Chapter 5
Peer-to-Peer Networks: Interdisciplinary Challenges for Interconnected Systems

Nicolas Christin
Carnegie Mellon University, Information Networking Institute & CyLab, U.S.A.

KEY WORDS

Peer-to-peer networks; Copyright infringement; Incentives; Social norms.

ABSTRACT

Peer-to-peer networks are one of the main sources of Internet traffic, and yet remain very controversial. On the one hand, they have a number of extremely beneficial uses, such as open source software distribution, and censorship resilience. On the other hand, peer-to-peer networks pose considerable ethical and legal challenges, for instance allowing exchanges of large volumes of copyrighted materials. This chapter argues that the ethical quandaries posed by peer-to-peer networks are rooted in a conflicting set of incentives among several entities ranging from end-users to consumer electronics manufacturers. The discussion then turns to the legal, economic, and technological remedies that have been proposed, and the difficulties faced in applying them. The last part of the chapter expands the scope of ethical issues linked to peer-to-peer networks, and examines whether existing laws and technology can mitigate new threats such as inadvertent confidential information leaks in peer-to-peer networks.

INTRODUCTION

Since their inception in 1999 with the Napster file-sharing service, peer-to-peer networks have grown to become a predominant source of Internet traffic (Karagiannis et al., 2005; Basher et al., 2008). One of the reasons behind the success of peer-to-peer networks is that they have many uses. For instance, in contrast to a centralized server that would have to bear over a awarm of hosts, peer-to-peer networks facilitate information dissemination by spreading the load, thereby reducing infrastructure costs. Applications that take advantage of the cost reduction offered by peer-to-peer infrastructures include software distribution, for example open-source software such as the Linux kernel,[i] or proprietary software such as World of Warcraft patches.[ii]

As another societal benefit, peer-to-peer systems offer increased censorship-resilience thanks to their decentralized organization. Once a file is in a peer-to-peer network, it is extremely difficult, if not impossible, to completely remove that file from the network, due to both the sheer number

This chapter appears in Information Assurance and Security Ethics in Complex Systems: An Interdisciplinary Perspective, edited by Melissa J. Dark.

of machines in the network that may host a copy, and the rate at which users join and leave the network.

For all their advantages, peer-to-peer systems pose considerable ethical and legal challenges, which stem from a conflicting set of incentives among the different network participants. As a case in point, a significant share of peer-to-peer traffic has historically consisted of copyrighted materials. Indeed, for end users, the ability to download "free" content often proves tempting, particularly when considering that most consumers are either unaware of, or have a basic misunderstanding of copyright law. As a response, copyright holders, most notably the music and movie industries, have been aggressively investigating legal and technological means to reduce, disrupt, or even abolish peer-to-peer network traffic.

To add further confusion, Internet service providers (ISPs) have adopted a more ambiguous position, due to the economic conundrum they face. On the one hand, peer-to-peer applications are a driver for consumers to purchase higher levels of broadband connectivity, which translates into higher revenue for ISPs. On the other hand, the explosion of peer-to-peer traffic puts a severe strain on network infrastructure, which results in increased costs for ISPs. Consequently, some service providers have been known to treat peer-to-peer traffic as undesirable, e.g., by downgrading its priority when it enters their network, without necessarily advertising this fact to their customers.

In this chapter, we will first examine in greater detail the incentive misalignment among the actors in peer-to-peer networks. We will then briefly summarize the legal issues associated with peer-to-peer networks, especially the questions of contributory infringement and vicarious liability. In this context, we will provide an overview of the legal and economic remedies that the content industry and service providers have entertained to tackle challenges posed by peer-to-peer networks.

In the third part of the chapter, we will describe the technological arsenal that content industry and Internet service providers have been using to limit peer-to-peer traffic, as a complement to legal recourse. We will present "interdiction technologies" for which patent applications have been filed. We will distinguish between methods that target content (i.e., files) from those that target peer-to-peer hosts (i.e., actual machines). We will use this distinction to inform our discussion on the ethical and legal dilemmas that the application of these interdiction technologies presents.

In the fourth, and final, section, we will explain how the problem of controlling information flow in peer-to-peer networks far exceeds the mere realm of copyright enforcement. We will show that the assumption that the content present in the network is voluntarily introduced by end-users may be flawed. Studies indeed document that private data (e.g., credit card numbers) are often accidentally leaked due to end-user misconfiguration. Even more perniciously, recent viruses and worms have been seen to exploit peer-to-peer infrastructures to leak and disseminate private information on a large scale.

We will conclude by discussing whether or not existing interdiction technologies can mitigate these new threats. We will use these recent developments to highlight the modern ethical challenges that society faces in dealing with peer-to-peer networks.

THE ROOT OF THE PROBLEM: CONFLICTING INCENTIVES

We argue that the root causes of the rapid rise of peer-to-peer filesharing of copyrighted materials belong more to the economic realm, than to the technical realm. To be sure, technology has acted as a primary catalyst in the development of peer-to-peer filesharing – but, far from being slanted toward nefarious purposes, digital technology has also facilitated economies of scale on the content provider side. In other words, technology has been agnostic, in that it has favored both sides equally. Conversely, economic tensions between end-users, software manufacturers, Internet Service Providers (ISPs) and content providers have given rise to current conflicts.

Technology as a Facilitator

It may be useful to recall what has fueled the massive dissemination of information we see today. First, digital storage has become extremely cheap: a 4 GB USB drive, which can roughly contain a full movie in DVD format, today sells for less than U.S. $15; optical storage, e.g., DVD, is about 100 times cheaper – $20 for 100 blank 4.7 GB DVDs was a common price at the time of this writing.

Second, replication and compression technologies have benefited from massive improvements in computing power and technology over the last decades. Average processor speeds have jumped from 100 MHz in the mid 1990s to 2-3 GHz in the late 2000s, an increase over 20-fold. At the same time, the development of compression technologies, such as MPEG Layer 3 (MP3) and MPEG-4 (and its derivatives such as the DivX or Xvid codecs), have made it possible to quickly transcode, compress, and store digital content using commodity hardware. For instance, compressing a full DVD into DivX format and storing the resulting video on a CD-ROM can now be done in mere hours using off-the-shelf hardware. Added to this, network access speeds have increased – another important factor that we will discuss in more detail later – and indeed, copying and transmitting content, including copyrighted content, has never been as easy as it is today.

Content providers have also benefited from these technological advances. Industry consortia indeed encouraged the digitization of most musical and video content. The Compact Disc format, which signaled a departure from more than a hundred years of analog recordings, was pushed to the forefront by consumer electronics manufacturers, and was enthusiastically adopted by content providers due to the potential economies of scale they could realize in the manufacturing process. In 2009, replicating 1,000 records (LPs) currently costs about $1,850; producing 1,000 CDs is about half as costly,[iii] and these estimates likely include digitization of part of the process even in the case of LP pressing. Likewise, DVDs are considerably cheaper to produce than videotapes, yet are being sold at comparatively higher prices (even when adjusting for inflation) thanks to the inclusion of bonus contents and special features, made possible by the additional storage available on the medium.

In short, the development of new technology to copy, store, and transmit information rapidly has had economic benefits for all parties involved, by reducing content providers' manufacturing costs, and making it more practical for consumers to create data back-ups.

Digital Replication and the Change in the Rules of Engagement

While digitizing media and replicating digital content has become possible for the average user since the mid- to late-1990s, it could be argued that, at that point, nothing had changed much since the days of the Sony Betamax recorder. The Betamax allowed individuals to record programs of their choosing and replay them at a later time, presaging the considerable development of VCR devices.

Initially, the content industry was strongly opposed to Betamax technology, and in fact went all the way to the U.S. Supreme Court to make the case that the technology was a vital threat to their revenue model (Boyle, 2008). However, the U.S. Supreme Court disagreed, the Betamax paved the way for VCRs, and the content industry eventually realized that, far from threatening its business model, improved replication technology could in fact lead to new revenues, i.e., by creating a secondary market for movies through their video releases. In 2009, the size of the home video market is estimated to be about $7.5 billion,[iv] down from about $9.2 billion in 2001 (Dana & Spier, 2001), but still considerable. On the other hand, the risk posed by the increased ease of unauthorized replication of content proved considerably exaggerated and was estimated by the industry itself at around $100 million in the 1980s.[v] This number is significant but it is small compared to the size of the whole industry, even when accounting for inflation and for the differences in the market size between the late 1980s and the early 2000s.

Thus, one would think that the development of digitization and peer-to-peer technology would be more warmly received by the content industry, which could certainly take advantage of novel revenue opportunities the way it took advantage of the ubiquity of VCRs in the 1990s. However, this has not been the case, and, as we will discuss later in the chapter, peer-to-peer technology in particular has been at the center stage of a number of legal battles. The key reason is that, while digitization in itself did not fundamentally change the rules of the game,[vi] the advent of the Internet provided a massive distribution channel that did not exist before.

At the time of the Betamax recorder, copyright infringement was extremely limited in scale. While replicating a movie or a TV program was made much easier, *mass diffusion* of the replica was impractical for all but the most determined people. Sending copies of a movie for instance, required a person to manually replicate the movie, and then to either send the replicas by mail, or to sell them "under the table." Mailing the copies was costly and time consuming. Selling them under the table posed non-negligible risks such as police raids and easy prosecution. By providing a way to interconnect millions of computer systems, the Internet lowered these diffusion costs to zero. What was missing was a technology that could harness this ubiquitous connectivity and turn it into a massive diffusion channel for digital content; which is precisely what peer-to-peer software provided.

The Modern Five-way Tussle

With replication *and* distribution costs of content nearing zero even for the average user, the content industry perceived a threat to its existing business model. Industry set out to eliminate this threat, both by fiercely combating the development of peer-to-peer technology in courts (and, as collateral effect, changing social norms), and by trying to limit its impact through technology. To paraphrase Clark et al. (2005), peer-to-peer filesharing technology set the ground for one of the major "tussles in cyberspace" between the different actors involved.

It would, however, be a mistake to reduce the peer-to-peer tussle to a tension only between content providers standing to lose their business revenue, and unscrupulous end users interested in obtaining free content. Specifically, the tussle taking place involves at least five actors, all with different incentives: content providers and end users, as noted earlier, but also consumer electronics manufacturers, software manufacturers, and Internet service providers.

End users have, in general, a fairly simple objective. They want to obtain the content they are interested in, at the smallest cost possible. Cost, here, however, does not solely mean pure monetary cost. For instance, obtaining a movie for free, but having to wait for 10 weeks for the download to complete may be much less satisfactory than paying a nominal price in order to obtain the same movie immediately, a phenomenon economists call "instant gratification."

As a case in point, in Europe and Asia, a large portion of peer-to-peer traffic consists of popular U.S. television shows: These shows are available on peer-to-peer networks the day after they are broadcasted, but, outside of the United States, are not available for legal downloading for several months. Rather than wait for the legal alternative, a large number of users prefer immediate downloads at the potential expense of copyright infringement.

To summarize, most users' preferences appear to be driven by a combination of monetary factors and convenience.[vii] Economists and technologists have been trying to express user preferences through utility functions, which formalize how users value different situations. Precisely characterizing utility functions for peer-to-peer users may be quite challenging (Golle et al., 2001; Feldman et al., 2004; Christin and Chuang, 2005), but understanding what determines the evolution of the utility function may be all that is needed for a business to be economically successful. Indeed, adapting to the end-users' utility function proved successful in the past. At a time where sales of VHS recordings were prohibitively expensive for most (about $40-45 on average in the early 1980s), video rental stores emerged and were extremely successful thanks to their affordability and convenient service.

Particularly relevant to the discussion in this book is the question of ethical behavior as opposed to, or in intersection with, utility. Most users generally refrain from engaging in unethical activities, which rules out a large number of possible strategies. In particular, despite the potential for instant gratification and low cost, very few people walk out of a video store with stolen DVDs. On the other hand, most users do not view it as necessarily wrong to download music or movies from a peer-to-peer network. While the recording and movie industry have been arguing that downloading copyrighted materials is akin to stealing from a store, a large number of users have a different perception (Shang et al., 2008) – and instead view peer-to-peer downloads as closer to recording contents from a radio program or a TV program (Easley, 2005). As a result, users do not rule out downloading as a possible strategy in evaluating their own utility. Why do some

users perceive walking out of a store with a DVD as stealing, but downloading a movie from a peer-to-peer network as acceptable?

Content providers such as movie studios or the recording industry argue that the rapid rise of peer-to-peer software combined with the ubiquity of digital media results in massive infringement, which in itself yields staggering revenue losses, evaluated in the vicinity of U.S. $1 billion per year by Liebowitz (2006). However, this estimate needs to be viewed with caution. First, others, such as Oberholzer-Gee and Strumpf (2007) reach different conclusions that do not seem to support the existence of sizeable losses. Second, even if we accept the premise of considerable losses, loss calculations are based on business models currently in place, and do not factor in new revenue streams made possible by the development of peer-to-peer technology. As a less controversial proposition, we postulate that, regardless of the magnitude of the realized losses, there is a certain perceived risk to the revenue stream of content providers, caused by the significant economic shift that peer-to-peer infrastructure has induced. The recent development of new business models, such as iTunes, which harness new diffusion capabilities, has been spearheaded by consumer electronics manufacturers like Apple and software companies such as RealNetworks or Microsoft, not by the content industry. The content industry still worries about the risks associated with a change in business model, and there is, in fact, still some reluctance from content providers to embrace such new business models, as evidenced by the much-publicized rift between Apple and Universal.[viii]

The plot thickens when we consider *consumer electronics manufacturers*. These companies, which, as we mentioned, were opposing the content industry in the Sony Corp. v. Universal City studios "Betamax case," actually have a much more conflicted set of incentives. On the one hand, they fully stand to benefit from digitization and peer-to-peer networks. Indeed, digital replication requires faster computers and more disk space; at the same time it provides content, which makes portable hardware such as portable MP3 players useful. In that respect, the incentives of consumer electronics manufacturers and *software manufacturers* are aligned. As a particularly striking example, Apple's commercials ("Rip, mix, burn") put digitization as the new "killer" application to sell more software; but this software itself demands state-of-the-art hardware. In this particular case, the consumer electronics manufacturer and the software manufacturer are one and the same company, but more generally there may be a certain level of collusion (voluntary or not) between both entities. At the same time, hardware manufacturers cannot antagonize the content industry completely. Content is indeed, as discussed above, what most consumers want in the end. Hardware without content is of little interest to most segments of the population, and as such, hardware that does not support enough content is doomed to fail, regardless of its intrinsic qualities. Examples abound, but, perhaps ironically, a good example is the Betamax system itself, which was edged out of the market by the VHS standard, as soon as content providers decided to primarily use VHS.

Software manufacturers have a similarly complex set of incentives. We define peer-to-peer software in a large sense, comprising both programs running on personal computers, and indexing services running on remote servers, such as offered by YouTube, MiniNova or The Pirate Bay. Peer-to-peer software has value to customers, as it can reduce the cost of obtaining content (copyrighted or not) and that makes it a "must-have" program for end users as shown in surveys (Good et al., 2005). As such, peer-to-peer software can be relatively easily monetized. Indexing

services can be made accessible through paying subscriptions, or can derive revenue for online advertising. Likewise, software programs can be sold (but then, at the risk of itself ending up being copied and distributed over peer-to-peer networks), or can produce advertising revenue to its manufacturer, as was, for instance, the case of the KaZaA peer-to-peer system. While a "free" version of KaZaA existed, it came with bundled software, including adware, whose providers paid Sharman Networks, the company that developed KaZaA, a fee to tie their programs to KaZaA. Interestingly enough, KaZaA itself showed that software, much like hardware, is only viable if content is available. As soon as content started to be difficult to access in KaZaA (due both to the legal threats, and the technological countermeasures, which we will discuss later), users switched to different peer-to-peer systems. To sum up the incentives of software manufacturers, content brings visibility; visibility can then foster business models even if the software or service itself is free. Outside the peer-to-peer realm, Google is the primary example of a service that has been immensely successful in generating derived revenue despite offering its core search service free of charge.

Finally, *Internet Service Providers* (ISP) have an even more ambiguous position. Peer-to-peer software requires consequent network capacity to be useful, particularly in the case of applications, such as BitTorrent, that reward contributions to the network. The situation was different in earlier decades. Relatively low bandwidth applications such as web access and email made up the bulk of Internet traffic. Higher bandwidth broadband connections like DSL and cable were in general not necessary for web or email, so that, by the end of the 1990s, Internet Service Providers had a hard time driving their customers to purchase broadband connectivity. Hence, although video compression technology was available, the lack of access network capacity led to very poor quality at the end-hosts, which in turn, hindered the development of a market for Internet-based video content.

Peer-to-peer networking changed the game as people bought increased network access in order to be able to download peer-to-peer content. As a matter of fact, the ability to download movies and multimedia content was heavily advertised. While stopping short of encouraging users to participate in copyright infringement, ISP advertisements have been, at the very least, ambiguous; encouraging users to sign up for "fast downloads," including access to "movies and music."[ix]

As we will discuss in the next section, the content industry felt that Internet Service Providers were actually partly responsible for copyright violations, and tried to hold them liable for it. While attempting to hold Internet Service Providers liable for copyright violations only had mixed success in court, one outcome was clear. It made ISPs aware that an unconditional support of peer-to-peer infrastructures would essentially alienate the content industry.

In addition, from an economic standpoint, peer-to-peer systems are not necessarily lucrative for Internet Service Providers. If, indeed, peer-to-peer partly fostered the adoption of faster access links, this adoption finally allowed video services (e.g., YouTube) and other multimedia applications (particularly photo sharing) to soar, and led to a greater demand for access capacity from end users. In other words, peer-to-peer systems were great as a *bootstrapping* mechanism to initially drive the demand for bandwidth, but their impact on bandwidth demand is no longer as predominant as it once was. Furthermore, a profitable ISP is an ISP that sells access capacity that is not used. Assume that a given ISP sells 100 Mbps access to 100 customers. If all users transmit

simultaneously up to the promised rate, the ISP needs 10 Gbps of capacity. If, however, most users are idle most of the time, the ISP is able to multiplex and significantly reduce its infrastructure costs. If customers use, on average 10% of their access capacity, the ISP should be able to only provision a 1 Gbps link. Web traffic lends itself well to multiplexing, because it consists mostly of short bursts of data transfers. On the other hand, peer-to-peer traffic, which typically consists of large file transfers[x], tends to heavy and prolonged usage. This is partly due to the TCP congestion control protocol, which tends to grant larger shares of network capacity to long-lived flows (Katabi et al., 2002). So, while peer-to-peer systems might have played an interesting role to attract new customers, their extensive use actually hurts the economic bottomline of Internet Service Providers, because the bandwidth usage of peer-to-peer networks does not fit the provisioning model used by ISPs for network planning.

To add to the ambiguity, recent years have seen a blurring of the roles of content providers and Internet Service Providers. For instance, consider the case of an entity that is both a cable service provider, as well as a broadband service provider. As an ISP, this entity stands to gain from the promise of free content offered by peer-to-peer systems; yet, if the content in question includes popular TV series, the cable division is at risk of losing customers. Such a situation is far from unique due to the increased integration of data services: Comcast in the United States, Orange in Europe, and Virgin Media in the UK are among the many operators that offer both services.[xi]

With their complex incentive structure and possible conflicts of interest, it is not surprising that ISPs' positions have not generally been transparent to customers. In a well-publicized incident (Eckersley et al., 2007), it was discovered that Comcast was aggressively filtering BitTorrent traffic without having previously informed its customers. Worse, Comcast was using strategies akin to denial-of-service attacks. Public outrage persuaded Comcast to eventually relent. Other ISPs (e.g., Verizon) have been forced to turn over names of potential copyright infringers to law enforcement agencies despite fighting court injunctions to do so. In these cases, end users felt betrayed by the ISP, whom they thought was "on their side" when in fact the reality was markedly different.

A BRIEF REVIEW OF LEGAL ISSUES AND REMEDIES IN P2P NETWORKS

In the five-way tussle just described, all participants try to tip the scale in their favor using different strategies. Of particular interest to the discussion in this chapter, are the strategies employed by the content industry to achieve its goals. Defensive strategies against the perceived threat of peer-to-peer networks can be categorized into economic, legal, and technological strategies. Economic defenses are often overlooked as an old saying goes, "You can't beat free." That is, if content replication and diffusion has a negligible cost to end users, it will be difficult for the content industry to provide economic incentives against copyright infringement. While such propositions neglect the possibility that new business models, such as models in which revenue comes from advertising derived from content, they have colored the perception of the content industry that legal and technological recourse are the most promising avenues to win the tussle. In this section, we provide a brief summary of copyright law, and discuss the liability issues at hand, before summarizing the legal actions pursued by the content industry. We point out that this section is heavily colored with U.S. law. Other countries present differences that, in

the interest of brevity, we do not discuss here. For an abridged treatment of international copyright laws, we refer the interested reader to Gassner (2005).

Relevant Elements of U.S. Copyright Law

In the United States, the Copyright Act of 1976 defines copyright. Copyright applies to works of authorship, such as literary works, including software, which is considered as literary work in this context, pictorial work (e.g., movies, photographs), and musical works. The copyright owner is usually the author of the work, with the caveat that the employer of the author may be entitled to copyright as well, if the work has been performed "for hire." In the United States, copyright automatically applies, by default, to any work of authorship that is (1) an original work, and that is (2) fixed in a tangible medium. However, authors must register to be able to sue. Copyright owners are, among other things, granted exclusive rights to prepare derivative works (including translations), and to publicly perform, display or broadcast the work. Directly related to the discussion in this chapter, copyright owners are also granted *exclusive rights to reproduce work in copies and to distribute copies to the public*. Furthermore, copyright owners are also provided with some rights to control the acts of those who facilitate or contribute to copyright infringement.

The Digital Millennium Copyright Act does not fundamentally change the rights of the copyright holders, but amends the Copyright Act of 1976 to (1) criminalize production and dissemination of technology that can circumvent measures taken to protect copyright (e.g., it becomes illegal to defeat copy-protection mechanisms for instance) and to (2) heighten the penalties for copyright infringement on the Internet.

Limitations and Fair Use Doctrine

There are limitations on the set of exclusive rights granted to copyright owners. These include personal backup copies and library-archival copying along with specific exceptions such as the right to play a radio in a restaurant without it being considered a "public performance." Most importantly, the concept of "fair use" defines a large category of exceptions to copyright. The main challenge is that what constitutes fair use is relatively open to interpretation. Indeed, prior to the Copyright Act of 1976, fair use was only defined by common law in the United States. The Copyright Act of 1976 made the definition a bit clearer, by requiring a balancing test consisting of at least the following four criteria to determine whether a specific usage of a work characterizes as fair use:

1. the purpose and character of the use, including whether such use is of a commercial nature or is for nonprofit educational purposes;
2. the nature of the copyrighted work;
3. the amount and substantiality of the portion used in relation to the copyrighted work as a whole; and
4. the effect of the use upon the potential market for or value of the copyrighted work. (Copyright Act of 1976, p. 107)

In general, fair use applies primarily to scholastic citations for reference or critique. The last criterion proves particularly interesting to our discussion, because it ties in economic value to the notion of copyright.

Primary and Secondary Liability

Most important to the issue at hand in this chapter is the question of liability. When a copyrighted work is illegitimately distributed over the Internet, who is legally responsible? While this question, as a whole, remains an open problem as argued by Taipale (2003), it is useful to distinguish between two classes of liability in copyright infringement: primary liability, where the perpetrator directly violates one of the exclusive rights granted to the copyright holder; and secondary liability, where the perpetrator does not directly abuse these rights, but either facilitates infringement by providing active support to the person infringing (*contributory infringement*), or indirectly exerts benefits, e.g., obtains monetary rewards, from infringement (*vicarious infringement*).

This distinction raises an ethical problem. As our brief economic discussion showed, one can make the case that, to some extent, a large number of parties indirectly benefit from peer-to-peer traffic and could therefore be liable for secondary infringement. Unfortunately, the law only imperfectly specifies whether punishment is needed, who should be punished, and to which extent.

Liability and P2P File-sharing

Copyright holders have two legal remedies at their disposal: (a) going after the primary infringers, i.e., the end users who exchange copyrighted material over peer-to-peer networks in violation of the exclusive rights of the copyright holders, and/or (b) going after entities that can be held liable under secondary liability criteria. These include software manufacturers (including designers of online content indexing services such as The Pirate Bay, MiniNova, YouTube), consumer electronics manufacturers, and Internet Service Providers.

At first glance, it seems that identifying primary infringers is relatively easy. After all, end users downloading or uploading music to peer-to-peer networks do so without authorization, and could easily be held liable for it. However, the issue is not as simple.

First of all, some, including Charles Nesson, a Harvard Law School professor, have argued that peer-to-peer filesharing essentially falls under fair use, as the four criteria are only illustrative and not a complete test.[xii] This argument has been used in defense of Joel Tenenbaum, a student who was sued for sharing 30 songs on the KaZaA network. This defense has gained little traction, and has been ultimately unsuccessful – Tenenbaum was convicted and fined $675,000,[xiii] and reportedly plans to appeal the verdict.

A much thornier issue lies in the notion of control, i.e., whether willful infringement can be proven easily. On the uploading side, suppose a user participates in a peer-to-peer network, and mistakenly shares a portion of his or her hard drive where legitimately purchased music is stored. Should the user be held liable? As we will discuss in a subsequent section, misconfigurations are far from being an exception, as users frequently share data unknowingly, as shown by Good & Krekelberg (2003). If we allow the premise that users may not be held liable for configuration mistakes, the question becomes one of showing active intent to share data – which is considerably more difficult to prove.

If, on the other hand, we consider that users should be in complete control of their systems at all times, we then make users responsible for all possible security vulnerabilities present on their machines, as a single security vulnerability can give full access to the system to a remote unauthorized party. However, most security vulnerabilities are due to faulty software, rather than user mistakes, and liability is, here too, difficult to assess. If Software A is exploitable, e.g., because no software patch exists or because installing a patch is beyond the expertise of most casual users, is it the user's responsibility to ensure Software A does not run on their computer? What if Software A is actually a critical piece of software, e.g., the operating system?

On the downloading side, one could think that the responsibility should be clearer. After all, downloaders are copying work that they do not own. Yet, the same issues of control that apply to uploading entities can apply to downloading entities. That is, it is very difficult to prove with absolute certainty that a specific user was indeed in control of the machine when a download took place. For example, one can easily show that a machine on Bob's computer network downloaded a song – but that does not necessarily mean that Bob was the person downloading the song. It could as well have been Alice, his wife, using Bob's computer, Mallory, a malicious third-party who gained unauthorized access to Bob's computer, e.g., through a virus, or even Eve, a neighbor who simply took a free ride on Bob's unprotected wireless network, and made it look like, from the outside world, Bob was actually downloading the file. Actually proving that Bob was responsible for the download is considerably difficult.

In the context of secondary liability, the discussion is even more complex. Consumer electronics manufacturers have, thus far, by and large escaped the ire of content producers in the context of peer-to-peer networks. They argue, for instance, that it is difficult to say that a vendor of MP3 players vicariously benefits from peer-to-peer systems, as there are plenty of other, legal, sources of MP3 music, including personal backups of CDs and online music stores. On the other hand, Internet Service Providers and software manufacturers have faced legal pressures from the content industry. Both have been accused of contributory and vicarious infringement. As discussed in our section on incentives, ISPs profit to some extent from peer-to-peer networks by enticing users to purchase additional broadband connectivity. Peer-to-peer software manufacturers stand to gain from the use of their software or service. And, as we have mentioned before, content is the main driver for use. Here again, however, showing actual intent to benefit or contribute to infringement is relatively difficult, and has not always been successful.

Legal Strategies against Peer-to-Peer Networks

The content industry initially focused its legal strategies against the peer-to-peer software providers. Shortly after the Napster service was started in 1999, several companies, including A&M Records, sued Napster on grounds of vicarious and contributory infringement. Napster was eventually shut down following the decision of the 9th Circuit Court of Appeals.[xiv]

Subsequently, various content providers sued software companies that were providing peer-to-peer software. Some of these cases gained extreme visibility. MGM v. Grokster/Streamcast,[xv] went all the way to the United States Supreme Court. In Universal Music Australia v. Sharman Networks,[xvi] about thirty copyright holders sued the makers of the KaZaA software. The Federal Court of Australia ultimately decided the case.

The main lessons are that, the courts have generally sided with copyright holders, and have held software manufacturers responsible for secondary infringement. However, the court decisions have also outlined the difficulties of rendering judgment in these cases. For instance, the MGM vs. Grokster decision rejected the fair use argument from Grokster and its allies, stating "We hold that one who distributes a device with the object of promoting its use to infringe copyright, as shown by the clear expression or other affirmative steps taken to foster infringement, is liable for the resulting acts of infringement by third parties."[xvii] Rather than bringing finality to the debate, this decision actually shifted the problem to defining "clear expression" and "affirmative steps." If Streamcast and Grokster had been less aggressive in their marketing, would the outcome have been the same?

Also, courts have repeatedly refused to declare any specific piece of software illegal. That is, the use of a peer-to-peer system may be illegal, and may in fact cause the designer of the peer-to-peer system to be held liable for secondary copyright infringement in some cases (e.g., if they have taken the "affirmative steps" to promote infringement), but writing peer-to-peer software, in itself, is not illegal. Thus, when faced with decentralized networks supported by free software like eMule or Gnutella, the content industry could not simply sue the individuals that had contributed to the software (many of whom are outside U.S. jurisdiction). They instead applied legal pressure on the other two actors in the tussle: ISPs and end users.

In the United States, ISPs are conditionally protected from secondary liability by a safe harbor clause (17 U.S.C. § 512, 1998); namely, they cannot be held liable for data that merely pass through their networks, as long as they do not make a profit from it, are unaware of the violation before being notified, and respond immediately to "take-down" notices. We note that this law predates the rise of peer-to-peer systems, and was intended to target services where, for instance, an ISP providing web hosting had some of its customers post copyrighted materials on their webpages. In the context of peer-to-peer networks, this law has led the content industry to send take-down notices to ISPs, asking them to provide a list of infringers. Indeed, the only information available publicly about infringers is their IP addresses. Mapping an IP address to an actual individual is not possible without the help of the ISP that owns the IP address. In many cases, ISPs have fought off such requests, citing privacy concerns, but most have been forced to comply.

Armed with lists of potential infringers, copyright holders have then taken steps to directly sue individuals as primary infringers. This strategy is interesting, in that it is more of an economic strategy than a legal one. Indeed, the key idea is to impose significant monetary damages on infringers in order to deter them from using peer-to-peer software. For instance, Universal requested that Marie Lindor pay $750 in punitive damages per track she downloaded, when the actual damages were estimated at about 70 cents per song.[xviii] Most of the cases were settled out-of-court for a few thousand dollars in total. Among the cases that did in fact go to court, was that of Jammie Thomas, who was fined $222,000 for sharing over 1,500 tracks on the KaZaA network. The judgment was later overturned and a mistrial declared. The new trial, finally held in June 2009 resulted in an increase of the penalty to U.S. $1.92 million, or about $80,000 per song.[xix] Thomas has appealed. Both Thomas and Tenenbaum, whom we mentioned earlier, face bankruptcy if the penalties are held on appeal.

While the magnitude of the potential fines has had a deterrent effect on some users, Bhattacharjee et al. (2006) question the actual effectiveness of the deterrent on users sharing small numbers of files. They claim that legal action had overall no noticeable impact on file availability. In addition, the content industry also suffered tremendously from bad publicity linked to suing "grandmothers,"[xx] and has, so far, only won two court cases (the aforementioned Thomas and Tenenbaum cases), whose appeals were pending at the time of this writing.

Finally, a more recent line of legal attack has been to put pressure on owners of websites, such as The Pirate Bay,[xxi] that support searches for peer-to-peer networks. We grouped these content indexing service providers with software manufacturers earlier, but while legal action against program designers has diminished, content indexing services have been increasingly targeted. The problem with this strategy is that such websites can be replaced relatively easily, and suing the owners, as in the case of The Pirate Bay can actually turn into a public relations nightmare. As a case in point, in response to the lawsuit against The Pirate Bay, a "Pirate Party" was created in Sweden and ran in the 2009 European Parliamentary Elections on a platform advocating copyright and patent law reform.

Having a political party solely running on a copyright reform platform outlines one key concern with the legal issues outlined thus far: the ethical dilemma they present. First, users are still not convinced that downloading copyrighted materials off the Internet is "wrong." In fact, whether the strategy of suing individuals for primary copyright infringement has been a success is open to interpretation, as peer-to-peer transfers have not slowed down much (Kargiannis et al., 2004), and the legal strategy came at a public relations cost. The fact that peer-to-peer activity persists despite the increased legal threats to users, and repeated attempts at education through commercials and advertisements, shows that the public does not perceive copyright infringement as a serious crime (Easley, 2005), and that social norms are actually not playing in favor of copyright holders. This ethical question is the primary reason behind the rise of the Pirate Party and similar movements that argue that, instead, it is the law that is "wrong" and must be changed. Second, as discussed above, many different parties can technically be held liable to some extent. The law does not distinguish between the different degrees of liability, so it is up to the copyright holders and the courts to decide whom to punish, and to which extent. Compared to the legal argument of what constitutes infringement, both ethical discussions of where the responsibilities lie, and to which extent should each party be held responsible, have received little attention. We do argue, however, that, in absence of stronger social norms, the situation is unlikely to change. That is, only the most risk-averse users will refrain from exchanging copyrighted materials over peer-to-peer networks; others will run a simple cost/benefit analysis as described before, and may not even necessarily be convinced that the activities in which they are engaging are of questionable legality.

USING TECHNOLOGY TO MITIGATE THE P2P "THREAT"

Well aware of the limits of a defensive strategy resting on pure legal grounds, copyright holders have also been aggressively pushing technology to alleviate the perceived threat posed by peer-to-peer networks. Here, the objective is to prevent the diffusion from copyrighted materials already present in the network by either making content difficult to find, or difficult to download.

We purposely do not discuss Digital Rights Management technologies aiming at preventing content from being copied in the first place. Indeed, such technologies are generally flawed, in that no DRM technology is 100% effective. As argued by Felten[xxii], all DRM technologies are vulnerable to the fact that content has to be decoded at some point to be usable, and that decoded content can be replicated. However, with the advent of peer-to-peer technology, DRM has to be foolproof to be viable as a protection mechanism for copyright holders, as a single copy can be propagated to millions of users immediately. As a result, while content providers have been using DRM technologies in an effort to reduce copyright infringement, they have been relatively unsuccessful thus far, and have, by and large, combined any DRM protection with defenses directly targeted at peer-to-peer infrastructures.

There are two major components to any peer-to-peer infrastructure: indexing, which tells peers where desired content is located, and transmission, which is the mechanism by which peers acquire content once it has been located. As illustrated by the Napster episode, compromising either of the components (in Napster's case, the indexing infrastructure) can result in a quick demise of a given peer-to-peer network.

In this section, we first discuss technologies that primarily aim to incapacitate specific hosts, in charge of indexing or serving data that can violate copyrights. The interesting aspect of this group of techniques is that some implementations are essentially a form of denial-of-service (DoS) attack, whose legality is itself suspect, as we discuss in more details below.

We then look at an alternative proposal, also in use, which targets search mechanisms by making content hard to find, but without assaulting specific hosts. This technique, called "poisoning", and studied in depth by Christin et al. (2005), relies on drowning peer-to-peer searches in a large amount of useless information. Rather than compromising network participants, poisoning is usually done by voluntarily injecting extraneous data in the peer-to-peer network. As such, poisoning may circumvent some of the legal doubts surrounding denial-of-service, and can additionally permit selective targeting (e.g., of a specific artist, song, or movie), but it comes at the expense of potential overload in the network.

A Primer on Peer-to-Peer Structure

Before we delve into the technical details of interdiction technologies, it may be useful to recall how most peer-to-peer services operate. As discussed in Christin et al. (2005), a large number of popular peer-to-peer networks, including Gnutella, eDonkey, and FastTrack, use two-tiered hierarchical topologies. In these hierarchies, nodes are split between leaf nodes and hubs – called "ultrapeers" in Gnutella, "supernodes" in FastTrack, and "servers" in eDonkey. Leaf nodes only maintain connections to a handful of hubs, while hubs maintain connections with hundreds or thousands of leaves, and with many other hubs.

Figure 1 shows the relationship between hubs and leaves. Each hub node serves as a centralized index for the leaf nodes to which it is connected. Whenever a leaf node issues a query, the leaf node first sends the query to the hub(s) to which the leaf node is connected. If the item requested is not present in the index maintained by the hub(s), the hub forwards the query to other hubs. BitTorrent, which was not studied in Christin et al. (2005), has a similar two-tiered structure, where hubs are referred to as "trackers." Originally, BitTorrent supported a centralized tracker,

which was hosted on a single machine. The more recent versions of the BitTorrent protocol support "decentralized trackers," which spread the tracking tasks over several hosts that communicate with each other.

Figure 1. Peer-to-peer network structure.

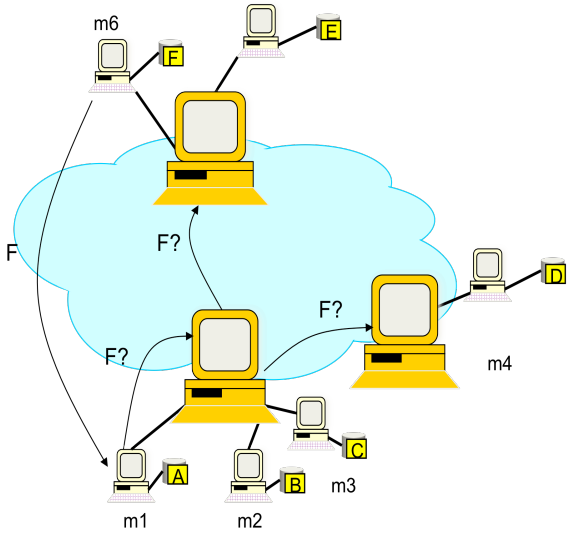

Most current peer-to-peer networks use this kind of two-level hierarchy. A notable exception are peer-to-peer networks solely based on distributed hash tables (DHT), such as Kademlia (Maymounkov & Mazières, 2002). DHT networks are decentralized, flat networks, which rely on a hashing function associating each item (e.g., search term, file) to a hash value. Items are then stored over one or several nodes according to their hash values, which greatly increases the efficiency of look-up and indexing. The specifics of the hashing and repartition functions used depend on the specific network under consideration, and have as an objective to provide better load balancing and faster search compared to hierarchical networks. An example of a DHT network is the Overnet network, but, interestingly, Overnet shares content with the eDonkey network so that the hierarchical indexing is indirectly used in Overnet as well.

Proposed Interdiction Technologies

Among the many companies engaging in "interdiction" methods, Macrovision Corporation submitted a number of interdiction technologies to the U.S. patent office (Moore et al., 2004, Levin & Disher, 2004). Rather than disserting on the merits or originality of the patent applications, we point out that the interesting feature of these documents is that they formally express what many users suspected was taking place, without having firm evidence of it actually happening. We summarize the technical content of these patent applications briefly considering the potential impact of the described techniques on existing peer-to-peer networks.

We can categorize the proposed mechanisms using the three types of attacks they implement: man-in-the-middle attack, a poisoning attack, similar to what Christin et al. (2005) and Liang et al. (2005) described in their academic works, and a partitioning attack. These terms are borrowed from the network security jargon, and are not used *as is* in the patent applications.

Man-in-the-Middle Attack

A man-in-the-middle attack is an attack on two entities communicating, whereby an intruder independently makes connections with both entities, and impersonates the other entity to both communicants. A similar principle is behind the first set of mechanisms proposed to control peer-to-peer traffic. The key idea here is to infiltrate the network with malicious nodes, which reroute query traffic to a third-party controlled server and database. When responses to a query match a

record indicating a protected (copyrighted) item in the database, the server can modify the query results to:

- point to invalid peers, which can be either non-existent IP addresses, or IP addresses of hosts not participating in the network, or
- point to incorrect files, that is, either decoys hosted by the server, or files present in the network, which do not match the requested content; e.g., someone trying to download a song by Madonna could end up receiving a picture of a dog instead.

Details on how, in practice, traffic can be rerouted unbeknownst to the users are discussed in a related patent application (Levin & Disher, 2004).

Partitioning Attack

A second set of mechanisms aims to isolate nodes that serve infringing content, by surrounding them with malicious nodes in the peer-to-peer topology. The malicious nodes can then "quarantine" each offending node by, for instance, dropping all query traffic coming to that node, which results in making the infringing content unavailable to the rest of the network.

The interesting aspect of this attack lies in the implementation details the patent application discloses. In practice, one has to first disconnect existing connections between peers to insert malicious nodes at the appropriate locations in the network. For the disconnection phase, the partitioning attack includes techniques such as (Moore et al., 2004):

- "overwhelming the capacity of [the targeted] node[…]'s connection to [its neighbor] by bombarding it with messages or requests it must parse", or
- "eliminating or disconnecting N1 [a node adjacent to the targeted node] […] by exploiting a known defect in the client software application."

While the specific techniques used to achieve these goals are not disclosed in the patent application, they implement what is commonly known as a denial-of-service attack. For instance, the former could, in practice, be implemented by techniques such as TCP SYN flooding similar to those described in Lemon (2002). It is quite surprising to find such statements in a U.S. patent application, as they seem to be in violation of U.S. law. More precisely, the Computer Fraud and Abuse Act (CFAA, 18 U.S.C. § 1030 (1984)) makes it an offense to "intentionally access a protected computer without authorization, and as a result of such conduct, cause damage," which makes the first technique of flooding a host with unsolicited requests illegal. The CFAA also makes it illegal to "knowingly cause […] the transmission of a program, information, code, or command, and as a result of such conduct, intentionally cause […] damage without authorization, to a protected computer," (CFAA, section (a)(5)(A)) which seemingly makes the second technique equally illegal.

If the peer happens to be in the United States, both of the techniques described above would therefore appear to be in violation of CFAA. Perhaps more puzzlingly, peers might not have to serve infringing files to be targeted, according to the description of the attack. Being "in the wrong place, at the wrong time," i.e., too close topologically to an infringing peer is sufficient.

Poisoning Attack

A third category of mechanisms detailed in the patent application implements poisoning attacks. A poisoning attack is characterized by nodes serving bogus, synthesized decoys whose metadata matches the description of infringing content, but which contain unsuitable data (e.g., white noise). Such attacks were first proposed in a different patent (Hale & Manes, 2004).

Whether the poisoning attack discussed is really new, especially compared to the technique discussed in (Hale & Manes, 2004), is unclear, but some of the proposed variants are interesting:

- "File transfer attenuation" consists of malicious nodes serving (potentially correct) files with a transmission speed decreasing as a function of time, ultimately halting completely right before the whole file has been uploaded. Here, the goal is obviously to maximize user frustration.
- "Hash spoofing" consists of advertising files (or file chunks) with a hash corresponding to a valid file, when in fact the file (or file chunk) is corrupted, or contains bogus content. Contrary to a common misconception, hash spoofing does not necessarily require breaking strong cryptographic primitives. In most peer-to-peer networks, the hash is indeed merely advertised along with the metadata of the file. Hence, simply advertising fake metadata suffices to implement hash spoofing. Since the advertised and the actual hash do not match, hash spoofing generally results in a peer-to-peer client inferring that a network error occurred during transmission of the file. In the case of a naive implementation, hash spoofing can lead the client to request the corrupted file again, with the same results. Here again, the primary goal is to maximize user frustration.

Comments

Most of the mechanisms described in Moore et al. (2004) rely on the ability to infiltrate the peer-to-peer network with malicious hubs, in order to manipulate the results peers obtain in response to search queries. In highly distributed peer-to-peer filesharing systems (e.g., KaZaA/FastTrack, and Gnutella), where promotion to hub is relatively frequent and generally not monitored, the proposed interdiction mechanisms could be particularly effective.

However, man-in-the-middle attacks are much more difficult to carry out in more centralized networks, such as eDonkey, which rely only on a few dozens of supernodes ("servers"). Indeed, the modest number of servers makes newcomers particularly obvious, and network infiltration problematic. In contrast to FastTrack and Gnutella, effectively infiltrating networks like eDonkey is further complicated by the fact that users can manually choose to which server(s) their peers connect.

Likewise, setting up partitioning attacks in a highly centralized network reduces to launching a denial-of-service attack on the infringing node or the server to which it connects. In the extreme, in a peer-to-peer filesharing system such as BitTorrent, where a centralized tracker is in charge of coordinating the progress of all transfers, taking down the tracker might be the most effective strategy.

On the other hand, it is interesting to note that some of the mechanisms we classified as "poisoning attacks" (file transfer attenuation and hash spoofing) can apply to most (if not all) existing peer-to-peer networks. Their effectiveness, however, once again depends on the ability of the poisoning entity to infiltrate the peer-to-peer network.

Also, except for mentioning that malicious nodes avoid expulsion from the network by frequently changing IP addresses, the patent application (Moore et al., 2004) contains very little detail concerning the defenses that malicious peers need to implement to evade detection. With the proliferation of automated blacklists of peers known to be malicious, and ever improving distributed reputation systems used to rate the integrity of other peers and the contents they transmit, e.g., (Walsh & Gün Sirer, 2006), one may wonder for how long the techniques presented in this patent application will be effective.

Finally, the proposed technological defenses also present considerable ethical dilemmas. Even setting aside the questionable legality of some of the defenses presented here, how do we ensure that these defenses are used only to target "undesirable" traffic? More specifically, what defines "undesirable" traffic? Our discussion so far has mostly revolved around copyright issues, but the techniques presented here can apply to any kind of peer-to-peer traffic. Defining which traffic is "undesirable" is likely to be based on social norms and legal grounds, which differ from one country to the next, while on the other hand peer-to-peer networks are usually not contained within geopolitical borders.[xxiii] For instance, contents considered subversive in China may be protected under the First Amendment in the United States. So, is it legitimate for a Chinese computer to poison networks hosted in the United States? Likewise, copyright laws in the United States may protect some materials that are not protected in some other countries, or whose protection has a much narrower scope. Is it then legitimate for a U.S. based entity to attack hosts situated in a different country on grounds that they are facilitating infringement on U.S. soil?

BEYOND COPYRIGHT ENFORCEMENT

Peer-to-peer networks are not limited to exchanging copyrighted materials. As we discussed in the introduction to this chapter, they have many beneficial uses, including promoting free speech, or reducing infrastructure costs for software providers.

But, beyond copyright infringement issues, peer-to-peer systems also create puzzling challenges for the community. With so many hosts connected to each other, information leaks can propagate at extremely high speed. Indeed, in Japan for instance, nuclear plant floor maps or databases of crime victims ended up on the Winny peer-to-peer network,[xxiv] after a virus targeting the Winny software was found to upload the victims' entire hard drive on the network.[xxv] Similarly, Good & Krekelberg (2003) showed that inadvertent information sharing on peer-to-peer networks was very common, with some peer-to-peer users' credit card statements or mortgage documents being shared on the KaZaA network.

Unfortunately, once data start being replicated on a peer-to-peer network, it is all but impossible to "take it back." Thus, the only possible defense is to limit its diffusion as much as possible. On the bright side, most users of peer-to-peer networks have no interest in accessing such confidential data, so that propagation on the peer-to-peer network is relatively limited even in the

absence of counter-measures.[xxvi] As a result, techniques like the poisoning techniques discussed above may be particularly effective to further impede the diffusion of confidential data.

CONCLUSIONS

In this chapter, we have discussed how economic incentives led to the establishment of a "tussle" between the different actors involved in peer-to-peer networks. We have shown that, to try to win the tussle, the content industry has used a pretty extensive legal arsenal, with relatively limited success thus far. We have also shown some of the technological countermeasures in place to attempt to limit dissemination of copyrighted materials in peer-to-peer networks. We have shown, however, that some of these technologies appear themselves to break the law. In addition, the proposed countermeasures may work to some extent, but are likely most applicable when targeted at specific items with small propagation (e.g., private information leakages) rather than widely distributed movies or songs.

While we have tried to paint a picture of the tussle as neutral as possible, our conclusion is that the tussle probably will not come to a resolution without the intervention of legislators. None of the legal or technological remedies against copyright infringement appear to be working satisfactorily. Peer-to-peer traffic is more predominant than ever on the Internet, driven by the widespread availability of hardware and network connectivity. Copyright reform seems necessary and could lead to new business models, in the same way the "legalization" of the Betamax technology paved the way for the home video industry.

On June 7, 2009, the Pirate Party of Sweden gained 7% of the votes in the European Parliamentary elections and will have one representative seated in Brussels. More significantly, 19% of the voters younger than 30 years old cast a vote in favor of the Pirate Party.[xxvii] While the name of the party may elicit a smile, these new developments show that the issue of copyright reform has now been put in the spotlight more than ever. Likewise, the surprising initial rejection by the French National Assembly of the Hadopi law, which advocates "gradual response" against copyright infringers culminating in a loss of right to use the Internet, grabbed the headlines. While the law was eventually adopted, these events again demonstrated the importance to society of the debate about copyright reform.

Besides copyright reform, the emergence of peer-to-peer networks presents us with considerable ethical challenges. Defining acceptable use, assigning liabilities, and establishing acceptable punishments for misuse all pose legal and ethical questions that we need to answer. Given the large amount of peer-to-peer traffic that is infringing copyright, there seems to be a large disconnect between acceptable social norms and legal behavior. The content industry is trying to reduce this disconnect through aggressive advertising campaigns, including messages on DVD informing customers that illegal downloads are a crime. But the fact that peer-to-peer traffic persists tends to show this education strategy is not successful, as most users remain unconvinced that copyright infringement is a serious offense (Easley, 2005). Even if we agree that copyright infringement is a criminal offense, we have shown that peer-to-peer networks involve a large number of actors, besides end users, that stand to indirectly profit from infringement. Who should then be held liable? And last, in terms of punishment, both the legal and technological remedies invite reflection. As far as legal sanctions are concerned, appropriately quantifying the amount of

monetary losses suffered and fairly defining the punitive damages are both challenging. Is it acceptable to use large monetary punishments for the sake of deterrence? From a technological standpoint, is it ever acceptable to launch denial-of-service attacks against some users? Is suspected copyright infringement enough of a justification? When considering leaks of private or sensitive data in peer-to-peer networks, should we consider the level of confidentiality of the information in our decision to deploy countermeasures? Can we authorize such attacks while at the same time guaranteeing freedom of speech may not be threatened?

We do not claim to have answers to these questions, at this stage. While we can point out the conditions in which certain technical countermeasures may be more effective than others (Christin et al., 2005), we firmly believe that the issue of information flow control that has been posed by the emergence of peer-to-peer networks, will require considerable thought, and will be one of the more important technological challenges of the 21st century.

ACKNOWLEDGMENTS

This chapter greatly benefited from the feedback and suggestions of several anonymous reviewers, and from discussions with John Chuang, Andreas Weigend, Alexandre Mateus, Joe Hall, Jens Grossklags, Keiji Takeda, and the students at Carnegie Mellon CyLab Japan, among many others. Part of this presentation is derived from notes the author contributed to the blog of Pam Samuleson's peer-to-peer seminar at UC Berkeley, which was held in the Spring of 2005.

REFERENCES

Basher, N., Mahanti, A., Mahanti, A., Williamson C., & Arlitt, M. (2008). A comparative analysis of web and peer-to-peer traffic. In *Proceedings of the 2008 WWW Conference*. Beijing, China.

Boyle, J. (2008). *The public domain: Enclosing the commons of the mind*. New Haven, CT: Yale University Press.

Bhattacharjee, S., Lertwachara, K, Gopal, R., & Marsden, J. (2006) Impact of legal threats on online music sharing activity: An analysis of music industry legal actions. *Journal of Law and Economics 49*; (1), 91-114.

Christin, N. & Chuang, J. (2005). A cost-based analysis of overlay routing geometries. In *Proceedings of IEEE INFOCOM'05*, Vol. 4., pages 2566-2577. Miami, FL.

Christin, N., Weigend A., & Chuang, J. (2005) Content availability, pollution, and poisoning in peer-to-peer file sharing networks.. In *Proceedings of the Sixth ACM Conference on Electronic Commerce* (pp. 68-77). Vancouver, BC, Canada.

Clark, D., Wroclawski, J., Sollins, K., & Braden R. (2005). Tussle in cyberspace: Defining tomorrow's internet. *IEEE/ACM Transactions on Networking*, 13; (3) 462-475.

Copyright Act, 17 U.S.C. § 101-122, Pub. L. 94-553 (1976).

Computer Fraud and Abuse Act, 18 U.S.C. §1030 (1984).

Eckersley, P., von Lohmann F., & Schoen S. (2007). Packet forgery by ISPs: A report on the comcast affair. Electronic Frontier Foundation online report. Retrieved from http://www.eff.org/wp/packet-forgery-isps-report-comcast-affair.

Dana, J. & Spier K. (2001). Revenue sharing and vertical control in the video rental industry. *The Journal of Industrial Economics. 49* (3) 223-245.

Digital Millennium Copyright Act, 17 U.S.C. §§ 512, 1201–1205, 1301–1332; 28 U.S.C. § 4001, Pub. L. 105-304 (1998).

Easley, R. (2005) Ethical issues in the music industry response to innovation and piracy. *Journal of Business Ethics*. 62,163–168.

Feldman, M., Lai, K., Stoica, I., & Chuang, J. (2004). Robust incentive techniques for peer-to-peer networks. In Proceedings of the *Fifth ACM Conference on Electronic Commerce (EC'04)* (pp. 102-111). New York, NY.

Gassner, U. (2005). Copyright and digital media in a post-Napster world: International Supplement. Retrieved from SSRN: http://ssrn.com/abstract=655391 or DOI: 10.2139/ssrn.655391

Golle, P., Leyton-Brown, K., Mironov, I., & Lillibridge, M. (2001). Incentives for sharing in peer-to-peer networks. In *Proceedings of the Second International Workshop on Electronic Commerce*. L. Fiege, G. Mühl, and U. G. Wilhelm (Eds.), Lecture Notes in Computer Science, Vol. 2232. (pp. 75-87). Springer-Verlag, London.

Good, N., Dhamija, R., Grossklags, J., Thaw, D., Aronowitz, S., Mulligan, D., & Konstan. J. (2005). Stopping spyware at the gate: A user study of privacy, notice and spyware. In *Proceedings of the First Symposium on Usable Privacy and Security (SOUPS)*. Pittsburgh, PA, U.S.A.

Good, N. & Krekelberg, A. (2003). Usability and privacy: A study of Kazaa P2P file-sharing. In *Proceedings of the ACM Symposium on Human Computer Interaction (CHI)*. Fort Lauderdale, FL, U.S.A.

Hale, J. & Manes, G. (2004) Method to inhibit the identification and retrieval of proprietary media via automated search engines. U.S. Patent number: 6732180. Filing date: Aug 8, 2000. Issue date: May 4, 2004

Karagiannis, T., Broido, A., Brownlee, N., Claffy, K.C., & Faloutsos, M. (2004). Is P2P dying or just hiding? In *Proceedings of IEEE Globecom 2004*. Dallas, TX, U.S.A.

Karagiannis, T., Rodriguez P., & Papagiannaki, K. (2005). Should internet service providers fear peer-assisted content distributions? In *Proceedings of the 2005 ACM/USENIX Internet Measurement Conference.* Berkeley, CA, U.S.A.

Katabi, D., Handley, M., & Rohrs, C. (2002). Congestion control for high bandwidth-delay product networks. In *Proceedings of the 2002 ACM SIGCOMM Conference.* Pittsburgh, PA, U.S.A.

Lemon, J. (2002) Resisting SYN flood DoS attacks with a SYN cache. In *Proceedings of USENIX BSDCON 2002*. San Francisco, CA.

Levin, S. & Disher, J. (2004). System and methods for communicating over the Internet with geographically distributed devices of a decentralized network using transparent asymmetric return paths. U.S. Patent Application number 10/869,208, publication number 2005/0089014.

Liang, J., Kumar, R., Xi, Y., & Ross. K. (2005) Pollution in P2P filesharing systems. In *Proceedings of IEEE INFOCOM'05*, Miami, FL.

Liebowitz, S. (2006). File-sharing: Creative destruction or just plain destruction? Center for the Analysis of Property Rights Working Paper No. 04-03. Retrieved from SSRN: http://ssrn.com/abstract=646943 or DOI: 10.2139/ssrn.646943

Limitations on liability relating to material online. 17 U.S.C. § 512 (1998).

Maymounkov, P. & Mazières, D. (2002) Kademlia: A peer-to-peer information system based on the XOR metric. In *Proceedings of the International Workshop on Peer-to-Peer Systems (IPTPS)*. Cambridge, MA.

Moore, J., Bland, W., Francis, S., King, N., Patterson, J., Srinivasan, U. & Widden P. (2004) Interdiction of unauthorized copying in a decentralized network. U.S. Patent Application number 10/803,784, publication number 2005/0091167.

Oberholzer-Gee, F. & Strumpf, K. (2007). The effect of file sharing on record sales: An empirical analysis. *Journal of Political Economy. 115*(1) 1-42.

Shang, R.-A., Chen, Y.-C., & Chen, P.-C. (2008). Ethical decisions about sharing music files in the P2P environment. *Journal of Business Ethics*, 80(2): 349–365.

Taipale, K. (2003) Secondary liability on the internet: Towards a performative standard for constitutive responsibility. Center for Advanced Studies Working Paper No. 04-2003. Available at SSRN: http://ssrn.com/abstract=712101

Walsh, K. & Gün Sirer, E. (2006). Experience with an object reputation system for peer-to-peer filesharing. In *Proceedings of the Symposium on Networked System Design and Implementation*. San Jose, CA, U.S.A.

---

[i] See for instance http://www.slackware.com/torrents/, last accessed June 9, 2009.
[ii] See http://www.blizzard.com/us/legal-bfd.html, last accessed June 9, 2009.
[iii] From http://www.dynamicsun.com/, last accessed June 2, 2009.
[iv] See http://www.enterprisenews.com/archive/x1225086200/MASS-MARKET-Nostalgia-alone-won-t-keep-movie-rental-stores-afloat, last accessed June 8, 2009.
[v] See "Disc Piracy: it costs more than you think," by Dan Daley, http://www.linkdata.dk/linkpress/TDB-pir.htm, last accessed June 8, 2009.
[vi] An exception worth noting is that in Asia, sales of illegitimate copies of movies on VCD and CD-ROMs have soared using traditional distribution channels (e.g., brick and mortar video stores) and have been made solely possible by the shift from analog to digital media.
[vii] Formally, convenience can be formulated in monetary terms, e.g., by estimating how users value their time. Formal economic modeling of actors' incentives, however, is beyond the scope of the discussion in this chapter.
[viii] See http://www.nytimes.com/2007/07/02/business/media/02universal.html, last accessed June 8, 2009.
[ix] See for instance http://www.plus.net/support/broadband/products/archive/bbyw/features.shtml, last accessed June 9, 2009.
[x] Even if complete file transfers are broken down in smaller pieces, they remain generally much longer than typical web connections.

[xi] It is worth noting that the risk of cannibalization of one's own business is not limited to peer-to-peer vs. digital TV, for these operators. Internet telephony, which is another "killer application" driving the demand for more bandwidth, essentially competes with traditional telephone access which are offered by the same service providers. A desire to avoid self-cannibalization has lead to the emergence of "triple play" offerings – discounts aimed at securing users' patronage in all services simultaneously.
[xii] See http://arstechnica.com/tech-policy/news/2009/05/harvard-prof-tells-judge-that-p2p-filesharing-is-fair-use.ars, last accessed June 8, 2009
[xiii] See http://tech.yahoo.com/blogs/null/146827, last accessed August 7, 2009.
[xiv] See A&M RECORDS, Inc. v. NAPSTER, INC., 239 F.3d 1004 (9th Cir. 2001).
[xv] See MGM Studios, Inc. v. Grokster, Ltd. 545 U.S. 913.
[xvi] See Universal Music Australia Pty Ltd v. Sharman License Holdings Ltd. FCA 1242 (2005).
[xvii] See MGM Studios, Inc. v. Grokster, Ltd. 545 U.S. 913.
[xviii] See http://www.betanews.com/article/RIAA-Piracy-Damages-Questioned-in-Ruling/1163182272, last accessed June 9, 2009.
[xix] See http://arstechnica.com/tech-policy/news/2009/06/jammie-thomas-retrial-verdict.ars, last accessed July 28, 2009.
[xx] See http://www.sfgate.com/cgi-bin/article.cgi?file=/chronicle/archive/2003/09/25/BUGJC1TO2D1.DTL&type=business, last accessed June 9, 2009.
[xxi] See http://thepiratebay.org/, last accessed June 9, 2009.
[xxii] See http://www.freedom-to-tinker.com/blog/felten/why-unbreakable-codes-dont-make-unbreakable-drm, last accessed June 8, 2009.
[xxiii] Some of them may, however, be somewhat contained within linguistic and cultural borders, due to the type of contents they are hosting. The next section discusses one such example with the Winny network.
[xxiv] Winny is one of the most popular peer-to-peer networks in Japan.
[xxv] See Nihon Asahi Shimbun, "Leaks spur splurge for new SDF computers," March 8, 2006.
[xxvi] "Security and privacy in file-sharing networks", presentation given by Nicolas Christin at Tokyo University, June 18, 2007.
[xxvii] See http://www.thelocal.se/19928/20090607/, last accessed June 9, 2009.