyrighted titles detected in DATCoM and of filenames detected in Metadata for each type of content (columns add up to 100%).

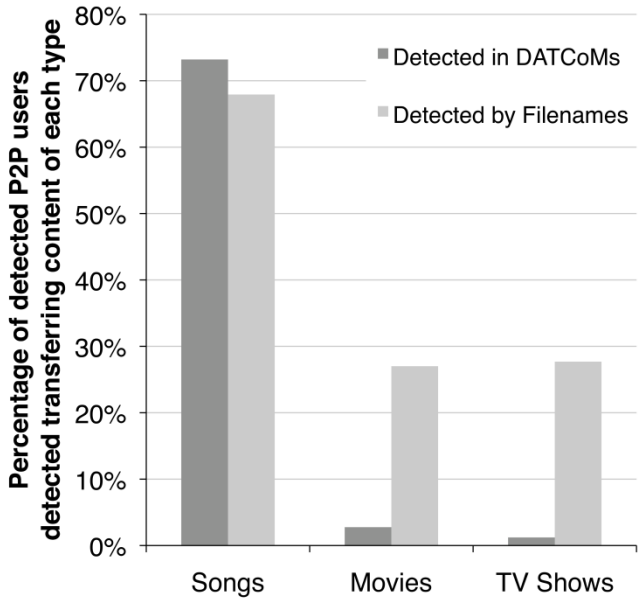

Fig. 2. Average for the three monitoring periods of the percentage of detected P2P users detected transferring songs, movies or TV shows by means of DATCoMs and by means of filenames.

It is conceivable that the differences shown in Fig. 2 are the result of a database that is more comprehensive for songs than for movies and TV shows. However, it can be shown that this is not the case by confining the analysis to the 100 songs and 100 movies with the most DATCoMs. Fig. 3 shows the average percentage of P2P users detected transferring the top 100 song and movie titles by means of DATCoMs, out of all users detected transferring the titles by DATCoMs or by filenames for each monitoring period. The higher percentages achieved with songs shows that DPI with content matching is more effective with songs than with video.[3] This is probably the result of technical differences in detecting audio vs. video. For example, a single packet of audio can contain seconds of content, whereas with video it is generally only a small fraction of a second.

[3] A formal test of the hypothesis that the percentage of people detected transferring each copyrighted song by DATCoMs is greater than the percentage detected transferring each copyrighted movie by DATCoMs, against the null hypothesis that they are equal, yields statistically significant differences in mean percentages for songs against movies in all periods, ranging from a low of 26% (15% to 36%) in Fall of 2007 to a high of 48% (36% to 59%) in Spring of 2007.

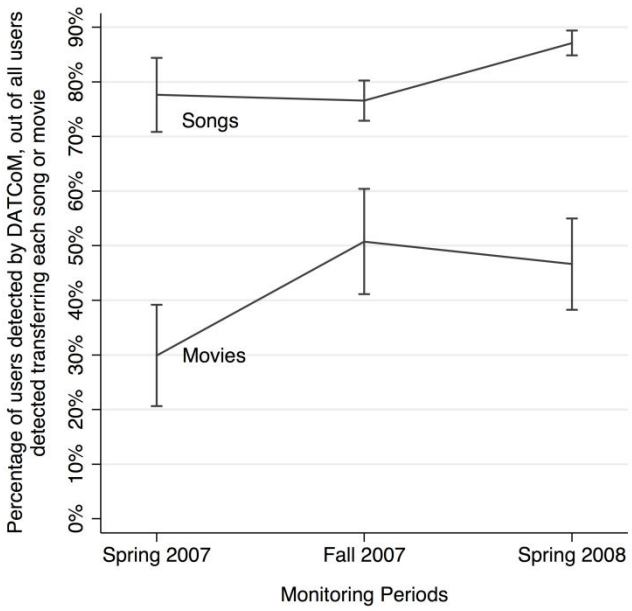

Fig. 3. Average percentage of users detected by DATCoM transferring each title in the top 100 out of all users detected transferring each title in the top 100 (by either DATCoM or filename).

5.2 Detecting users who transfer copyrighted content without encryption

If the goal is to identify users who transfer copyrighted material rather than the specific titles transferred, then it is not necessary to detect every transfer. When given enough time, which is apparently on the order of weeks, the DPI tool used at ISU was able to detect most individuals who attempted to transfer copyrighted content as long as they did not consistently conceal their transmissions (e.g. with encryption).

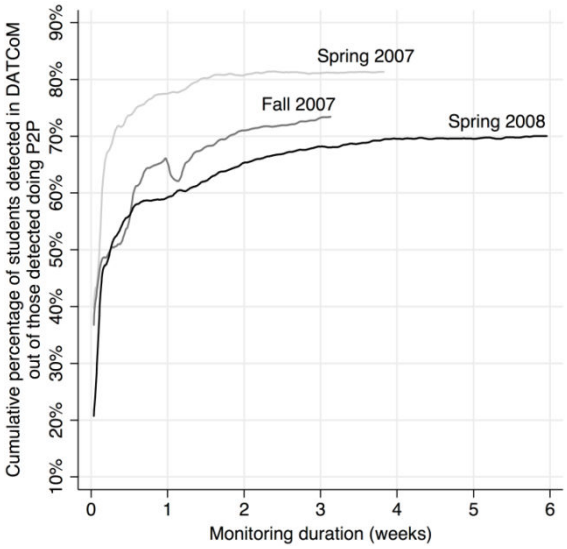

This is shown in Fig.

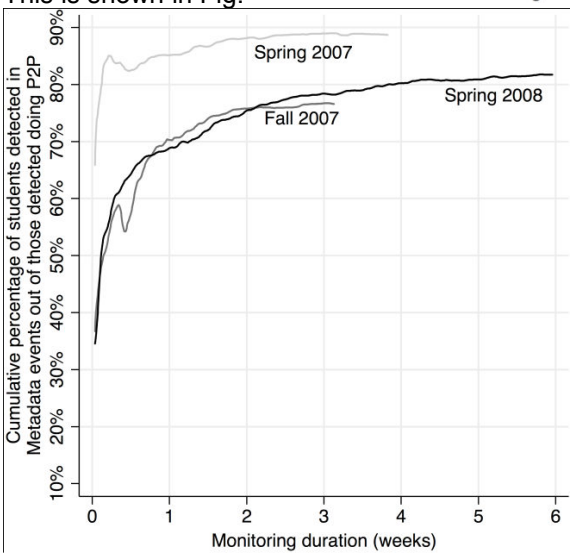

a. DATCoM users b. Users detected in Metadata events

Fig. 4.a, which shows the cumulative ratio of detected DATCoM users out of detected P2P users as a function of monitoring time. Initially, the percentage of detected P2P users observed in DATCoM is low, but after a couple weeks of monitoring that percentage tended to stabilize. Presumably, this percentage stops increasing because most DATCoM users have been observed at least once. Moreover, in each of the three monitoring periods, that percentage ultimately included a substantial majority of detected P2P users. To further support the hypothesis, the results are similar when using metadata rather than

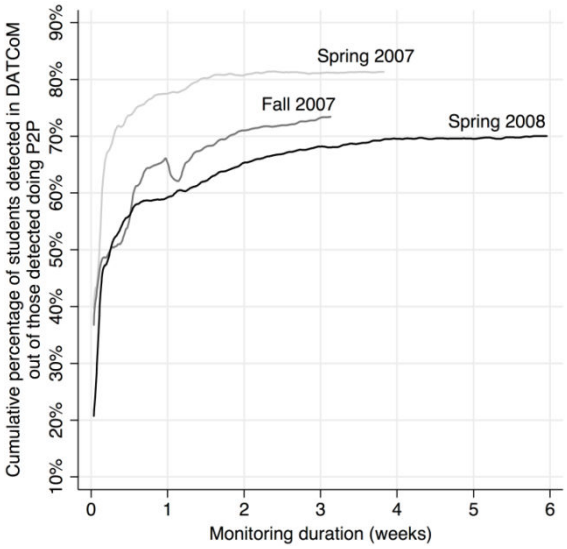

DATCoMs, as shown in Fig.

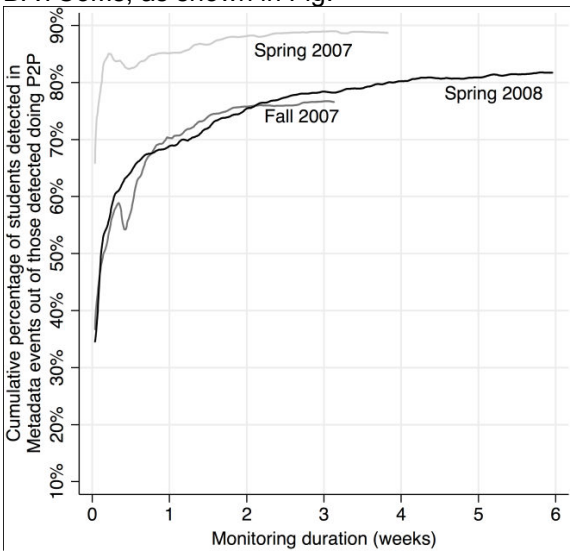

a. DATCoM users b. Users detected in Metadata events

Fig. 4.b which plots the same ratio using detection based on metadata.

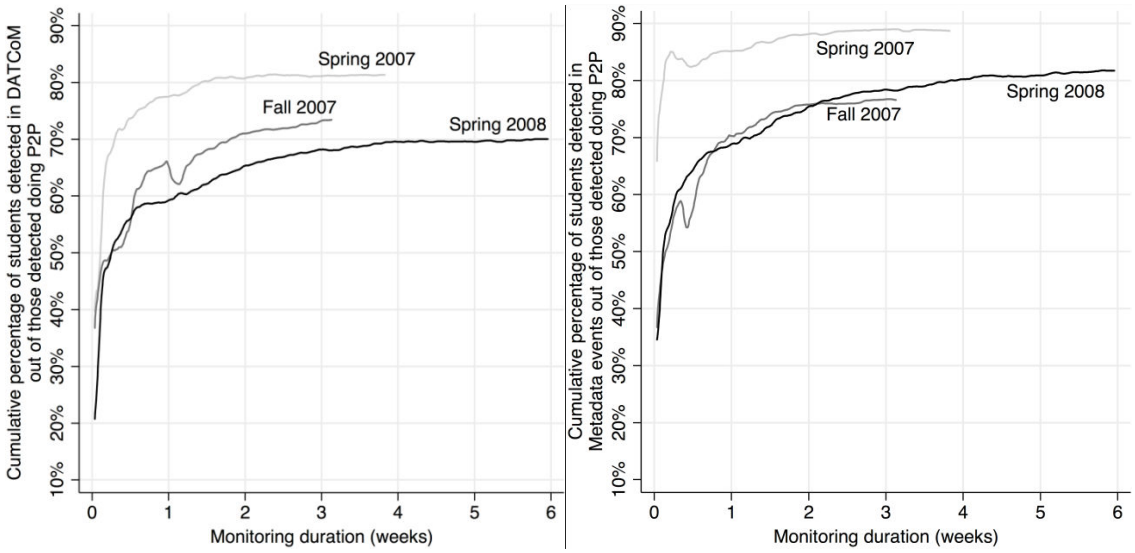

a. DATCoM users b. Users detected in Metadata events

Fig. 4. a) Cumulative ratio of detected DATCoM users out of detected P2P users as a function of the number of hours given to monitoring in each period. b) Cumulative ratio of users detected in events with Metadata out of detected P2P users as a function of the number of hours given to monitoring in each period.

The reason measurement can identify most DATCoM users in the first week is that those who are observed with one DATCoM tend to be observed with many. Indeed, perhaps even more copyrighted files are transferred without detection. For example, in the first 25-day monitoring period, DPI detected an average of 333 DATCoMs per user that was observed attempting to transfer copyrighted content at least once. This means that it is possible for a network-layer detection system to miss many more transfers while still observing most DATCoM users. Fig. 5 shows how the results of 25 days of monitoring would change if the DPI system had missed a given percentage of the DATCoMs, randomly selected with equal probabilities. If DPI kept just half of the DATCoMs from the first trial, it would observe 70% of the copyrighted titles, and nearly all of the DATCoM users. Indeed, if DPI caught only 15% of these DATCoMs, it would still observe 90% of DATCoM users.

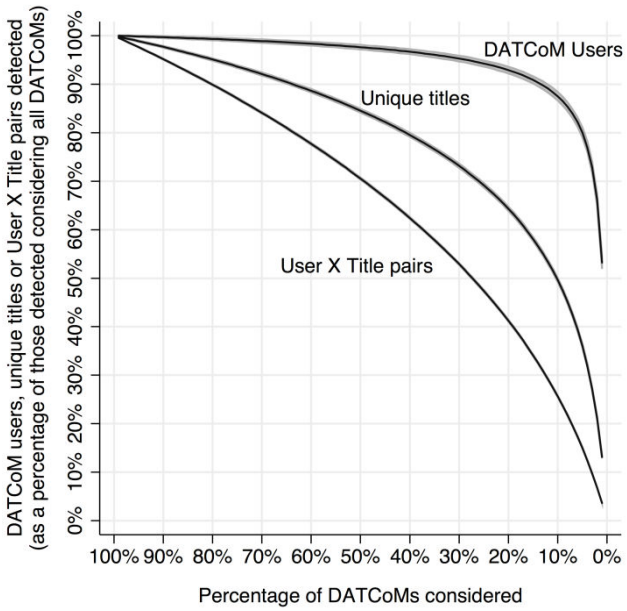

Fig. 5. Percentage of DATCoM users, unique copyrighted titles and user×title pairs that would be detected under smaller percentages of detected DATCoMs. Lines represent the mean percentages and shaded areas represent two standard deviations from the mean.

6 Growth of encrypted P2P

Even though most users of unencrypted P2P are probably observed after a sufficiently long monitoring time, users of encrypted P2P could evade detection indefinitely. This section will consider the extent to which this was observed at ISU, how it was changing over time, and how some of the other network-layer detection techniques discussed in Section 3 can be used to identify P2P, even if increased use of encryption undermines DPI's effectiveness. However, note that encryption is highly effective at concealing which content is transferred and whether copyright law might have been broken, so these other approaches can only be effective if the goal is to identify P2P irrespective of content (e.g. for resource allocation or to send out warnings).

Encrypted or not, there are ways to detect P2P (Fisher & Feyen, 2006; Collins & Reiter, 2006). Results are shown here for BitTorrent, which is the most popular P2P protocol. When BitTorrent is encrypted, the tools at ISU would label it as "unclassified." Some of the unclassified traffic does appear to be encrypted BitTorrent As discussed in Section 3, the simplest way to identify BitTorrent is port number. BitTorrent clients initially used port 6881 consistently. This is no longer necessary for most clients. However, at the time of this study, even a client that would not choose port 6881 on its own is likely to communicate with a peer that does still choose port 6881 sooner or later. (This approach will become less effective over time.) Thus, when "unclassified" sessions are observed that use port 6881, these are interpreted as encrypted

BitTorrent sessions.[4] After 25 days of monitoring in Fall 2007, for every host that was identified by DPI as a P2P user there were 0.45 hosts that were not identified by DPI as P2P users but were seen exchanging unclassified traffic using Port 6881. The majority of these are probably using BitTorrent with encryption, which indicates that a large number of college P2P users were safely immune from being associated with a DATCoM.

From a policy perspective, the more important question is probably whether use of encryption was increasing. Measurements around the world show that traffic identified as P2P is no longer increasing rapidly, and is falling as a percentage of total Internet traffic (Cisco, 2012). Nevertheless, the percentage of traffic identified as P2P fell even more rapidly at ISU, and it fell in absolute terms as well, as shown in Fig. 6 which characterizes traffic between the university and the rest of the Internet. Moreover, on the portion of the network serving university dormitories there was a decrease in detected P2P users, detected DATCoM users, and detected DATCoMs per DATCoM user, as shown in Fig. 7. Did this occur because students were using P2P less and exchanging fewer copyrighted files less, or because network-layer detection tools were less effective because users were adopting countermeasures? The answer is important, both to understand usage trends and to understand the impact of detection technology and reactions to it. ISU's Digital Citizen Project was publicly announced, and plans to monitor P2P use had been reported on university websites, campus newspapers, and elsewhere (Froeming, 2006; Smith, 2006; Swasko, St. John, Strand, & Yurgil, 2007). It is plausible that this discussion convinced some P2P users to quit, or that it convinced some users to adopt encryption. It appears that both were occurring.

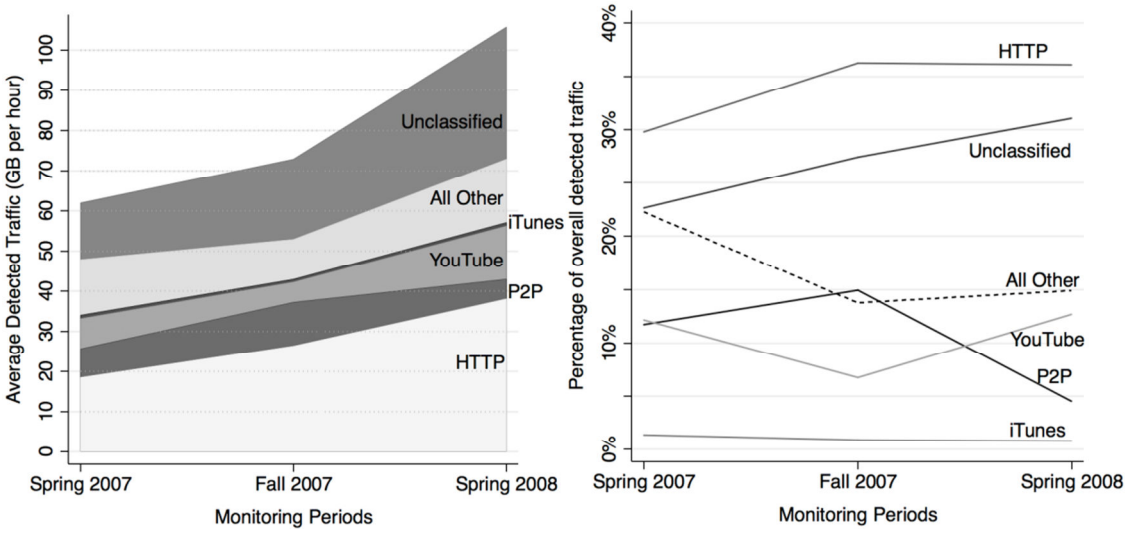

a. Hourly average exchanged traffic

b. Percentage of overall traffic

Fig. 6. Breakdown of detected traffic among main detected protocols. a) Hourly average exchanged traffic in each period. b) Overall percentage of detected traffic in each period.

[4] Although there is a chance that an application other than BitTorrent would use port 6881, it is unlikely to occur often given that no other application commonly uses this port and that Windows does not include 6881 in its list of volatile ports. More importantly for the validity of this particular comparison, it is equally likely in the two time periods compared.

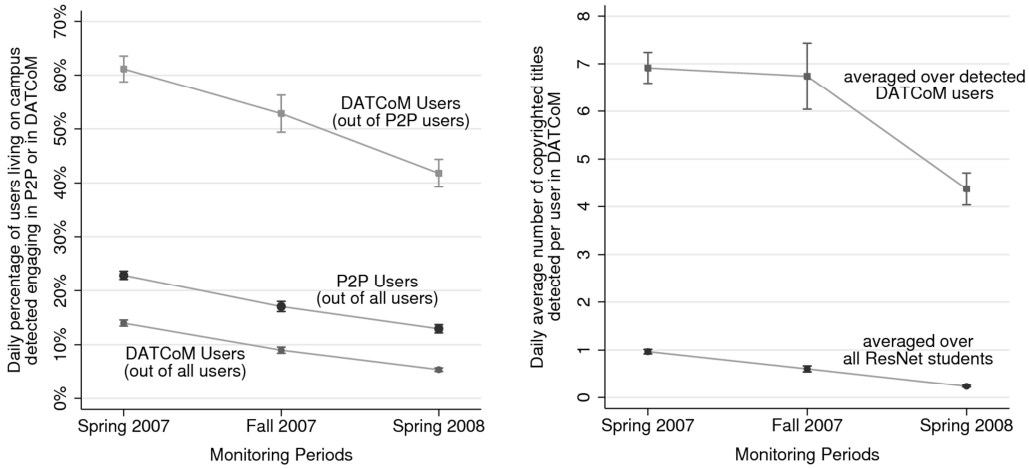

Fig. 7. (a) Average daily percentage of students detected engaging in P2P out of all students living on campus, and of students detected engaging in DATCoM, out of all students living on campus and out of detected P2P users in each day. (b) Daily number of copyrighted media titles detected in DATCoM, averaged over all students living on campus and over students detected engaging in DATCoM in each day. Caps show 95% confidence intervals.

One sign of the growing use of encryption is that the fastest-growing traffic category in Fig. 6 is unclassified. As further evidence, Fig. 8 shows the average fraction of all hosts on campus (or more specifically, all fixed IP addresses observed during that monitoring period) that were identified via DPI as engaging in P2P in a given hour, or observed using the port favored by BitTorrent, or both. Note that this includes all hosts on campus including servers, rather than just those computers sitting in student dorms, so the percentage engaging in P2P is considerably lower than in Fig. 7. Fig. 8 shows that the fraction of hosts identified as P2P users did fall during this period, but so did DPI's ability to detect them. Of the hosts seen using port 6881, DPI identified 91% of them as P2P users in Fall 2007, but only 75% of them in Spring 2008.

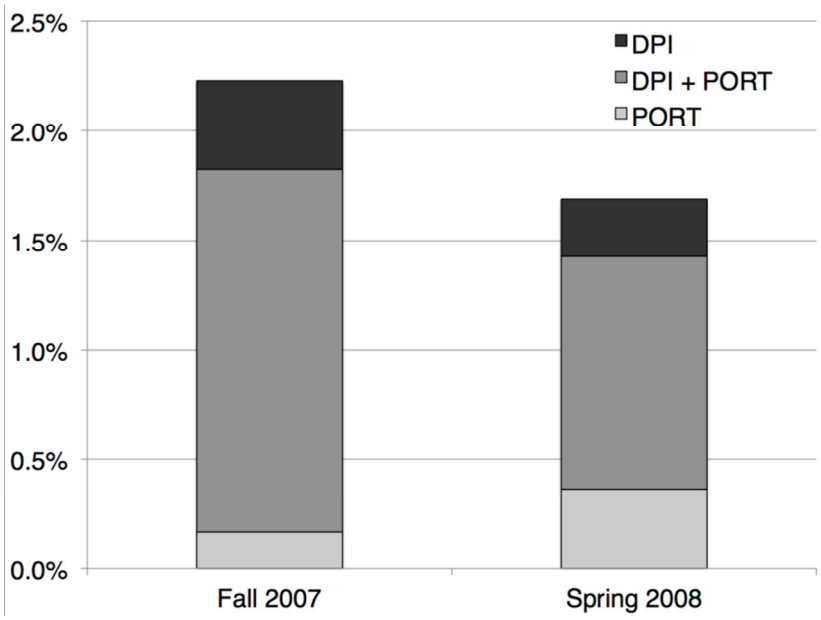

Fig. 8. Average percentage of all hosts observed engaging in P2P per hour via DPI, port number, or both.

7 Legal implications of technical differences

In addition to any differences in effectiveness, there may be legal differences between application-layer detection and network-layer detection approaches such as DPI and use of behavioral indicators. Sections 7.1 and 7.2 consider two such issues: liability of network operators and privacy.

7.1 Liability of network operators

Since network operators have limited ability to control or even observe what users do online, network operators are shielded to some extent from responsibility for user actions. This protection often depends on what network operators know about the traffic they carry, so detection technology matters. Indeed, it has been argued (Frieden, 2009) that under US law, any ISP that uses DPI may put these Safe Harbor protections from liability in jeopardy, potentially deterring useful applications of DPI.

In the U.S., two laws are of particular importance: the DMCA (US Congress, 1998) and the Communications Decency Act (CDA) (US Congress, 1996). Under the DMCA, ISPs are not liable for copyright infringement provided that they are not aware of the infringement or of "facts or circumstances from which infringing activity is apparent." If they are made aware of infringement, they have a number of responsibilities, including taking down the material when possible and terminating service for "repeat infringers." In the case of P2P file sharing, it is generally not possible to take down the allegedly infringing content, but ISPs can still identify the customer involved in the alleged infringement. An ISP can lose its safe harbor protection by failing to cooperate.

In many ways, the CDA offers much greater protection, and for a wide range of illegal activities. The CDA declares that ISPs cannot be treated as authors even when they have actual knowledge of offending material, provided that the ISP plays no role in selecting the content or determining who receives it. However, US courts have found that ISPs can be held liable if they exert any editorial control. For example, the ISP Prodigy was held liable for material that a customer posted on a public electronic bulletin board on the grounds that Prodigy had removed objectionable material in the past, effectively putting itself in the role of editor (Carter, 2008).

Safe harbor protections should not treat a network operator as if it is the author of all content that traverses its networks simply because the operator uses DPI or other network-layer detection. First, DPI is sometimes used for purposes that have nothing to do with copyright or other content issues, such as detection of malware (Peha, 2007). Copyright law should not deter this. Second, as shown by Section 5, even when technology detects some transfers of copyrighted material, it probably misses many more, and it is counterproductive to hold ISPs accountable for transfers they cannot detect. Third, even if an illegal transfer can be detected with absolute certainty, which is a difficult standard to achieve for copyright violations, this does not necessarily mean that network operators can stop the transfer. Blocking is sometimes possible, but it is more difficult than detecting, which may be why most vendors were reluctant to see their "filtering" technology quantitatively assessed in trials sponsored by copyright-holders, even when devices were allowed to filter all P2P transfers and not just those involving copyrighted material (Mueller, Kuehn, & Santoso, 2012; Rossenhovel, 2008). Comparing content captured from one or more P2P packets with content in a database of material to be blocked is a complex task. After it completes, the device could attempt to drop all subsequent packets in the stream or terminate the TCP session with a man-in-the-middle attack, but the transfer may have completed by this time. Even if that particular transfer can be blocked, the recipient may be able to find another source, and the subsequent transfer could escape detection. Regardless of what technology is deployed, it would be unreasonable to hold network operators accountable for ever allowing certain kinds of transfers to happen in the absence of highly effective technical solutions.

7.2 Privacy

This section addresses some implications for privacy policy that are rooted in detection technology as categorized in Section 3. (A broader legal analysis of which detection methodologies violate existing privacy laws around the world and which do not is beyond the scope of this paper.) Privacy laws that govern the personally identifiable information that public or private entities can maintain on individuals tend to be technology-agnostic, but sometimes the technical approach does matter.

DPI technology has been a particular focus for those concerned about privacy (Daly, 2011; Fuchs, 2012; Ohm, 2009; Katyal, 2004; Collins, 2009), more so than other detection approaches. Certainly a

superficial analogy to pre-Internet communications media could make one conclude that DPI is more intrusive than behavioral network-layer detection technologies. For example, in the U.S., law enforcement agents seeking judicial permission to listen to the content of a telephone conversation (i.e. a wiretap) face a much higher burden of proof than agents who only wish to find out what phone numbers a suspect has called (i.e. a pen register) (Samuel, 2008). Somewhat similarly, post offices in most countries are not allowed to look inside letters, but we all assume post offices can look at the addresses of sender and recipient. (This is part of the common carrier tradition, although this may be a distraction since ISPs are not viewed as common carriers in all countries.) By simple analogy, there might be less privacy protection for the packet header information that is monitored in some behavioral approaches as compared to the "content" accessed in DPI. One might argue that application-layer detection approaches are also like wiretapping, in that they also examine content, although where the agent doing the wiretapping is also on the call.

As an aside, note that while legal discussions of DPI sometimes assume that transmitted bits are either content or header, and examination of headers does not constitute DPI, engineers know that the term header is ambiguous and demands further clarification. Packets are encapsulated, one inside another, much like Russian dolls. As a result, bits in the header of a TCP packet are considered content of an IP packet, and bits in the header of an IP packet are considered content of an ethernet packet. An important example is port number, as discussed in Sections 3 and 5. Port is in the header of a TCP packet and the content portion of an IP packet. Since TCP headers are intended for the final recipient on the network, it is more appropriate to view examination of port by a network device as DPI, but firewalls have been filtering network traffic based on port for years.

While the simple extension of telephone privacy policy to the Internet has intuitive appeal, courts and policymakers could instead turn to the original logic behind telephone policy, which was based on the idea that people have a *reasonable expectation* that the content of their calls will be private, but their expectations for who they call should be lower. As Justice Blackman pointed out, a list of long-distance numbers called appear on monthly phone bills, so consumers know these numbers are logged (Samuel, 2008). This article will not speculate as to what level of privacy Internet users expect when it comes to recording the time between successive packets or the percentage of failed TCP sessions, but finding out may be complicated by the fact that very few users know what packets and TCP sessions are. Thus, it is less clear how the level of legal privacy protection may depend on detection technology if judgements are based on "reasonable expectation."

An alternative framing is to place a higher degree of protection against those detection techniques that are available to network operators because consumers entrust them with communications. With this perspective, the most important distinction would be application-layer approaches versus network-layer approaches as defined in Section 3. Use of DPI for law enforcement or for commercial purposes such as

behavioral advertising would then be more restricted than swarm infiltration, on the grounds that anyone can engage in swarm infiltration, but there would be no legal distinction between DPI and network-layer behavioral techniques which do not examine "content."

8 Conclusions and policy implications

This paper has shown that a wide range of technologies can be employed to detect P2P usage, identify the content being shared, and determine who is involved. This detection is at the root of policies and practices adopted in response to widescale copyright violations online, and of network neutrality policies which govern resource allocation techniques that discriminate based on application and content.

To a large extent, the choice of technology determines who is responsible for detection, which should be reflected in policy, regardless of whether that policy is derived from law, regulation, or agreements among stakeholders. Application-layer approaches such as swarm infiltration and tracker monitoring require no involvement from a network operator to determine which content is most popular and the IP addresses of users who are sharing that content. The only role of a network operator is to map an IP address to the individual responsible for that IP address (and perhaps to interact with that individual). In contrast, network-layer approaches such as DPI and various behavioral techniques require access to the network, and therefore much greater involvement by the network operator. Unless they are also content providers (and some are), network operators may choose not to pay for detection technology for copyright protection unless public policy demands it, but there is only reason to consider such policies if network-layer detection is found to be preferable to application-layer detection. Moreover, P2P users (or more likely those who design software for P2P users) can choose among many technical countermeasures. Thus, rather than simply choosing the detection technology that yields the "best" results in isolation, one must consider how a choice will change the outcome of this complicated multi-player game.

Which detection approach is best depends greatly on the purpose. Moreover, the existence of countermeasures creates some ironic conflicts. For example, some proposals designed to protect content creators involve compensating copyright-holders in proportion to the number of times their content is shared using P2P. Other proposals designed to protect the same group would attempt to detect illegal transfers so as to punish those responsible and deter further copyright violations. However, these goals are in conflict; punishment creates incentives for users to adopt countermeasures such as turning on encryption and preventing untrusted agents from joining swarms that make it harder to determine which content is being shared, complicating attempts at compensation. Similarly, a university may wish to determine which students are using P2P and then make sure these students are knowledgeable about copyright law, or to identify P2P traffic (regardless of who sent it) so they can allocate network resources accordingly, but these tasks become harder if students are motivated to conceal their traffic (Lehr, Sirbu, Gillett, & Peha, 2007). A comprehensive strategy is needed that takes potential countermeasures into account.

Since the best approach to detection depends on the purpose, each purpose will be considered here individually, beginning with copyright enforcement that includes some form of punishment. This requires reliable identification of the individual involved, and a determination that the content is copyrighted. (Even this is necessary but not sufficient for punishment, since correctly determining that a given person has attempted to transfer a copyrighted file does not necessarily mean that this person has violated copyright law.) Achieving these detection goals does not, however, require that every illegal transfer be detected. Based on the experience at ISU, network operators could use network-layer detection very effectively if and only if no countermeasures are adopted. Those who shared copyrighted files were observed doing so often, and most of these individuals can be identified with current network-based technology in just a few weeks, even though some transfers are clearly missed. (Especially video.) However, if P2P users turn on encryption, it would be impossible to tell whether the files being shared are copyrighted, and threats of punishment may motivate many to do so. Indeed, use of encryption during ISU's Digital Citizen Project was found to be significant and increasing over time. (More research is required to determine whether this trend is common throughout the Internet, or a response to Digital Citizen's publicly announced plans to monitor.) Alternatively, agents of copyright-holders have long used application-layer detection. This approach has limitations, but it will remain effective even if more users turn encryption on, or P2P clients begin encrypting by default. Moreover, because application-layer detection has more context, it is less likely to identify users who are transferring copyrighted files within the bounds of fair use. Thus, even for ISPs that adopt "graduated response" and therefore have a role in punishing violators of copyright law, it is probably better to use application-layer detection agents rather than require ISPs alone to determine which users are guilty. In this sense, some calls to hold ISPs accountable for copyright violations are overly simplistic, and possibly self-defeating. Given these advantages, there are reasons to applaud the shift from network-layer to application detection methodologies for graduated response policies that has occurred in recent years in Belgium, France, the U.K., and the U.S. (Mueller, Kuehn and Santoso, 2012).

With any of these detection approaches, when punishment is involved, it is necessary to give those who are accused an opportunity to defend themselves, for several reasons. First, technology can detect the sharing of copyrighted content, but technology cannot determine whether the individual has a right to make that transfer; not all DATCoMs are copyright violations. Thus, those who are accused need an opportunity to show that they were acting within their rights. Second, there are instances in which DATCoM detection systems reach an incorrect result. Anecdotes of obviously incorrect DMCA take-down notices have been reported, including notices asserting that a Mac computer was running file-sharing software that is only available on Windows (Katyal, 2004), notices that a printer is downloading content (Piatek, Kohno, & Krishnamurthy, 2008), and notices that an RTF (text) file entitled "Harry Potter Book Report" violated the copyright associated with a Harry Potter movie (Katyal, 2004). This issue has been at

the heart of the debate in the European Union over graduated response. As Horten (2011) put it, "the story began with copyright and ended a discussion on the right to a fair trial."

While giving users a chance to object to allegations is necessary, it is not sufficient. There has been too little attention on the underlying detection mechanisms, and the technical details matter. False allegations in the past are not simply clerical error; accusations have been made based on methodologies that are known to be fatally flawed (Piatek, Kohno, & Krishnamurthy, 2008; Freedman, 2009), Systemic problems like this seriously undermine fairness and credibility of copyright enforcement. However, this problem can be greatly alleviated by moving to a policy of transparency. Whenever there is a possibility of punishment, allegations should be supported with evidence, and evidence is worthless without a detailed description of how that evidence was gathered. In the U.S., copyright-holders have typically sent notices accusing a user of copyright infringement with no indication of how the alleged infringement was detected (Virginia, 2012). Perhaps designers of the detection mechanisms cut corners because they wanted to increase speed and lower costs, or simply because they were not aware of the problems. They may have used a source IP address that they observed without verifying it, or relied on metadata that is inherently unreliable. Even where the technology has correctly identified the transfer of copyrighted content, an assessment of fair use claims requires knowledge of whether this was from network-layer or application-layer detection, and this was not stated. To solve this problem, any accusation of infringement should be deemed invalid if it does not both provide some of the evidence that allegedly indicates guilt and specify the mechanism used to gather that evidence. Moreover, this mechanism must have been publicly disclosed prior to the allegation, so outside observers have the opportunity to check for flaws.

There is now a risk that, under the memorandum of understanding signed by copyright-holders and ISPs that will bring graduated response to the U.S. (Center for Copyright Information, 2011), this serious lack of transparency will be perpetuated. Although the MOU does guarantee users accused of repeated copyright violations the ability to argue for their innocence, the MOU includes no provisions to give those accused the forensic evidence or a description of the detection methodology.[5] Nevertheless, the MOU need not be the last word; it is always possible to put such provisions in place.

A related purpose of identifying P2P users is simply to inform them about copyright policy and the risks associated with violation. As long as punishment is not involved even after multiple warnings, it is tolerable to send warnings to individuals who have not actually violated copyright law, which simplifies the task. Users are less likely to adopt countermeasures in response. Indeed, warnings might be sent to P2P users even when it is not determined whether the content is copyrighted, so encryption is less problematic. Thus, the limitations of network-layer detection are less significant, and the advantages are

[5] The MOU does include provisions for a review by technical experts, but these experts are selected by industry and not the accused. Moreover, even if an expert finds that the detection methods are completely unreliable, the expert is prohibited from revealing this finding to the accused or the public, and the ISPs and copyright-holders are entitled to ignore the finding.

helpful. Network operators are also well positioned for this task because they already have the ability to identify the IP addresses of P2P users, and the ability to map IP address to individual and contact that individual. For mere warnings, there may not be sufficient reason to believe that a law has been violated to give copyright-holders the private or proprietary customer information needed to contact the user. Some network operators such as universities may see this as a useful service, since warnings may prevent copyright-holders and law enforcement from taking more serious legal actions against their users, although their level of interest will depend on equipment cost. Only experimentation can determine the extent to which warnings without punishment can influence user behavior.

A third reason to detect transfers of copyrighted files is to attempt to block them in real time, which can only be done with network-layer detection. Compared to the uses above, real-time blocking is more challenging from a technical perspective and problematic from a policy perspective. Moreover, because not all DATCoMs are copyright violations, users should be given the opportunity to claim that their transfer is legal, and to continue file sharing at least until claims can be assessed, which is a problem with real-time blocking. Otherwise, legitimate traffic will be blocked. The Court of Justice of the European Union recently reached a similar conclusion (European Union, 2011). In 2004, the Brussels Court of First Instance had ordered a Belgian ISP to deploy technology at its own expense that would block acts of copyright infringement (Mueller, Kuehn, & Santoso, 2012). In November of 2011, the EU Court of Justice overturned this order for a variety of reasons ranging from the cost incurred by ISPs to the privacy risks to consumers, but one factor is especially relevant here: this requirement "could potentially undermine freedom of information since that system might not distinguish adequately between unlawful content and lawful content, with the result that its introduction could lead to the blocking of lawful communications" (European Union, 2011)

Another purpose of detection is to determine the extent to which each copyrighted title is transferred so as to compensate copyright-holders accordingly. This turns out to be a surprisingly hard problem for several reasons, and further research is needed to determine its practical viability before these policies can be seriously considered. First, even roughly accurate measurements are challenging. There have been studies of the extent to which different titles are available in swarms indexed by public trackers (Felton, 2010; Layton & Watters, 2010; Envisional, 2011; Mateus & Peha, 2012b), and first-order estimates of the extent to which these titles are actually transferred (Mateus & Peha, 2012b). However, not all files are indexed on public trackers, and while first-order estimates are certainly useful to compare entire content categories and monitor trends, further research is needed to determine whether an even higher level of accuracy would be needed to compensate copyright-holders for individual titles. More importantly, these methodologies quantify P2P globally, when usage by country is probably more useful for this purpose. It is unlikely that any country will establish a policy whereby funds are collected within the country, but distributed based on P2P transfers among people who are outside the country. It is possible to infer a user's country from IP address, but gathering IP addresses and verifying their validity

vastly increases complexity and cost. In contrast, network-layer detection performed by ISPs could easily be directed to P2P within the country of interest, thereby solving this particular problem, but this brings other challenges. At ISU, the percentage of song transfers detected by current DPI technology was probably several times the percentage of video transfers detected, and there may be other similar distortions. To make compensation fair, methods must be found to identify and explicitly account for any differences in probability of detection.

An even bigger problem is that some copyright-holders may deliberately download their own content to increase their compensation. This can be done on a large scale by enlisting a botnet, i.e. an army of computers whose security has been compromised, and is now under the partial control of a botmaster. Botmasters routinely rent out these compromised computers to spammers for the generation of email, so why not rent them out to copyright-holders for the exchange of P2P traffic? For example, a compromised computer can cost just 15 cents to rent (Goodman, 2010), and might more than pay for itself by downloading a single copyrighted song. Nevertheless, it can download dozens of songs per hour while consuming such a small fraction of that computer's broadband connection that the owner would never notice. Until a solution can be found to this problem, or the economics of botnets change drastically, no compensation approach based on monitoring downloads can be viable.

An alternative is to compensate copyright-holders based on data provided by content providers who do not use P2P for distribution rather than based on what can be determined by monitoring P2P in some way. Many of those who sell content legally on physical media (e.g. CDs and DVDs) or from legal download sites would probably cooperate. There are challenges and limitations here as well. First, such an approach means that no matter how popular a song becomes via P2P, its creator can never make a profit without going through traditional channels, such as producing and marketing CDs or gaining access to iTunes. Thus, there would still be financial disincentives to take full advantage of the efficiencies of P2P distribution, especially for more obscure content. Also, even when content is easily available through both P2P and traditional outlets, its popularity is not always similar in both markets. For example, it has been observed that the number of times a song is transferred illegally via P2P for every time it is sold legally is much higher for content that appeals more to teenagers (e.g. Justin Bieber) than for content that appeals more to an older audience (e.g. Susan Boyle) (Mateus & Peha, 2012b). Perhaps compensation calculations should be based on a mixture of both online and offline data. A compensation system does not have to be perfect for it to be effective, but some of the challenges described above must be addressed.

A fourth use of P2P detection technology is for resource allocation, or in this case, reducing the QoS of P2P as a means of protecting the QoS of other applications that are considered to be more important, or more sensitive to poor performance, or both. Results indicate that current technology can perform this function, at least in an enterprise network. Moreover, there are benefits. However, in the hands of

commercial ISPs with significant market power, there can be strong incentive to use these capabilities in ways that harm consumers by extracting oligopoly rents, unless there are "Open Internet" policies in place that limit such behavior (Peha, 2007). Moreover, some attempts to extract oligopoly rents may involve limiting customers' access to music or video content by blocking transfers or degrading QOS, which may superficially look like legitimate copyright protection. Scrutiny is sometimes needed.

This similarity demonstrates a potential conflict between these Open Internet and copyright protection policies that is partially rooted in the limitations of detection technology. The FCC was careful to write Open Internet rules that only protect "lawful network traffic" (FCC, 2010). Thus, ISPs would clearly be within their rights if they blocked transfers known to violate copyright law, and perhaps to selectively degrade the QoS of some applications for users known to repeatedly violate copyright law, as many U.S. ISPs plan to do in their graduated response (Center for Copyright Information, 2011). But what if graduated response punishes users who were falsely accused? In reality, network operators cannot always know with certainty whether copyright law has been violated, and this is not reflected in policy as written. Moreover, ISPs owned by film studios or music labels have less incentive to give consumers the benefit of the doubt. The regulator must decide when ISPs are acting "reasonably." If done well, cooperation between ISPs and copyright-holders could yield a system for reviewing accusations of copyright violation that is vastly more efficient and more just than the costly and contentious civil cases of the past, but if done poorly, such an arrangement could eliminate the traditional protections of a legal process. At minimum, any ISP that degrades service based only on an unfounded accusation should be found in violation of Open Internet rules, and any accusation that does not reveal its detection methodology and related details (as unfortunately appears to be the current plan (Center for Copyright Information, 2011)) should be viewed by the regulator as unfounded.

9 References

Baumgartner, J. (2009, Oct. 15). Cox shuts down net congestion tests, *Light Reading.* Retrieved from http://www.lightreading.com/policy-control/cox-shuts-down-net-congestion-tests/240107546

Bridy, A.. (2010). Graduated response and the turn to private ordering in online copyright enforcement. *Oregon Law Review 89*, 89-132.

Carter, T. B. (2008). Who is safe in this harbor? Rethinking section 230 of the Communications Decency Act, *Forum on Public Policy*. Retrieved from http://www.forumonpublicpolicy.com/archivespring08/carter.pdf.

Center for Copyright Information (2011, July 6), Memorandum of understanding. Retrieved from www.copyrightinformation.org/sites/default/files/Momorandum%20of%20Understanding.pdf

Cheng, J. (2009, March 26). The pirate bay to roll out secure €5 per month VPN service, *Ars Technica*. Retrieved from http://arstechnica.com/telecom/news/2009/03/the-pirate-bay-to-roll-out-secure-vpn-service.ars

Cisco (2012). *Cisco visual networking index: forecast and methodology, 2011-2016*. Retrieved from http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-481360.pdf

Collins, R. T. G. (2009). The privacy implications of deep packet inspection technology: Why the next wave in online advertising shouldn't rock the self-regulation boat, *Georgia Law Review*, *44*, 545-79.
Collins, M. P. & Reiter, M. K. (2006). Finding peer-to-peer file-sharing using coarse network behaviors, *Proceedings of ESORICS*, 1-17.
Daly, A. (2011). The legality of deep packet inspection, *International Journal of Communications Law and Policy, 14*, 1-12.
Dingledine, R. Mathewson, N. & Syverson, P. (2004). Tor: The second-generation onion router, *Proceedings of 13th Usenix Security Symposium.*
Electronic Frontier Foundation (EFF) (2008, April 1). A better way forward: Voluntary collective licensing of music file sharing. Retrieved from https://www.eff.org/wp/better-way-forward-voluntary-collective-licensing-music-file-sharing
Envisional (2011). *Technical report: An estimate of infringing use of the internet*. Envisional ltd. Retrieved from http://documents.envisional.com/docs/Envisional-Internet_Usage-Jan2011.pdf
Ernesto (2008, August 12). Download torrents anonymously with torrentprivacy, *TorrentFreak.com*, 2008. Retrieved from http://torrentfreak.com/download-torrents-anonymously-with-torrentprivacy-080812
Ernesto (2009, Nov. 24). More BitTorrent users go anonymous, *TorrentFreak.com*, June 22, 2009. Retrieved from http://torrentfreak.com/more-bittorrent-users-go-anonymous-090622
European Union Court of Justice (2011). Judgment of the court (Third chamber), Scarlet Extended SA v Société belge des auteurs, compositeurs et éditeurs SCRL (SABAM). Retrieved from http://curia.europa.eu/juris/document/document.jsf?docid=115202&doclang=EN&mode=&part=1
Federal Communications Commission (2008, Aug. 20). in the matter of formal complaint of Public Knowledge and Free Press against Comcast Corporation for secretly degrading peer-to-peer applications, Memorandum Opinion And Order, File No. EB-08-IH-1518 and WC Docket No. 07-52. Retrieved from http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-08-183A1.pdf
Federal Communications Commission (2010, Dec. 23). in the matter of preserving the open Internet and broadband internet practices, Report And Order, GN Docket No. 09-191 and WC Docket No. 07-52. Retrieved from http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-10-201A1.pdf
Felton, E. (2010, Jan. 29). *Census of files available via BitTorrent,* blogpost on *Freedom to TInker*. Retrieved from http://freedom-to-tinker.com/blog/felten/census-files-available-bittorrent
Finnie, G. (2009). ISP traffic management technologies: State of the art, Canadian Radio Television and Telecommunications Commission, Retrieved from http://www.crtc.gc.ca/PartVII/eng/2008/8646/isp-fsi.htm
Fisher, W. F. (2004). *Promises to keep: Technology, law, and the future of entertainment*. Stanford, California: Stanford University Press.
Fisher, A. & Feyen, M. (2006, Aug.). Allot Communications NetEnforcer is first to detect and manage encrypted BitTorrent traffic, Allot Communications, Press release.
Folsom v. Marsh, Massachusetts (1841), In *Primary sources on copyright (1450-1900)*, L. Bently & M. Kretschmer (Eds.). Retrieved from http://copy.law.cam.ac.uk/cam/tools/request/showRecord.php?id=record_us_1841
Freedman, M. (2009, Nov. 23). Inaccurate copyright enforcement: Questionable 'best' practices and BitTorrent specification flaws, *Freedom to Tinker.* Retrieved from https://freedom-to-tinker.com/blog/mfreed/inaccurate-copyright-enforcement-questionable-best-practices-and-bittorrent-specificatio
Frieden, R. (2008). Internet packet sniffing and its impact on the network neutrality debate and the balance of power between intellectual property creators and consumers, *Fordham Intellectual Property, Media & Entertainment Law Journal*, *18*(3), 633-675.

Froemling, T. (2006, May 2). ISU develops new file sharing options. *Daily Vidette at Illinois State University*. Retrieved from http://is.gd/eFciC

Fuchs, C. (2012). Implications of deep packet inspection (DPI) Internet surveillance for society, *Privacy & Security Research Paper Series*, *1*. Retrieved from http://www.projectpact.eu/documents-1/%231_Privacy_and_Security_Research_Paper_Series.pdf

Goodman, D. (2010, April 28). Texas man cops to botnet-for-hire charges, *The Register*. Retreived from www.theregister.co.uk/2010/04/28/botnet_for_hire_guilty

Horten , M. (2011). The copyright enforcement enigma: Internet politics and the 'telecoms package,' London: Palgrave Macmillan.

Illinois State University (2006). *Digital Citizen Project, Summary of project*, Retrived from http://www.digitalcitizen.ilstu.edu/summary

Katyal, S. K. (2004). Privacy vs. piracy, *Yale Journal of Law and Technology*, *7*, 222-345.

Layton, R. & Watters, P. (2010, April). *Investigation into the extent of infringing content using BitTorrent networks*. ICSL - Internet Commerce Security Laboratory. Retrieved from http://www.afact.org.au/research/bt_report_final.pdf

Lehr, W. H., Sirbu, M. A., Gillett, S.& Peha, J. M. (2007). Scenarios for the network neutrality arms race, *International Journal of Communication*, 607-643. Retrieved from http://ijoc.org/ojs/index.php/ijoc/article/view/164

Lincoff, B. M. (1995). Everyone covered by blanket licenses. *Billboard 107*(50), 3-4.

Mateus, A. M. & Peha, J. M. (2012a, Winter). P2P on Campus: Who, What, and How Much," *I/S: A Journal of Law and Policy for the Information Society, 7*(2). Retrieved from http://moritzlaw.osu.edu/students/groups/is/archives/volume-72

Mateus, A. M. & Peha, J. M. (2012b). Characterizing global transfers of copyrighted content using BitTorrent, in press.

Mueller, M. L. & Asghari, H. (2012) Deep packet inspection and bandwidth management: Battles over BitTorrent in Canada and the United States. *Telecommunications Policy 36*(6). 462-475

Mueller, M. L, Kuehn, A. & Santoso, S. M. (2012) Policing the network: Using DPI for copyright enforcement. *Surveillance and Society, 9*(4), 348-364.

Ohm, P. (2009) The rise and fall of invasive ISP surveillance, *University of Illinois Law Review,* 2009I(5), 1417-1496.

Page, W. & Touve, D. (2010, Dec. 7). Moving digital Britain forward, without leaving creative Britain behind. *Economic Insight, 19*.

Parsons, C. (2008). Deep packet inspection in perspective: Tracing its lineage and surveillance potentials, Working Paper. Retrieved from http://www.sscqueens.org/sites/default/files/WP_Deep_Packet_Inspection_Parsons_Jan_2008.pdf

Peha, J. M. (2007). The benefits and risks of mandating network neutrality, and the quest for a balanced policy, *International Journal of Communication*, 644-668. Retrieved from http://ijoc.org/ojs/index.php/ijoc/article/view/154

Peha, J. M. (2008, April 4). Misstatements on Comcast P2P practices, and the implications for network neutrality, Comments in the matter of broadband industry practices, Federal Communications Commission WC Docket No. 07-52. Retrieved from http://apps.fcc.gov/ecfs/document/view?id=6519870758

Piatek, M., Kohno, T. & Krishnamurthy, A. (2008). Challenges and directions for monitoring P2P file sharing networks or why my printer received a DMCA takedown notice, *Proceedings of the 3rd conference on Hot topics in security*.

Posner, R. A.. (1992). When is parody fair use?, *Journal of Legal Studies*, *21*(1), 67-72.

Recording Industry Association of America (RIAA) (2011). *Piracy online*. Retrieved from http://www.riaa.com/physicalpiracy.php?content_selector=piracy_details_online

Rossenhovel, C. (2008, March 27). Peer-to-peer internet filters: Ready for Internet prime time? *Internet Evolution.* Retrieved from http://www.internetevolution.com/document.asp?doc_id=148803
Samuel, I. J. (2008). Warrantless location tracking, *New York University Law Review*, *83*, 1324-52.
Samuelson, P. (1994). Copyright's fair use doctrine and digital data, *Communications of the ACM*, *37*(1), 21-27.
Smith, J. (2006, Oct. 3). ISU staff speaks in Washington on issues of campus piracy. *Daily Vidette at Illinois State University*. Retrieved from http://is.gd/eFbpI
Storch J. and Wachs, H, (2011). A legal matter: Peer-to-peer file sharing, the Digital Millennium Copyright Act, and the Higher Education Opportunity Act: How Congress and the entertainment industry missed an opportunity to stem copyright infringement, *Albany Law Review, 74*, 313-360.
Swasko, M., St. John, D., Strand, E. & Yurgil, M. (2007, Feb. 21). Being a pirate is against the law, *Daily Vidette at Illinois State University.* Retrieved from http://is.gd/eFbUA
U.S. Congress (1996), Communications Decency Act, Title V, Section 230, Telecommunications Act of 1996. Retrieved from http://www.law.cornell.edu/uscode/text/47/230
U.S. Congress (1998, Oct. 28). Digital Millennium Copyright Act. Retrieved from http://www.copyright.gov/legislation/hr2281.pdf
U.S. Congress (2008, Aug. 14), Higher Education Opportunity Act. Retrieved from http://www.gpo.gov/fdsys/pkg/PLAW-110publ315/pdf/PLAW-110publ315.pdf
U.S. Court of Appeals (2010, April 6). On petition for review of an order of the Federal Communications Commission, Comcast Corporation v Federal Communications Commission. Retrieved from http://www.cadc.uscourts.gov/internet/opinions.nsf/EA10373FA9C20DEA85257807005BD63F/$file/08-1291-1238302.pdf
Virginia, University of, (2012) DMCA copyright information for UVa students and employees, Sample DMCA copyright violation notice. Retrieved from http://www.virginia.edu/informationpolicy/copyright/sample.html

A **Appendix: Data collection methodology**

This article presents results from research conducted using data collected in the scope of the Digital Citizen Project (DCP) at Illinois State University (ISU). The aim of the DCP was "to significantly impact illegal piracy of electronically received materials, using a comprehensive approach to confront pervasive attitudes and behaviors in peer-to-peer downloading of movies, music, and media" (Illinois State University, 2006). In February 2007, a team of engineers and social scientists from Carnegie Mellon University (CMU) began conducting research on the dissemination of copyrighted material on the ISU campus, which comprised the collection of network monitoring data. This section describes the methodology utilized for collection and anonymization of the network data used in this article.

A.1 Network monitoring

Monitoring data was collected from the ISU network, which serves the entire campus population. The campus network connects to the Internet by means of two Internet Service Providers (ISPs) and Internet 2. In order to maintain appropriate levels of service for applications considered crucial to the university's everyday operation, ISU uses traffic shaping at the point where the campus network connects to ISPs.

Shaping works by allowing specific classes of application to use up to a predefined percentage of the available bandwidth, but no limit is imposed on the amount of traffic generated by each network user.

The campus network is divided into several sub-networks. ResNet is the sub-network that students connect to in their dormitories. ResNet users purchase network access from ISU, which allows them to connect to one wired connection each in their dorm room. Each such connection has is assigned an IP address which is fixed for the entire semester. Students are not allowed to setup wireless 802.x networks in the residence halls.

Network monitoring was performed by three types of commercially available monitoring appliances, two were different types of deep packet inspection (DPI) appliances, and the third type consisted of regular Netflow collectors. The DPI appliances were Packeteer PacketShaper (Packeteer) and Audible Magic CopySense (AM). Both devices log relevant attributes of transmissions between users inside the campus network and outside parties, for traffic routed using the ISPs. Packeteer had already been deployed before the DCP project to perform traffic shaping as described above. It classifies communication sessions according to the type of traffic that they carry in over 500 classes, without retaining any actual content of the communications sessions.

AM was deployed on the network to enforce ISU policy before CMU started collaborating with the DCP. It uses header information to identify P2P communication sessions, and then tries to identify copyrighted media transferred in those sessions in real time. Identification of copyrighted media is performed by trying to match the content transferred in detected P2P sessions against a database of audio fingerprints of copyrighted media titles, or against hash codes of files known to contain copyrighted content. The hash code method works provided that this particular file has been seen before, and a hash of the file's content is already in the database. The audio fingerprinting approach is more challenging technically, but works even on a newly-produced copy of the content. With fingerprinting, a device collects a sample of the audio track, extracts relevant and unique characteristics, and compares these characteristics with the those in the database. The device records which copyrighted titles were matched in the database without permanently retaining any portion of the transmission. In cases when the content of P2P sessions cannot be matched against anything in the database, and when the sessions contain metadata describing the content being transferred (typically the name of the file being transferred), AM records such metadata.

AM logs information on communications in the form of *events*. An event corresponds to one or more consecutive TCP or UDP[6] sessions between a pair of peers in a P2P network. All the TCP or UDP sessions in an event are either identified as being associated with the same copyrighted media title, or cannot be associated with any media title in AM's database. Hence, an AM event means that two peers in

[6] UDP sessions are actually pseudo-sessions, with consecutive UDP packets being aggregated in the same pseudo-session if they occur within a time interval that is lower than a predefined threshold.

a P2P network, one inside the ISU campus and another one outside, were detected by AM exchanging or attempting to exchange information (either belonging to an identified copyrighted media title, or information that could not be identified as belonging to any copyrighted media title present in AM's database) over a set of consecutive TCP sessions or consecutive UDP sessions.

A.2 Connection of monitored activity to users and devices

All monitoring appliances produced logs of network activities that contain an IP address in use by a device on campus. In the case of data collected by AM, identification of devices and users associated with such IP addresses was implemented using data from several network management databases also collected from the ISU network. For each data record, the identification of the device (i.e., the device's MAC address[7]) from the recorded IP address was performed via lookups in the DHCP[8] lease logs using the IP address of the monitored activity and the time when the activity occurred.

To identify the user that performed each activity different network management databases had to be queried, depending on the type of network connection used to perform the activity. Such queries returned a unique identifier known as the University Login Identification (ULID) of the user that registered the device detected performing each online activity, thus it is assumed that such user was the one responsible for the activity. The ISU directory database provided the remaining information about the user associated with each ULID, namely the user's birth year, gender, major, role (student, staff, faculty), and university title (freshman, sophomore, junior, etc.). Netflow data and data collected by Packeteer did not go through either of these processes, so it contains only the recorded IP addresses.

A.3 Privacy protection

The collection of monitoring data was performed in accordance with the DCP policy guidelines. Measures to protect the privacy of monitored users include, but are not limited to, the following. Data collection was performed at ISU by ISU staff. The only output from monitoring appliances provided to researchers at CMU was an anonymized version of the collected data. To make it impossible to unveil personally identifiable information such as the ULID of a person, an IP address, or a MAC address, such fields were removed. Some were replaced by pseudonyms generated using a one-way 256-bit hashing function[9]. Both the data collection process and the generation of pseudonyms were performed in an automated fashion without human intervention, so no human would ever saw the raw data, and the keys used in the hashing function were destroyed. The monitoring and anonymization processes were controlled by the network management team at ISU, which could have access to the raw data anyway, and they were

[7] Media Access Control address, a 48-bit identifier that is (virtually) unique to every device that connects to an IP network.

[8] Dynamic Host Configuration Protocol, a protocol used by devices in a network to obtain a lease for a unique IP address and information about several other parameters necessary to connect to the network. IP addresses are assigned to requesting devices for a period of time and the lease information is typically stored in a log.

[9] Function F(K,X) → Y that, given a key K and an argument X, generates Y, a 256-bit long representation of X. F minimizes the probability that different X arguments will return the same Y. Furthermore, it is impossible in practice to map back from Y to X.

precluded from analyzing the anonymized data. CMU researchers who analyzed the resulting data were not allowed to observe raw data prior to anonymization, thus being unable to connect any of the data to a specific person, computer, or location on campus. Both the ISU Institutional Review Board (IRB) and the CMU IRB approved the research described in this paper.

A.4 Summary of collected data

Monitoring appliances collected data for about one month in each of the Spring 2007, Fall 2007, and Spring 2008 academic terms. In each of the periods, AM collected a log of the *events* described above. In Spring 2007 Packeteer collected hourly summaries with total amount of bytes and number of communication sessions entering and exiting the ISU network, broken down by protocol/application. In the latter periods, it collected one individual record per detected communication session, which is a Netflow v.5 record augmented with identifiers of the protocol/application used in the communication. Table 2 presents a brief summary of the data collected by each appliance in each period.

	Spring 2007	**Fall 2007**	**Spring 2008**
Number of people living on campus	6,544	6,764	6,763
Time span of AM data	03/31 to 04/30	09/01 to 10/04	02/12 to 04/27
Full hours / days with AM data	648 / 25	654 / 26	1,747 / 60
Number of AM events collected	24.6 million	22.2 million	58.1 million
Time span of Packeteer data	04/01 to 04/30	08/30 to 10/01	03/07 to 05/01
Full hours / days with Packeteer data	720 / 30	735 / 29	858 / 31
Number of Packeteer events collected	hourly summaries	3.3 billion	4.3 billion
Hours / days with both AM and Packeteer data	642 / 25	541 / 20	770 / 24

Table 2. Summary of data collected in the three monitoring periods by AM and Packeteer.

Data collected through network monitoring is always dependent on the vantage point where monitoring devices are deployed. In this case, both AM and Packeteer were deployed at the point where the campus network connects to commercial ISPs, which means that only communication sessions in which one party is inside the campus network and another party is in the external Internet were detected. Hence, none of the data collected by AM or Packeteer contain records of intra-campus communication sessions nor of communications routed through Internet2. Some data on intra-network traffic was gathered in Fall 2007 and Spring 2008 using Netflow collectors deployed at the entrance of several residence hall buildings. Such data was analyzed using CMU-developed software.