Forcing the Net Through a Sieve: Why Copyright Filtering is Not a Viable Solution for U.S. ISPs

Mehan Jayasuriya, Jef Pearlman, Robb Topolski, Michael Weinberg, Sherwin Siy

www.PublicKnowledge.org

Acknowledgements

The authors of this paper would like to thank the following advocates for their assistance and suggestions during the editing process: Public Knowledge president Gigi Sohn, Public Knowledge staff attorney Rashmi Rangnath, Public Knowledge legal consultant Adam Thomas, Public Knowledge intern Katy Tasker and Electronic Frontier Foundation staff technologist Peter Eckersley.

Table of Contents

Executive Summary 1

1. Introduction 2

2. Technological Analysis: The Anatomy of a Copyright Filter 7

I. Determining Filter Policy 8
- *A. Strict Policy Definition* *8*
- *B. Permissive Policy Definition* *9*

II. Identifying Types of Content 10
- *A. Traffic Inspection* *12*
  - i. Identifying Content Through Associated Transactions or Through the Identification of Connection Patterns 12
  - ii. Identification Through Analysis of Traffic Patterns 13

III. Content Analysis 13
- *A. Identifying Content Through the Use of Metadata* *14*
- *B. Identification of Content Through Digital Watermarks* *15*
- *C. Identification of Content Through Acoustic or Visual Characteristics ("Fingerprinting")* *16*
- *D. Instrumentation of the End-User's Device* *17*

IV. Policy Enforcement 18

V. Examples of Copyright Filters 20
- *A. Audible Magic Copysense* *20*
- *B. Red Lambda Integrity* *22*
- *C. Vobile VideoDNA* *23*

3. Limitations and Consequences of Copyright Filtering 25

I. Technological Limitations of Copyright Filters 25
- *A. Architecture is a Poor Indicator* *25*
- *B. Protocol is a Poor Indicator* *25*
- *C. Media Type is a Poor Indicator* *26*

II. Filter Processing Adds Latency 26
- *A. Filtering Technologies Can Expose Networks to Security Risks* *27*
- *B. A Copyright Filter Could be Intentionally Misused for Censorship Purposes* *27*

III. The Implementation of Copyright Filters Will Result in a Technological Arms Race 29
- *A. Encryption* *31*
- *B. Protocol Obfuscation* *33*

IV. The Ramifications of the Arms Race 35

4. Economic Analysis 38

I. Who Will Pay for Copyright Filtering? 38

II. Copyright Filtering Holds the Potential to Disrupt the Internet Economy, Our Most Promising Engine for Economic Growth 42

5. Legal Analysis 47

I. Mandatory Copyright Filters Would Impose an Unconstitutional Burden on Free Expression, Contrary to the Principles of Copyright Law 47

II. Copyright Filtering Could Undermine the Safe Harbor Provisions Granted to ISPs Under the Digital Millennium Copyright Act (DMCA) 49

III. ISPs That Engage in Packet Inspection Risk Violating the Electronic Communications Privacy Act (ECPA) 52

6. Conclusion 55

Executive Summary

Copyright filtering, the latest proposed "magic bullet" solution from the major music and movie studios and industry trade groups, poses a number of dangers to Internet users, legitimate businesses and U.S. federal government initiatives to increase the speed, affordability and utilization of broadband Internet services. The following whitepaper presents a number of reasons why the use of copyright filters should not be allowed, encouraged or mandated on U.S. Internet Service Provider (ISP) networks. Among them:

1. **Copyright filters are both underinclusive and overinclusive.** A copyright filter will fail to identify all unlawful or unwanted content while harming lawful uses of content.
2. **Copyright filter processing will add latency.** Copyright filters will slow ISP networks, discouraging use, innovation and investment and harming users, businesses and technology policy initiatives.
3. **The implementation of copyright filters will result in a technological arms race.** Users will act to circumvent the filters and the architects of the filters will find themselves caught in a costly, unwinnable arms race.
4. **Copyright filters do not make economic sense.** The monetary costs associated with copyright filtering far outweigh any perceived benefits.
5. **Copyright filters will discourage investment in the Internet economy.** Copyright filters will disrupt the Internet ecosystem, severely undermining our most promising engine for economic growth.
6. **Copyright filters will harm free speech.** Due to technological limitations, copyright filters will harm lawful, protected forms of speech such as parody and satire.
7. **Copyright filters could undermine the safe harbor provisions that shield ISPs from liability.** Under the Digital Millennium Copyright Act (DMCA), ISPs are shielded from liability for their users' actions. Copyright filters could undermine these safe harbors, which have allowed the Internet to become the most important communications medium of the modern era.
8. **Copyright filtering could violate the Electronic Communications and Privacy Act.** Copyright filtering could constitute unlawful interception under the Electronic Communications and Privacy Act (ECPA).

1. Introduction

Ever since the advent of mainstream file-sharing networks in the late 1990s, the major music and movie studios and industry trade groups (henceforth referred to as "the content industry") have been searching for a "magic bullet" solution to the problem of online copyright infringement—one that would simultaneously eradicate the problem of online file sharing while breathing new life into pre-digital business models. Obviously, this magic bullet has eluded the industry thus far, though not due to a lack of effort. In the years since the emergence of Napster, the content industry has taken legal action against service providers and end users, shut down centralized file-sharing networks and flooded the web with bogus copies of commonly traded files. Yet, as even a casual observer can attest, illicit file sharing persists online.

The industry's latest silver bullet comes not in the form of a legal campaign or disruption strategy but rather in the form of a technology that falsely promises to automatically and effectively eradicate copyright infringement online. Copyright filtering, as it is called, is a method whereby network appliances use a technology known as Deep Packet Inspection (DPI) to inspect the data that travels over an Internet Service Provider's (ISP's) network, identifying content as it passes through the filter and then dealing with that content accordingly. Unsurprisingly, the content industry has become an outspoken advocate for the use of this technology, pushing governments the world over to pass legislation requiring its use on ISP access networks.

Here in the United States, the content industry was almost successful in forcing language into the American Recovery and Reinvestment Act of 2009[1] that would have allowed ISPs to engage in copyright filtering under the auspices of "reasonable network management," an act that would have had far-reaching consequences for citizens, businesses and the entire Internet ecosystem.[2] Meanwhile, based on information publicly available,[3] a secretive, multilateral, international trade agreement known as the Anti-Counterfeiting Trade Agreement (ACTA)[4] could also pave the way for copyright filtering in the United States, the European Union, Australia, Japan, Canada and a number of other nations.

Before we rush to implement a technology—much less require its use—we must objectively examine that technology and question not only its efficacy but also its costs, consequences and drawbacks. It is in this spirit that Public Knowledge, a non-profit public interest advocacy group, has embarked upon an analysis of copyright filtering technologies. The results of that analysis are presented in this whitepaper.

In the following paper, we will take a close look at the technology behind copyright filtering, analyze the legal ramifications of filtering and discuss the policy implications of condoning or mandating the use of copyright filters. We will consider the impact that filtering is likely to have on user privacy, examine the costs associated with filtering and contemplate the potential impact copyright filtering will have on network security for both users and service providers.

---

[1] *See* the Recovery.gov website (http://www.recovery.gov/?q=content/act).
[2] "Senator Feinstein Trying to Sneak ISP Copyright Filtering Into Broadband Stimulus Bill," *TechDirt*, Febuary 10, 2009 (http://www.techdirt.com/articles/20090210/1050313726.shtml).
[3] "Secret ACTA Treaty May Include ISP Filtering," *Ars Technica*, June 4, 2008 (http://arstechnica.com/tech-policy/news/2008/06/secret-acta-treaty-may-include-filtering-provisions.ars).
[4] For more information on ACTA, see the Electronic Frontier Foundation website (http://www.eff.org/issues/acta).

Regardless of how they are implemented, copyright filters will alter the fundamental behavior of the Internet and in so doing, will likely disrupt the Internet ecosystem in ways that we cannot predict. The Internet was designed to be an open system from end-to-end, which is to say, a system that moves content between hosts and clients as quickly as possible on a first-come-first-served basis—regardless of the nature of that content. Copyright filters will inject delay into this system, make automated judgments regarding the legality of content and will then degrade or discard that traffic accordingly. The Internet was not designed to support this type of activity and for this reason, the implementation of copyright filters will assuredly result in a variety of technical problems for all parties involved in the Internet ecosystem. Much like traffic lights on our interstate highway system, copyright filters on our open, high-speed networks will be a poor fit.

Technical issues aside, there are other pressing questions that must be answered before we can determine whether copyright filters should be installed on the networks of Internet service providers. Chief among them is the question of efficacy: will copyright filters solve the problem that they purport to solve? Based on our technical analysis, the answer appears to be no. By virtue of their design, Internet filters are doomed to be both underinclusive and overinclusive—they will fail to identify all illegal uses of content while simultaneously blocking legal content. The filters will be underinclusive because their technology is not advanced enough—and will likely never be advanced enough—to identify every instance of prohibited content on the network. Filters will also be overinclusive; as a filter will never be able to distinguish between fair, legal uses of content and illegal uses of content with 100 percent accuracy. Given that even legal scholars and courts are often unable to reach a consensus on questions of fair use, this is not surprising—the question of whether or not a piece of content

constitutes fair use is often a difficult one for even human beings to answer. Furthermore, as history attests, users will work to actively circumvent the filter, thereby luring the architects of the filter into a fruitless technological arms race.

Ultimately, we will demonstrate that copyright filtering does not constitute "reasonable network management," as some proponents of the technology would have policymakers believe; rather, copyright filtering is *content* management. Instead of making determinations based on how data moves over the network, copyright filters attempt to ascertain *what* that data constitutes, in order to block, degrade or delay certain types of content. Copyright filtering is not necessary for a network to operate properly or efficiently and therefore, should not be considered a form of network management at all.

The content industry would like to convince policymakers and the general public that copyright filtering is the most effective means by which to combat online copyright infringement and protect America's creative economy. This could not be further from the truth. In practice, copyright filtering is likely to harm innovators, end users, online service providers and Internet service providers alike. What's more, it will compromise the privacy of all American Internet users for the perceived benefit of one industry. As such, copyright filtering will discourage investment in the Internet economy—our most promising engine for economic growth—and will harm American competitiveness in the global market.

Finally, copyright filtering holds the potential to undermine the goals of the National Broadband Plan (NBP), a Federal initiative to increase the speed, adoption and affordability of broadband Internet services nationwide. As we will see, copyright filtering, if implemented on U.S. networks, will chip away at each of the NBP's stated goals. The network appliances that act

as filters will slow traffic on the networks on which they are installed, an unavoidable consequence of their traffic analysis. The costs associated with filtering—most notably, the purchasing and maintenance of filtering hardware and software—are likely to be passed on to consumers, decreasing the affordability of broadband services. And in compromising the privacy of Internet users, copyright filters will discourage Internet use, even as the Federal government works to promote the educational, economic and civic benefits of broadband access.

If we are serious about combating copyright infringement online, we should invest our time and resources in developing methods that will effectively discourage illegal activity, without harming fair users, innovators and the Internet economy. If we rush to require the use of a technology that we have not fully considered, we will find ourselves trapped by a burdensome and even dangerous mandate, one that will have far-reaching consequences that will affect the flow of information, knowledge, political discourse and capital. To date, the content industry has presented the debate surrounding copyright filtering as a choice between a specific technology and rampant online "piracy". In the paper that follows, we will demonstrate that this is a false choice and that copyright filtering should not be considered as a viable solution for U.S. ISPs.

2. Technological Analysis: The Anatomy of a Copyright Filter

The term "filter" refers to a technology that can be employed by an ISP or end-user to automatically detect a specific type of traffic or content and then take action based on the nature of the data detected. Historically, filters have been used at the network level on private networks to block content that is presumed to be illegal (child pornography, unlawfully traded movies and music, etc.) or unwanted (spam, personal communications, social networking content, etc.). In such cases, the filtering hardware or software is installed at the point where the private network connects to the public Internet. If a filter were installed on an ISP network, however, it would likely be installed at the transit provider level—either at points along the Internet backbone or at the point where the private ISP connects to the backbone.

The architecture of the Internet is such that in order to provide Internet service, a network operator needs only to concern itself with the Internet Protocol (IP) headers of data packets that traverse its network—that is to say, the outer layer of the most basic unit of Internet traffic. This is because all of the information required by a provider to do its job—receiving data packets and then forwarding them to their next stop en route to their destinations—is contained in the packet's IP header. ISPs are occasionally required to inspect the traffic of individual users, pursuant to regulatory requirements and the needs of law enforcement but otherwise generally do not analyze any part of the packet other than the header.

However, if an ISP chose to implement a copyright filter, that ISP would be electing to scrutinize all traffic from all subscribers, in real time. The provider would then act upon the information gleaned by the filter, in accordance with its policy aims. Needless to say, this system

represents a significant departure from the traditional role of the ISP, which is to forward traffic as quickly and efficiently as possible.

When designing an Internet filter of any sort, one must first determine the policy that will define the filter's function. In this section, we will consider two policy definitions of theoretical copyright filters: a "strict" policy definition and a "permissive" policy definition. We will then address the different mechanisms whereby a filter might identify different types of content. Next, we will explore the filter's policy enforcement mechanism, whereby it takes action once a piece of content has been isolated and identified. Finally, we will look at examples of technologies on the market that are currently marketed as real-time filtering solutions.

I. DETERMINING FILTER POLICY

A. Strict Policy Definition

As all works created in the United States are essentially "born" copyrighted[5], a strict copyright filter would have to operate under the assumption that all Internet traffic contains copyrighted work. As such, a filter based on a strict policy definition would analyze each packet, in an attempt to identify data that pertains to a copyrighted work. If a packet were determined to be a piece of such a work, the filter would then attempt to verify the identity of the work. Once the work's identity is determined, the filter would verify that both the server and client are properly licensed to distribute and receive the content in question and if not, would then presumably take action to delay, degrade or discard that packet.

[5] A work is "created" when it is fixed in a copy or phonorecord for the first time (17 U.S.C. § 101); "Copyright protection subsists, in accordance with this title, in original works of authorship fixed in any tangible medium of expression, now known or later developed, from which they can be perceived, reproduced, or otherwise communicated, either directly or with the aid of a machine or device." (17 U.S.C. § 102(a)).

This, of course, will be an extremely complicated process. Questions of content identification aside, the filter must consider a number of factors when making judgment calls regarding licenses. This is because certain content licenses (mechanical licenses, for example) might only authorize a certain number of reproductions of a work or might only allow distribution to certain users in certain geographies or within certain time frames. The complex nature of licensing in the United States adds an added layer of complexity to the copyright filter.

Here is an abstracted version of how a strict filter might operate:

```
SomeServer is sending SomeContent to SomeUser

{

if SomeContent is authorized on ContentAllowedList;

And SomeServer is authorized on the AuthorizedToSend list for
SomeContent to SomeUser;

Then, allow the transfer;

};

Otherwise deny the transfer;
```

B. Permissive Policy Definition

Unlike a strict copyright filter, a permissive copyright filter would operate under the assumption that any traffic passing through it does not represent an unauthorized transfer of copyrighted work. In this sense, a permissive filter could be said to take an "innocent until proven guilty" approach, rather than the "guilty until proven innocent" approach taken by the strict filter. Like the strict filter, the permissive filter would still have to identify works, the identities of the parties involved in the transfer of the works and any applicable licenses that authorize that transaction. Unlike the strict filter, however, the permissive filter would cross-check the content against a database of "blacklisted" content—that is to say, content that is not

allowed to be transferred on the web, in accordance with the wishes of the copyright owner. The permissive filter might also record the IP address of the end user and/or cross-check that IP address against a database of known offenders.

Here is an abstracted example of a permissive filter:

```
SomeServer is sending SomeContent to SomeUser

{

if SomeContent is not listed on RegisteredContentList;

Or SomeServer is authorized on the AuthorizedToSend list for
SomeContent to SomeUser;

Then, allow the transfer;

};

Otherwise deny the transfer;
```

Most copyright filters will fall into either the strict or permissive category, with regard to policy definition. It should be noted that neither type of filter is compatible with existing copyright law, as neither filter allows for the unauthorized uses permitted under copyright law. While the permissive filter appears to be a more practical implementation (as it would likely halt/slow the flow of traffic far less than the strict filter), it still fails to consider those uses that are not subject to the exclusive rights of the content owner and hence permitted under copyright law. These limitations are discussed in depth in the legal analysis section of this paper.

II. IDENTIFYING TYPES OF CONTENT

As part of the policy definition process, the designers of a copyright filter would have to craft a policy that would allow the filter to identify different types of content. In order to determine the type of content being transmitted, the filter, on behalf of the network operator, would analyze the bits that travel over the network and then make a determination as to what

type of content is in transit. Is it audio? Is it video? Is it both audio and video? Is it a document? If so, is it a book? Is it a magazine article? Is it sheet music? Is it a photograph?

Once the filter has made a first-level content determination (type of content), it will then attempt to identify the exact identity of the work being transferred. Needless to say, this part of the process will significantly increase the complexity of the determination process. A filter cannot simply block music as a broad category; it must only block specific pieces of music that match a certain profile. In generic terms, a filter might seek to block *SomeRestrictedPerformance* of *SomeRestrictedTitle* by *SomeSignedBand* which was released *SomeYear* by *SomeLabel*. It might also be required to block *SomePublicDomainMelody* as sung by *SomeSinger* as arranged by *SomeRestrictedArranger*, for example.

In order to have enough data to make these determinations, the filter must either permit the user to download a significant portion of the file in question (perhaps the entire file), or it must itself download a significant portion of the file, before passing that data on to the user, if that data is determined to not run afoul of the filter's policy. Given that the former method undermines the effectiveness of the filter, we must assume that the latter method would be used and that Internet traffic would be significantly slowed, as each file would effectively have to be downloaded twice—once by the filter and a second time by the user. The only obvious solution to this problem is to require a license that authorizes the transfer of a work in advance of the file transfer. Such a system would, of course, require that all legitimate transfers on the Internet be accompanied by a license—a solution that seems unrealistic at best.

Regardless of how transfers are regulated, the most basic task that the filter performs will be identifying a piece of content. Technologically speaking, there are a number of different ways

that this might be accomplished, many of which are discussed below. These methods have been divided into three categories: traffic inspection, content analysis and instrumentation of the end-user's device.

A. Traffic Inspection

The term traffic inspection refers to an identification mechanism that does not attempt to identify specific pieces of content online. Rather, a traffic inspection scheme would simply analyze the nature of traffic that travels over the network and would then make assumptions about the content carried by that traffic. This was the method used by Comcast, whereby it clandestinely degraded all traffic that used the BitTorrent protocol, under the flawed assumption that all of the traffic that uses that protocol is illicit in nature. Comcast has since been reprimanded by the Federal Communications Commission (FCC) for its actions, as it was found that Comcast's content management technique violated the four principles outlined in the FCC's Internet Policy Statement.[6] Comcast has since ceased degrading traffic related to BitTorrent.

i. Identifying Content Through Associated Transactions or Through the Identification of Connection Patterns

One traffic analysis method that might be used for content identification is the analysis of either packet exchange or geographic patterns. The phrase "packet exchange" refers to the conversational patter of an Internet connection—the back-and-forth exchange of data between the host and the client. Connection pattern analysis methods look at who sits on either end of a packet exchange transmission and how many connections those hosts and clients are opening and uses this data to speculate as to what the nature of the transmission might be.

[6] See "Formal Complaint of Free Press and Public Knowledge Against Comcast Corporation For Secretly Degrading Peer-to-Peer Applications" (http://www.publicknowledge.org/pdf/fp_pk_comcast_complaint.pdf).

While there are existing technologies that can analyze Internet traffic using these methods, some of which are marketed for copyright protection, the fact of the matter is that these technologies do not identify instances of copyrighted content. Rather, they block certain types of traffic—file transfers, for example—regardless of the nature of the content in transit. In other words, they interrupt all traffic, lawful and unlawful.

ii. Identification Through Analysis of Traffic Patterns

Similar to the technique of identifying works through packet exchange patterns, traffic pattern analysis identifies traffic patterns popular with—but not exclusive to—users who illegally share copyrighted works online. Like with packet exchange analysis, this method does not attempt to directly distinguish between legal and illegal content and instead, blocks all traffic that matches a certain pattern. However, unlike connection pattern analysis, which looks only at who is sending and receiving data, traffic exchange analysis also examines the size and timing of the packets exchanged. Commercial network products such as ArborNetworks eSeries[7] and RedLambda Integrity[8] use this method.[9]

III. CONTENT ANALYSIS

Unlike traffic inspection techniques, content analysis methods look beyond the headers of packets, in an attempt to analyze the content of those packets. While content analysis techniques avoid some of the pitfalls of traffic inspection, they exhibit a different set of shortcomings and, as such, are similarly flawed.

[7] See Arbor Networks' website (http://www.arbornetworks.com/index.php?option=com_content&task=view&id=1466&Itemid=693).
[8] See Red Lambda website (http://redlambda.com/integrity_overview.php).
[9] Examples of filtering technologies will be discussed in-depth later in this section.

In order to identify a work using a content analysis technique, the filter would compare the data in hand to a registry containing the identifying characteristics of all known works or, failing that, all works of concern. Considering that it will be impossible to create a database containing information on all works, we must assume that any such database will contain a subset of those works, all of which will likely be works protected under copyright.

A. Identifying Content Through the Use of Metadata

Most digital works contain what is called metadata: data about data that is used for purposes of categorization. One of the most recognizable forms of metadata are the ID3 tags that come attached to many digital music files. In this case, the metadata provides information including the performer, songwriter, title, album, year of recording and file quality of a file to the software and hardware that decodes that file. This is how devices like iPods and software like iTunes are able to display the title and artist when playing a song. Most digital media files including movies, photographs and digital books, contain similar metadata.

Given this wealth of data, you might assume that identification would be a relatively simple task for a copyright filter. Unfortunately for the filter, this is not the case. Metadata can be easily edited by users and as such, often contains incomplete or imperfect information. If you have ever seen two songs by the same artist in your music library that were classified correctly but differently (for example, “John Lennon” vs. “Lennon, John”), you have already observed this phenomenon. If we allow for the possibility of mislabeling (“John Lennon” vs. “The Beatles”), this problem is compounded.

Even if the metadata in question were complete and reliable, however, a filter would still encounter a number of problems when attempting to identify a work based on that data. Much of

this has to do with the complex nature of the type of data that is found in metadata and the possibility for overlap. Take for example, a song entitled "Happy Birthday". Given only this title, a filter would be unable to determine whether the work in question is the copyrighted song "Happy Birthday," a different song of the same name or something else entirely that has been mislabeled. If the filter attempts to identify the song based on its listed performer, it will run into a similar set of problems. Does the artist listed perform the song in question or is it an amateur cover of a song by that performer? Or, is the file in question simply a song performed in the style of that performer?

In attempting to positively identify a piece of content based on its metadata, a copyright filter will likely have no choice but to view the metadata holistically. Given that none of the data will be certifiably reliable, however, the filter will never be able to identify a piece of content based solely on its metadata with any degree of certainty. You might say that the filter would have trouble judging a book by its cover.

B. Identification of Content Through Digital Watermarks

The term "digital watermark" refers to a technique whereby content producers embed an invisible digital signature into a file intended for physical or digital distribution. Once the watermark has been embedded, the watermarked content can be easily identified, tracked, managed and secured by technology that can read the watermark. Such technology is currently being used by a number of major movie studios, record labels, television broadcasters and enterprises with digital image assets.[10] While a copyright filter could theoretically filter only for watermarked content, this would result in an underinclusive filter. While an unauthorized

[10] See "Descriptions of Demonstrators at the ICAC 2007 Tech Exhibition," Congressional Internet Caucus Advisory Committee (http://www.netcaucus.org/events/2008/kickoff/demonstrators.shtml).

transfer of a file containing the watermark would be halted, a non-watermarked version of the same file would be allowed to pass. Considering that many of the unauthorized transfers of copyrighted files that take place online involve the transfer of content created by users through means such as camcording and the ripping of insufficiently secured CDs and DVDs, it is likely that a great deal of copyrighted content would pass through the filter undetected. Furthermore, if watermark filters were widely deployed, users would likely develop techniques for stripping the watermarks out of files.[11] This would only exacerbate the content-protection arms race (discussed at length in the "limitations and consequences" section of this paper) and inject additional costs into Internet services, increasing costs for the consumer and providing no discernable benefit.

C. Identification of Content Through Acoustic or Visual Characteristics ("Fingerprinting")

Metadata aside, every copyrighted work contains a great deal of identifiable information. A song is not known only by its title but also by a succession of notes and words. A book isn't simply known by its title but by its characters, setting and plot. The term "fingerprinting" refers to a technique whereby audio, video or literary patterns in the work are used to generate a unique fingerprint, which can then be used to identify other instances of that same work. Network products like Vobile VideoDNA[12] and Audible Magic[13] use this method to identify content.

While fingerprinting might be an effective technique for content identification, it is not a practical solution for Internet filtering. In order to positively identify a work, a fingerprint filter

[11] For an in-depth account of one research team's successful effort to crack digital watermarks, see Craver, Scott A. et al. "Reading Between the Lines: Lessons Learned from the SDMI Challenge," August 13, 2001 (http://www.usenix.org/events/sec01/craver.pdf).
[12] See Vobile website (http://www.vobileinc.com/technology.html).
[13] See Audible Magic website (http://www.audiblemagic.com/index.asp).

would have to download a considerable portion of the work in question. This would mean that either all Internet traffic would have to be considerably delayed or that, in many cases, most or all of a transfer would have to be completed before the filter would be able to make a determination regarding the nature of the content.

Even if it were practical to use digital fingerprinting in conjunction with a copyright filter, such a filter would fail to recognize legal uses of copyrighted content--for example, a film review containing brief movie clips, a transformative remix of a song or a parody or satire of a copyrighted work--and would therefore encroach on the rights of Internet users. This matter is discussed in depth in this paper's legal analysis section.

D. Instrumentation of the End-User's Device

While this paper interprets the term copyright filter to mean an in-network technology (that is to say, hardware or software that is implemented by an ISP and which sits on the ISP's network), some have suggested a different method, whereby filtering software would be installed on the machines of end-users, ostensibly as the result of a legal mandate. Certain content industry representatives, including Recording Industry Association of America (RIAA) president Cary Sherman,[14] have advocated this method and some have interpreted this advocacy as a tacit admission that in-network copyright filtering will be ineffective.[15]

While a full analysis of filtering techniques that require instrumentation of the end-user's device is outside the scope of this paper, it should be apparent that such instrumentation would raise a number of questions regarding legality, practicality and feasibility. While instrumentation

[14] "RIAA boss: Move copyright filtering from ISPs to users' PCs," *Ars Technica*, February 7, 2008 (http://arstechnica.com/old/content/2008/02/riaa-boss-spyware-could-solve-the-encryption-problem.ars).
[15] *Ibid.*

of the end-user's device would offer some advantages to those attempting to filter, such a solution would still suffer from most, if not all of the shortcomings discussed in this paper. Additionally, the installation of software on end-users' machines raises an entirely new set of security questions.[16] When the Chinese government issued a mandate that the "Green Dam" filtering software be installed on all machines sold in the Chinese market, major security vulnerabilities were discovered within days of the software's release.[17] Some of these flaws could have been exploited remotely, allowing malicious programmers to compromise the security of any machine running Green Dam over a network connection. As a result, the Chinese government rescinded the mandate, making the use of Green Dam optional.[18]

IV. POLICY ENFORCEMENT

Once a piece of content has been isolated and identified, the filter must act pursuant to the policy definition that it was created to enforce. This stage of the filter's operation is referred to as policy enforcement.

By default, the sending and receiving of all content is allowed on the Internet. Currently, the exceptions to this rule are filters installed by users, businesses and governments on their private networks, some of which allow no traffic to pass as a default and then allow data through, bit-by-bit, after it has been determined to meet a certain set of criteria.

---

[16] See also the Sony BMG CD copy protection scandal, wherein Sony encoded certain audio CDs with a copy protection mechanism that would surreptitiously install a "rootkit" on users' PCs, thereby compromising the security of those PCs while sending user information to Sony. As a result, lawsuits were filed against Sony BMG in the states of California, New York and Texas. For more information, see "Real Story of the Rogue Rootkit," *Wired*, November 11, 2005 (http://www.wired.com/politics/security/commentary/securitymatters/2005/11/69601).
[17] "Big Vulnerabilities in China's Mandatory Filtering Software," *Ars Technica*, June 1, 2009 (http://arstechnica.com/tech-policy/news/2009/06/big-vulnerabilities-in-chinas-mandatory-filtering-software.ars).
[18] "China Caves, Says Green Dam is Optional," *ChannelWeb*, June 16, 2009 (http://www.crn.com/software/217900033).

Given that the Internet's default is to allow all traffic to pass, we might reasonably assume that a copyright filter would operate in the same manner, allowing any content that it cannot identify to pass. While it is possible that a copyright filter could be set up to operate like some security filters, allowing no content to pass as a default, such a policy definition would likely result in massive traffic problems on that network, as the filter would simply reject any bit that it did not understand. Given the speed with which new technologies and content types are created and propagated on the Internet and the interconnected nature of all Internet traffic, the consequences of deploying such a filter on a public ISP could be disastrous.[19]

Regardless, once the filter has captured and identified a piece of content, it must take action. The range of actions that can be taken are defined by the following broad categories:

1. Allow: The data is allowed to pass unmolested, as it would on the Internet.

2. Flag: Before being forwarded on to its destination, the traffic is 'flagged'. In most cases, this will mean that the traffic is reclassified to have a higher/lower precedence over other traffic, in order to slow the traffic, off-load the traffic to follow a different path, or change the traffic's source or destination. Flagging might also be used to track traffic (*i.e.* to identify the sender or receiver of the data).

3. Deny: The data is discarded and is not forwarded toward its final destination.

Within these three categories, there is much room for ingenuity. If the filter alters the path of traffic, the intended destination may never receive that traffic if the path of redirection is

[19] See Public Knowledge and Free Press' complaint against Comcast wherein it is stated that not only does Comcast's blocking of BitTorrent traffic affect Comcast users, it also causes a disproportionate amount of BitTorrent traffic to be offloaded on to other ISPs, thereby disrupting the usual balance of traffic using that protocol. "Formal Complaint of Free Press and Public Knowledge Against Comcast Corporation For Secretly Degrading Peer-to-Peer Applications" (http://www.publicknowledge.org/pdf/fp_pk_comcast_complaint.pdf).

congested or has no route to the destination host. Alternatively, that traffic might simply be dropped, as if the end host had gone offline. For example, a filter designed to falsely return empty search results may accomplish this task by routing the search queries down dead-end paths. Or, the designers of a filter could create a virtual tunnel restricted to a very low speed and could then reroute all unwanted traffic to that tunnel. In this latter scenario, while end-to-end communication would still be possible, its purpose (rapid file transfer, for example) would be defeated.

In all of these categories, the ISP would retain the ability to log what is being observed and acted upon by the filter. In some cases, the action taken by the filter might vary depending on network conditions—a policy that some U.S. ISPs claim to have adopted with regard to their network management techniques.[20]

V. EXAMPLES OF COPYRIGHT FILTERS

While copyright filtering is rarely used at the ISP level, filtering technologies have long been used by the administrators of academic, corporate and government networks as a means to regulate the content, applications and information that can be accessed by end-users. In this section, we profile a few of the most commonly used filtering technologies—solutions that might be used by an ISP looking to implement a copyright filter.

A. Audible Magic Copysense

[20] Both Comcast and Cox initially claimed that their network management scheme, which throttled BitTorrent traffic, was only used during times of "peak congestion". This was quickly disproven by network researchers who discovered that Comcast and COX were actually throttling that traffic 24-hours a day, regardless of network conditions. See "Study: Comcast, Cox Slowing P2P Traffic Around the Clock," *PCWorld*, May 15, 2008 (http://www.pcworld.com/businesscenter/article/145952/study_comcast_cox_slowing_p2p_traffic_around_the_clock.html).

Audible Magic develops, markets and sells content inspection filters that use fingerprinting technology. Audible Magic solutions have most commonly been used by destination websites that host user-uploaded content and are currently in use by a number of prominent Internet companies, including YouTube and Fox Interactive subsidiary MySpace.[21] As was previously described in the content identification section, fingerprinting filters analyze changes in the pitch, rhythm and relative sound level of multimedia files and then compare the results of that analysis to a database of known or "protected" works.[22]

Copysense, an on-the-network appliance sold by Audible Magic, can inspect data in transit on a live network. At present, Audible Magic markets Copysense to universities that are looking to police and/or discourage the use of peer-to-peer file trading software on their networks.[23] When a Copysense appliance detects a violation, it can take a number of actions, including logging that violation, interrupting the TCP protocol by forging reset (RST) packets,[24] or sending a series of increasingly severe notices to the user suspected of engaging in infringement (this regime, which has also been proposed for government use, is known as "graduated response" or "three strikes").[25]

Like all copyright filters, Copysense has a number of limitations. Most glaringly, it only monitors known P2P networks, which means that violators who download material using the

[21] According to the Audible Magic website (see http://audiblemagic.com/products-services/contentsvcs/customers.asp).
[22] See Audible Magic website (http://www.audiblemagic.com/index.asp).
[23] See the Audible Magic website (http://audiblemagic.com/products-services/copysense/copysense-university/).
[24] This technique, which was also employed by Comcast to sever BitTorrent connections, tricks both the host and the client machine into believing the other has reset the connection. See "Packet Forgery By ISPs: A Report on the Comcast Affair," Electronic Frontier Foundation, November 2007 (http://www.eff.org/wp/packet-forgery-isps-report-comcast-affair).
[25] For more background on "three strikes," see "Should Online Scofflaws Be Denied Web Access?," *The New York Times*, April 12, 2009 (http://www.nytimes.com/2009/04/13/technology/internet/13iht-piracy13.html). For more technical information on Copysense's specific implementation of three strikes, see Audible Magic's website (http://audiblemagic.com/products-services/copysense/copysense-university/graduatedresponse.asp).

FTP or HTTP protocols will remain undetected. Furthermore, Copysense is strictly reactive and requires a great deal of data in order to make a positive match—and this data must be provided by the copyright holder. This is to say that Copysense can only detect works that have been registered with Audible Magic.[26]

Copysense must also possess the P2P network's metadata in order to correctly identify a file as it travels over the P2P network. Not only does this metadata often vary from network to network, it also changes if any aspect of the file is altered (i.e. if a user changes the metadata contained within the ID3 tag of an MP3 file), which further undermines the filter's efficacy.

B. Red Lambda Integrity

Despite its shortcomings, Audible Magic's Copysense seems less overinclusive than other copyright filtering technologies on the market. Other solutions utilize traffic analysis as a detection mechanism, which is to say that they simply block content at the protocol or applications level, without even attempting to make a determination as to the legality of the content that is in transit. Red Lambda's Integrity is one such popular technology. Red Lambda boasts that Integrity "Uses advanced hybrid DPI/Behavioral IDS to monitor network traffic for undesired protocols and applications,"[27] including P2P apps, IRC, FTP, IM, Skype, proxy use and application tunneling over HTTP, HTTPS, DNS and ICMP.[28] Red Lambda blocks a wide range of file transfers that take place using a wide range of protocols, presumably resulting in the blocking of untold numbers of authorized and/or non-protected works, including open-source software, independent music releases and public domain material. Even as it casts an

[26] See Audible Magic's website (http://www.audiblemagic.com/products-services/registration/).
[27] See Integrity features as listed on Red Lambda website (http://redlambda.com/integrity.php?p=features).
[28] See Red Lambda press release, "Joint House and Senate Committee Endorse Red Lambda's Integrity Solution to Fight Music and Movie Piracy on University Networks" (http://www.redlambda.com/files/press080408_red_lambda_house_senate_committee_recommend_rl.pdf).

unnecessarily wide net, Integrity is unable to detect browser-based file-sharing services, which are becoming increasingly popular with file sharers[29]. In many ways, Integrity exemplifies the shortcomings of copyright filters: it blocks a great deal of legal content, while still allowing a considerable amount of illegal content to pass.

C. Vobile VideoDNA

Vobile's VideoDNA is a fingerprinting technology that is used to identify video and audio content. At present, Vobile does not market a copyright filtering solution; rather, its technology is used mostly for identification purposes, though it does offer a Software as a Service (SaaS) application that automates the process of identifying content and sending DMCA (Digital Millennium Copyright Act) takedown notices to the sites that are suspected of hosting infringing content.[30] According to Vobile's website, VideoDNA is currently being used by "all major Hollywood studios,"[31] to track the movement of content online and to identify opportunities for sending DMCA takedown notices. As we have seen, not only do these DMCA notices often target legal, fair uses of content, some studio representatives have explicitly stated that they don't consider fair use when sending such notices.[32] Recently, however, a court in the Northern District of California held that copyright holders are required to consider fair use when sending DMCA takedown notices.[33] On its website, Vobile hints at the complex legal issues that surround its products, though it makes no claim that its technology can be used without resulting in false positives over-indicating infringement. In fact, the language on the company's website

[29] "German Court Rules Against Rapidshare," *Billboard*, June 23 2009 (http://www.billboard.biz/bbbiz/content_display/industry/e3i6fad5a2a1d8e51328f91857dabe3e123).
[30] See Vobile's website (http://www.vobileinc.com/solutions.html).
[31] *Ibid.*
[32] "Universal Says DMCA Takedown Notices Can Ignore 'Fair Use'," *Wired*, July 18, 2008 (http://www.wired.com/threatlevel/2008/07/universal-says/).
[33] "Judge Rules That Content Owners Must Consider Fair Use Before Sending Takedowns," Electronic Frontier Foundation, August, 20 2008 (http://www.eff.org/deeplinks/2008/08/judge-rules-content-owners-must-consider-fair-use-).

suggests that the overinclusive nature of VideoDNA might be a desirable feature for content owners looking to eliminate all uses of their content online—whether legal or not ("In this digital age, content is often transcoded, mashed-up and transformed in a variety of ways before finding its way online. This creates a need for the content owner to actively monitor many online sharing sites to identify their content and decide whether to allow it to remain on the site or send a DMCA notice to the site operator asking for the copy to be taken down.").[34] It's also worth noting that U.S. ISP AT&T and entertainment companies Disney and NBC Universal are major investors in Vobile.[35]

[34] See Vobile's website (http://www.vobileinc.com/solutions.html).
[35] "AT&T, NBC and Disney invest in Vobile's VideoDNA," *Broadcast Engineering*, November 12, 2007 (http://broadcastengineering.com/news/att-nbc-disney-vobile-videodna-1112/).

3. Limitations and Consequences of Copyright Filtering

In the previous section, we discussed the different methods that might be used to automatically identify and filter out copyrighted content. In so doing, we also identified some of the limitations and shortcomings of copyright filters. As we have seen, depending on the technology used to identify copyrighted works, copyright filters will be underinclusive, overinclusive or both. The filter will fail to identify all copyrighted works that pass through it, will filter out legal, legitimate content or, as is the case with most filtering technologies currently on the market, the filter will fail on both counts.

In the following section, we will summarize some of the technological limitations of copyright filters and will discuss an unintended consequence that can be expected if copyright filtering is done at the ISP level: an encryption arms race.

I. TECHNOLOGICAL LIMITATIONS OF COPYRIGHT FILTERS

In this section, we will briefly summarize the technological limitations of copyright filters. The non-technical limitations of filters will be discussed in depth in the legal analysis section.

A. Architecture is a Poor Indicator

Both peer-to-peer and client-server architectures carry legitimate traffic. As such, it is impossible to block traffic based on architecture alone, without also blocking legal content.

B. Protocol is a Poor Indicator

Like with specific types of architecture, a great deal of legitimate traffic is carried over protocols that are popular with peer-to-peer users. Many open-source software developers, for example, use the BitTorrent protocol to deliver their software to users efficiently.[36] Without the use of BitTorrent, which allows anyone to distribute a large file online without paying for bandwidth costs, many small and independent software developers would not have the means to distribute their software online. In addition to such small entities, even large corporations like Blizzard Entertainment[37] and public broadcasting entities like the Canadian Broadcasting Corporation and the Norwegian Broadcasting Corporation have used BitTorrent to make files available to users.[38] Finally, the similarities between P2P protocols and client-server protocols can result in the blocking of traffic that is completely unrelated to illicit file sharing. For example, when Comcast attempted to block the BitTorrent protocol, the company also inadvertently blocked traffic related to Lotus Notes, a commercial software suite used by businesses for email, scheduling and file-sharing purposes.[39]

C. Media Type is a Poor Indicator

As is the case with architecture and protocol, media type (a category assigned to a certain type of media, *i.e.* film, music, image, etc.) is a poor indicator, as blocking any media type outright will result in the blocking of fair use, legally-shared and public domain content.

II. FILTER PROCESSING ADDS LATENCY

[36] See "BitTorrent," *Wikipedia* (http://en.wikipedia.org/wiki/BitTorrent_(protocol)#Software).
[37] See "Blizzard Downloader," *World of Warcraft Universe Guide* (http://www.wowwiki.com/Blizzard_Downloader).
[38] See "BitTorrent," *Wikipedia* (http://en.wikipedia.org/wiki/BitTorrent_(protocol)#Broadcasters).
[39] 'Comcast Is Blocking More Than BitTorrent, Including Lotus Notes," *InformationWeek*, October 22, 2007 (http://www.informationweek.com/blog/main/archives/2007/10/comcast_is_bloc.html).

No matter how sophisticated filtering technologies eventually become, filtering will always slow the speed of traffic that travels over the network, so long as it is used to prevent unwanted or disfavored content from reaching users.[40] While the Internet was originally designed as a system that forwarded packets on toward their destination as quickly as possible, filters alter this behavior by analyzing each packet before determining how it should be treated or flagged. This process introduces delay or latency into the normal packet delivery process and as such, a filtered network will always be slower than an unfiltered network.

A. Filtering Technologies Can Expose Networks to Security Risks

As we saw in the case of China's "Green Dam" filtering software, any piece of software that sits between users and the network or which is installed on a large number of PCs pursuant to a mandate will provide a highly visible target for malicious hackers.[41] Presumably, a great deal of resources will be required to defend the filter against those who wish to subvert, control or misuse it—a possibility that is discussed in the next sub-section.

B. A Copyright Filter Could be Intentionally Misused for Censorship Purposes

At present, content filtering is used by a number of governments to censor Internet content, so as to control the flow of information. According to Reporters Without Borders' annual report, "Internet Enemies," China, Iran, North Korea, Syria and Cuba are among the nations who most tightly control the flow of information using some form of active content

[40] While it is possible to create a filter that would not slow traffic on a network, that filter would not be able to block or degrade traffic in real-time. Such a filter would "clone" the traffic that passed through it and then send that data elsewhere for analysis, without interrupting the flow of the original traffic. This data could then be analyzed and action could be taken after the fact (infringement notices could be sent to users, etc.).
[41] "Big Vulnerabilities in China's Mandatory Filtering Software," *Ars Technica*, June 1, 2009 (http://arstechnica.com/tech-policy/news/2009/06/big-vulnerabilities-in-chinas-mandatory-filtering-software.ars).

filtering.[42] This list also lists two additional "democracies under surveillance": Australia and South Korea.[43] In 2008, a minister in the ruling Australian Labour Party attempted to institute a nationwide filtering mandate with the aim of blocking content relating to child pornography. In late 2008 and early 2009, a series of leaked "blacklists" revealed that the filter would also block several legal sites with no relation to child pornography, including sites hosted by Wikipedia, the online encyclopedia.[44] As a result, the filtering mandate has yet to be enforced, though members of the Labour party have continued to push for implementation, most recently discussing plans to use the filter to prevent any Australian, regardless of age, from accessing video game content that has been deemed inappropriate for someone who is 15-years of age or younger.[45] Meanwhile, in South Korea, the government uses filtering technology to block access to a number of sites containing political content, most commonly sites that are deemed to have a "pro-North Korean" agenda.[46]

The timeliest example, however, comes from Iran, where filtering technologies are now being used to block access to popular websites like Twitter, YouTube and Facebook, in the wake of widespread political unrest. According to a report that appeared in *The Wall Street Journal*,[47] Iran's communications system, which serves an estimated 23 million Internet users, was designed to allow government surveillance of any and all online communications that originate or terminate in Iran. As the Iranian government maintains a monopoly on telecommunications

[42] "Internet Enemies" (2009 edition), Reporters Without Borders, April 9, 2009 (http://www.rsf.org/IMG/pdf/Internet_enemies_2009_2_.pdf).
[43] *Ibid.*
[44] "Rudd's internet blacklist includes dentist, kennel, tuckshop," *The Courier Mail*, March 20, 2009 (http://www.news.com.au/couriermail/story/0,23739,25214413-953,00.html).
[45] "Web filters to censor video games," *The Sydney Morning Herald*, June 25, 2009 (http://www.smh.com.au/digital-life/games/web-filters-to-censor-video-games-20090625-cxrx.html).
[46] "South Korean Internet Censorship," *1StopKorea.Com*, April 2007 (http://1stopkorea.com/index.htm?korean-internet-censorship.htm~mainframe).
[47] "Iran's Web Spying Aided by Western Technology," The Wall Street Journal, June 22, 2009 (http://online.wsj.com/article/SB124562668777335653.html).

services, the goal of omnipresent Internet surveillance was easily met, as it simply required the deployment of filters, powered by Deep Packet Inspection (DPI) technology, at "a single choke point" in the government network, through which all inbound and outbound traffic passes.[48] Apparently, the Iranian government first installed this hardware for the purported purpose of blocking pornography, citing "lawful intercept"—an internationally-recognized concept that "relates to intercepting data for the purposes of combating terrorism, child pornography, drug trafficking and other criminal activities carried out online."[49] This example illustrates that the act of filtering is a slippery slope. While filtering technology might be deployed to serve a legitimate purpose—be it to stem the flow of child pornography or copyrighted content—when placed in the wrong hands, that same technology can become a highly effective instrument of private or governmental censorship. It is for this reason that the Open Internet Coalition urged Congress to convene hearings to address the use of these technologies domestically, warning that policymakers "must fully understand the implications of wide deployment of deep packet inspection technology [in order to make] decisions to prevent its misuse in the United States".[50]

III. THE IMPLEMENTATION OF COPYRIGHT FILTERS WILL RESULT IN A TECHNOLOGICAL ARMS RACE

Even if we could somehow design a filter that was 100 percent accurate, users would actively work to circumvent that filter, in order to access the content of their choice without exposing themselves to liability. The end result would be a technological arms race whereby users attempt to evade the filter while those maintaining the filter attempt to thwart those evasion attempts. As is the case in most arms races, this is a cat-and-mouse game where everyone loses:

[48] *Ibid.*
[49] *Ibid.*
[50] See Open Internet Coalition's letter to Congress, June 29, 2009 (http://www.openinternetcoalition.org/files/OIC_DPI_Iran_062909.pdf).

network protocols used by all users will be slower and less efficient and filters will become even more costly and less effective than they are today. Given that service providers are likely to pass on any additional costs to the consumer, not only will these networks be slower but subscription prices will also be higher.

A similar, if brief, arms race was observed during the FCC's proceeding regarding Comcast's blocking of the BitTorrent protocol. As the legal process wound its way through the FCC, the BitTorrent development and user community did not stand idly by. Rather, the community took matters into its own hands, devising a number of different workarounds that allowed users to bypass Comcast's traffic analysis filter.[51] Unfortunately, while these methods rendered Comcast's content management less effective, they did so at the cost of the BitTorrent protocol's efficiency.[52]

Another arms race is currently unfolding as a result of the Iranian government's politically motivated web filtering regime. Soon after the Iranian government started blocking websites in an attempt to censor news relating to the June 2009 election protests, web users outside of Iran made unfiltered connections available to users inside Iran, through the use of web proxies and Tor bridges.[53] A few weeks after the protests began, a group of software developers announced the upcoming release of Haystack, a software package for the Windows, Macintosh and Unix operating systems that was specifically designed to circumvent the Iranian

[51] "Beating Comcast's Sandvine On Linux With Iptables," *Slashdot*, June 30, 2008 (http://tech.slashdot.org/tech/08/06/30/0249249.shtml).
[52] "As Expected, BitTorrent Providers Planning To Route Around Comcast Barrier," *TechDirt*, February 18, 2008 (http://www.techdirt.com/articles/20080215/171450267.shtml).
[53] These techniques allow users to tunnel their traffic through gateways, so as to disguise that traffic's point of origin. See "Help Protestors in Iran: Run a Tor Bridge or a Tor Relay," Electronic Frontier Foundation, June 29, 2009 (http://www.eff.org/deeplinks/2009/06/help-protesters-iran-run-tor-relays-bridges).

government's web filters.[54] In contrast to previous methods, Haystack will significantly reduce the amount of technical expertise required to subvert Iran's filters, allowing any Iranian with basic computer literacy skills to access the web freely.

Meanwhile, in France, hackers have announced plans to release custom router firmware known as "the HADOPI router," named after French President Nicholas Sarkozy's famed "HADOPI" law, which would have required French ISPs to kick users off of their networks after receiving three allegations from copyright holders that that user had engaged in acts of online infringement (the HADOPI law was rejected by the French Constitutional Council).[55] The HADOPI router software allows a user to route his or her traffic at random through other nearby connections, effectively making that user's traffic untraceable by those who would enforce the HADOPI law.[56]

If an ISP attempts to implement a network-wide copyright filter, we will likely see a similar arms race, albeit on a far larger scale. There are a number of different ways that users and developers might subvert a copyright filter, thereby instigating an arms race. These methods are discussed in detail below.

A. Encryption

Encryption is the science by which communications are scrambled so that an observer who views that communications data—for instance, an ISP—cannot determine what its contents are. To a party without the encryption "key"—the Rosetta Stone that allows encrypted

[54] "Haystack Anti-Censorship Tool Specifically For Users in Iran, to Launch Soon," *Boing Boing*, July 6, 2009 (http://www.boingboing.net/2009/07/06/haystackanticensor.html).
[55] "Top Legal Body Strikes Down Anti-Piracy Law," *France 24*, June 10, 2009 (http://www.france24.com/en/20090610-top-legal-body-strikes-down-anti-piracy-law-hadopi-constitutional-council-internet-france).
[56] "French Hackers Unveil the HADOPI Router: Cracks Nearby WiFi and Makes Your Traffic Traceable to Your Neighbors," *Boing Boing*, July 10, 2009 (http://www.boingboing.net/2009/07/10/french-hackers-unvei.html).

communications to be decrypted—the data appears as little more than a random string of characters. Encrypting a communications stream makes it very difficult for an observer to not only view the data but to also determine that data's type (music, video, etc.) or protocol (HTTP, BitTorrent, etc.). This holds the potential to circumvent both content and traffic analysis filters, as neither type of filter will be able to make a determination when presented with an encrypted data stream. Unfortunately, encryption also has drawbacks—most notably, it increases the amount of computation required on both sides of a communication, subsequently increasing the time and cost associated with that transfer. It is for this very reason that even very large corporations like Google (which offers equipment that is used for security and privacy purposes) have been hesitant to enable encryption by default in their products.[57]

Encryption has long been used in conjunction with many Internet-based applications, and is increasingly being used with many more. HTTP, the protocol associated with most browser-based web traffic, has been available in a secure form since at least 2000.[58] The BitTorrent file transfer protocol first saw an implementation of header encryption in 2005[59] and in 2006, a standardized, full-protocol "Message Stream Encryption" method was devised.[60] Transport Layer Security, meanwhile, provides a security and encryption specification that can be used in conjunction with a wide variety of protocols and is available as an easily integrated software package on many platforms.[61]

---

[57] See "Gmail to go with HTTPS by Default," *WindowsITPro*, June 17, 2009 (http://windowsitpro.com/Articles/Index.cfm?ArticleID=102316).
[58] See Internet Engineering Task Force, "Request for Comment 2818, HTTP Over TLS," May 2000 (http://tools.ietf.org/html/rfc2818).
[59] See IPP2P, "News," April 2006 (http://www.ipp2p.org/news_en.html).
[60] See "Message Stream Encryption Specification," January 27, 2006 (http://www.azureuswiki.com/index.php?title=Message_Stream_Encryption&oldid=4197).
[61] See Internet Engineering Task Force, "Request for Comment 5246, The Transport Layer Security (TLS) Protocol Version 1.2," August 2008 (http://tools.ietf.org/html/rfc5426).

Encryption, for all its benefits, does not always render a communication fully opaque. For instance, even if a filter is unable to make a direct determination regarding the nature of the content or the protocol used, it might still be able to analyze the size, sequencing, and timing of data packets, in an attempt to make a protocol determination.[62] In this way, a filter may still be able to determine with some degree of reliability that a user is using a particular protocol, though it would not be able to determine if that protocol was being used to send a licensed movie, an infringing song, or an open-source software package.[63] As such, even if a filter were able to glean information from an encrypted data stream, it would still face limitations that would likely result in either the blocking of legal content, the allowing of illegal content or both. Furthermore, such analysis can also be easily circumvented.

B. Protocol Obfuscation

When filters target specific protocols (for example, BitTorrent), that protocol can be "obfuscated" in order to make it more difficult to detect both the protocol and what type of content is being carried by that protocol. This process generally does not require the sharing of keys (as is the case with encryption); rather, obfuscation merely rearranges data in order to create a more complicated layout that requires additional analysis if that data is to be decoded. If a user used protocol obfuscation on a copyright filtered network, the operator of that network would have to look more closely at that user's communications. This would cause the ISP to incur greater costs, would lower the efficiency of the network and would further compromise the privacy of the user, who in most cases would be engaging in only lawful activity in the first place.

---

[62] The particulars of such traffic and connection pattern analysis methods are discussed in this paper's technical section.

[63] "Company Cracks BitTorrent Protocol Encryption and Introduces Tracker Whitelists," *TorrentFreak*, April 27, 2007 (http://torrentfreak.com/company-cracks-bittorrent-protocol-encryption-and-introduces-tracker-whitelists/).

Protocol obfuscation can take many forms. The form of obfuscation used would likely vary depending on the type of filter and the particulars of the protocol in question. For instance, if a filter attempted to identify a protocol based on the size of the packets,[64] then an effective obfuscation technique might be to pad the data or packet headers with random amounts of extraneous blank space, in order to bypass the filter. While methods like data padding are fairly generic, others have been devised for use with specific protocols and the filters that detect those protocols. Some BitTorrent clients, for example, offer a "lazy bitfield" option, which intermittently sends phony messages suggesting that a piece of data is either incomplete or unavailable—a method devised specifically to evade certain filtering methods favored by cable ISPs.[65] Other, even more sophisticated methods masquerade targeted protocols as other data types, including email and web traffic.[66]

As filters advance in sophistication, so will the methods used for purposes of obfuscation. Filters that identify protocols by analyzing the timing between packets can be evaded by randomly altering the time at which the host sends data—a method that generally does not require the receiver to be aware of the change. Filters that work by correlating the use of multiple users who connect to the same host can be evaded by rerouting data through additional hosts.[67] In all cases, each additional level of complexity employed by a filter will result in more sophisticated methods of data obfuscation. This arms race will inevitably result in more data on the network, more resources being used by hosts attempting to communicate, greater resource

[64] This method, referred to as connection or traffic pattern analysis, is discussed in this paper's technical section.
[65] For an explanation of the use of this option to "[a]lways send a faked incomplete bitfield, and instead send remaining completed piece info via Have messages", see Azureus Wiki, "Command Line Options" (http://www.azureuswiki.com/index.php/Commandline_options). Certain ISP (i.e. Cablevision's Optimum Online) networks seem to block peer seeding via "complete" bitfield message filtering."
[66] See "A Robust Data Obfuscation Approach for Privacy Preserving Collaborative Filtering" by Rupa Parameswaran (http://etd.gatech.edu/theses/available/etd-05082006-193521/unrestricted/parameswaran_rupa_200608_PhD.pdf).
[67] See Tor website (http://www.torproject.org).

requirements and higher costs to those trying to filter and more invasive analysis of end-user communications.

IV. THE RAMIFICATIONS OF THE ARMS RACE

While the arms race between filter architects and end users will result in inefficiencies on both ends, history suggests that end-users will have a slight advantage. This is due to the fact that there is an entire field of study devoted to ensuring that users can communicate freely without interlopers being able to read or alter their communications. As Paul Kocher, President and Chief Scientist at Cryptography Research, Inc. once put it at a NIST Physical Security Testing Workshop, "Moore's Law favors the cryptographer."[68]

Effective, efficient, and even government-approved[69] encryption standards are commonplace, have low (though not insignificant) cost barriers to use and are impractical or impossible for ISPs to circumvent. Other data obfuscation techniques, meanwhile, require relatively simple changes on the user's end, yet require a far more detailed inspection of data, more computing resources, and a greater volume of sample data on the ISP's end. In many cases, users will need only to upgrade or tweak the settings of the software on their home PCs in order to evade the newest, most expensive filtering technology.[70] Considering that the filtering arms race will result in increased costs and decreased performance both for service providers and users, we can conclude that the cost of filtering will far outweigh the benefits—especially when

[68] See "The Increasing Complexity and Need for Validation," NIST, September 26, 2005 (http://csrc.nist.gov/groups/STM/cmvp/documents/fips140-3/physec/presentations/physecpre02.pdf).
[69] See National Institute for Standards and Technology, "Announcing the Advanced Encryption Standard (AES)," November 26, 2001 (http://csrc.nist.gov/publications/fips/fips197/fips-197.pdf).
[70] For examples of high-cost digital protection schemes that were quickly circumvented by users, see "Quick Guide: Cracked DRM Systems," *BBC News*, August 27, 2007 (http://news.bbc.co.uk/2/hi/technology/6944830.stm).

one considers the primary aims of the National Broadband Plan, which are to increase the speed and adoption of broadband services while decreasing costs.[71]

If service providers and content owners choose to go down this path, they might soon find themselves unable to identify the vast majority of data on their network as the result of users migrating to encrypted systems. In such a world, where service providers would be unable to make an accurate determination regarding the type of data traversing their networks, ISPs might decide to block or delay any and all data that cannot be identified. While this presents a predicament in and of itself, the problem will be amplified if users react by masquerading their targeted data as commonly encrypted web traffic (for example, email or traffic related to online commerce). If this happens, service providers will find themselves in an unenviable position: they will be forced to choose between blocking or delaying legitimate traffic and allowing through the very data they were attempting to filter in the first place.

In such a world, where providers might elect to block or degrade all unidentified or unreadable traffic, users will also be forced to make a difficult choice, one between privacy and access. If they choose to read their email over an encrypted connection, as nearly every service provider recommends,[72] they may find that their traffic is degraded or blocked. If users choose not to utilize encryption, they will be subjecting their data and communications to inspection by their ISP, content companies and anyone else who is able to view that traffic—including users in

[71] See "FCC Launches Development of National Broadband Plan," April, 8, 2009 (http://hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-289900A1.pdf).
[72] See Comcast Customer Central, "Configuring Outlook Express to send and receive email while traveling" (http://customer.comcast.com/Pages/FAQViewer.aspx?Guid=0a2e36ac-2737-412b-b0ab-d2bf6a0fe54d).

the same office, on the same wireless network, or elsewhere along the data's path.[73] This choice, between security and privacy, is one that Internet users should not be forced to make.

[73] See "Using Wireless Technology Securely," US-CERT, available at (http://www.us-cert.gov/reading_room/Wireless-Security.pdf).

4. Economic Analysis

While a full economic analysis of the theoretical impact of copyright filtering is outside of the scope of this paper, there are certain issues regarding the economics of filtering which cannot be ignored. These issues are briefly addressed in the sections below.

I. WHO WILL PAY FOR COPYRIGHT FILTERING?

The most obvious economic question that filtering raises is a pragmatic one. Who will pay for the deployment, maintenance and upkeep of the hardware, software and personnel that will be required to implement copyright filtering at the ISP level? Regardless of the technology used, ISPs will be required to deploy dedicated hardware devices on their network, which we can only assume will be purchased or rented from a hardware vendor. We can also assume that a certain number of manpower hours will be required to maintain this hardware and any associated software and keep it operating at maximum efficiency. And when this equipment needs to be upgraded, maintained or replaced, there will be an additional, recurring cost.

As this hardware and software will sit on the ISP's network, one might naturally assume that the ISP will cover the cost associated with filtering its network. However, seeing how there is no clear benefit and several harms to ISPs, there is no reason why a service provider should feel compelled to cover the costs associated with copyright filtering. Perhaps, as is the case with the Communications Assistance for Law Enforcement Act (CALEA), the party requesting the examination of communications data would foot the bill.[74] In the case of law enforcement

[74] See the Communications Assistance for Law Enforcement Act (CALEA) (http://en.wikisource.org/wiki/Communications_Assistance_for_Law_Enforcement_Act_of_1994).

wiretapping, that party is the United States government. In this case, that party would be the content industry.

In addition to the costs associated with filtering, ISPs have a number of non-monetary incentives to not engage in copyright filtering on their networks. Some of these incentives are legal in nature and will be discussed at length in the legal section of this paper. Others are purely technical. As we have seen, copyright filters will significantly slow any network on which they are installed, resulting in a decrease in the quality of the customer experience. This will place the ISP who engages in copyright filtering at a competitive disadvantage *vis-à-vis* its competitors.

At least one U.S. ISP has recognized that the drawbacks of copyright filtering far outweigh any perceived benefits. In January 2008, Verizon Communications publicly stated that it has no intention to install copyright filters on its network. "From a business perspective, we really don't want to assume the role of being police on the Internet," Verizon executive vice president Tom Tauke told the audience at a Washington Internet policy conference.[75] "We are leery of using these technologies on our networks."

Other U.S. ISPs, however, have expressed a strong interest in copyright filtering. At the 2008 Consumer Electronics Show (CES), AT&T senior vice president, external and legal affairs, James Cicconi revealed that AT&T was considering the possibility of installing copyright filters on its network and had engaged in discussions with technology companies, content companies like NBC Universal and content industry trade groups like the RIAA and the Motion Picture Association of America (MPAA).[76] This prompted Columbia Law School professor Tim Wu to

[75] "Verizon: We Don't Want to Play Copyright Cop on Our Network," *CNET*, January 30, 2008 (http://news.cnet.com/8301-10784_3-9861402-7.html).
[76] "AT&T and Other ISPs May be Getting Ready to Filter," *The New York Times*, January 8, 2008 (http://bits.blogs.nytimes.com/2008/01/08/att-and-other-isps-may-be-getting-ready-to-filter/).

write an opinion piece for *Slate* magazine, entitled "Has AT&T Lost its Mind?".[77] "…[T]he bizarre twist is that the proposal is such a bad idea that it would be not just a disservice to the public but probably a disaster for AT&T itself," Wu wrote in the article. "Has AT&T, after 122 years in business, simply lost its mind?"

Despite the lack of any clear motive on the part of the ISPs, some have floated theories as to why service providers might be amenable to go along with the content industry's copyright filtering agenda. In the case of AT&T, it is likely that the decision was motivated by the company's other, non-Internet offerings, namely its U-Verse TV television service. Since AT&T also acts as Multichannel Video Programming Distributor (MVPD),[78] copyright filtering would provide a means by which the company could demonstrate its commitment to copyright enforcement, thereby currying favor with the studios from whom it must license video content.

In the case of Comcast, a company that did not filter for copyright but which did block the BitTorrent protocol using a traffic inspection method of the type described earlier in this paper,[79] it was suggested that the company was simply hoping to boot high-bandwidth users off of its network, in the hope of increasing speeds on its oversubscribed network without having to invest in additional network capacity.[80] As has been noted elsewhere in this paper, this act earned Comcast a strong rebuke from the FCC.[81]

---

[77] "Has AT&T Lost its Mind?," *Slate*, January 16, 2008 (http://www.slate.com/id/2182152).
[78] The term "MVPD" includes cable television providers, direct broadcast satellite providers and wireline video providers, including those who sell IPTV services.
[79] See section 2.II.A. of this paper.
[80] See "Formal Complaint of Free Press and Public Knowledge Against Comcast Corporation for Secretly Degrading Peer-to-Peer Applications," November 1, 2007 (http://www.publicknowledge.org/pdf/fp_pk_comcast_complaint.pdf).
[81] See the FCC's Memorandum Opinion and Order, August 1, 2008 (http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-08-183A1.pdf).

Some have even gone further to suggest that these ISPs have an additional incentive to block video online for competitive reasons.[82] As was mentioned earlier in this paper, as a result of technological limitations, filters will be overinclusive. This means that a great deal of lawful video content--mashups, parodies, satire and non-copyrighted works--is likely to be blocked by any copyright filter that is implemented by an ISP. From a user's perspective, the end result will be that a significant amount of video content is no longer available online. This will, in effect, make watching online video a less attractive prospect for the user, who might turn to alternative sources for video entertainment as a result. This could offer an advantage to any MVPD service that seeks to compete with online video.

At present, the prospect of free streaming and downloadable online video clearly offers an attractive alternative to users who are used to paying an average monthly fee of $71 per month for cable television[83]—especially when one considers that the average monthly cost for a broadband Internet connection is only $39.[84] Indeed, one recent study found that 77 percent of Internet users watch online video and that that 20 percent of those viewers watch less programming on a television set as a direct result of their online viewing habits.[85]

Internet video distribution also provides an attractive proposition for content creators, by allowing those creators to beam content more directly to the consumer, cutting out the MVPD middleman while retaining the ability to derive revenue from that content through advertising. Indeed, some forms of programming are already commanding higher ad rates online than on

[82] "AT&T TOS Update Shows Network Management's True Colors," Public Knowledge Blog, September 11, 2008 (http://www.publicknowledge.org/node/1736).
[83] "Study: Average Cable TV Bill Is $71 Per Month," *Multichannel News*, April 16, 2009 (http://www.multichannel.com/article/196364-Study_Average_Cable_TV_Bill_Is_71_Per_Month.php).
[84] See "Home Broadband Adoption 2009," Pew Internet and American Life Project, June 17, 2009 (http://www.pewinternet.org/Reports/2009/10-Home-Broadband-Adoption-2009.aspx).
[85] "Magid Video Futures 2009: Opportunities in Online Video," Frank N. Magid Associates, Inc., June 2009 (http://www.magid.com/metacafe.pdf).

television. The popular prime-time television program "The Simpsons," for example, has a higher Cost Per Thousand (CPM) price on the streaming video site Hulu than it does on network television.[86]

In this light, it becomes clear why ISPs and their shareholders might want to invest in copyright filtering, especially considering how many service providers now offer video services either alongside Internet access or as part of a package containing Internet, telephony and video services (e.g. so-called "Triple Play" offerings). In fact, late last year, the RIAA announced that it was working with "major ISPs" to kick subscribers suspected of engaging in infringement off of those ISP's networks.[87]

II. COPYRIGHT FILTERING HOLDS THE POTENTIAL TO DISRUPT THE INTERNET ECONOMY, OUR MOST PROMISING ENGINE FOR ECONOMIC GROWTH

As we will discuss in the following section, the Internet economy is a delicate ecosystem, consisting of a number of different layers of hardware, software and service providers, all working in tandem. This is a fact of which Congress was acutely aware during the process of crafting the Digital Millennium Copyright Act (DMCA).[88] Any change made to how the Internet operates will have a ripple effect that will be felt widely, at all levels of the ecosystem. Tinkering with the basic paradigm of Internet access is extremely unwise, especially at a time when the Internet's ability to promote economic growth and ability to provide citizens with economic opportunity is needed more than ever.

---

[86] "D'oh! The Simpsons Worth More on Hulu Than on FOX," *NewTeeVee*, June 29, 2009 (http://newteevee.com/2009/06/25/doh-simpsons-worth-more-on-hulu-than-on-fox/).
[87] "Music Industry to Abandon Mass Suits," *The Wall Street Journal*, December 19, 2008 (http://online.wsj.com/article/SB122966038836021137.html).
[88] See section 5.II. of this paper, "Copyright Filtering Will Undermine the Safe Harbor Provisions Granted to ISPs Under the Digital Millennium Copyright Act (DMCA)".

As was discussed in the technological analysis section of this paper, one likely consequence of copyright filters is that they will inject delay into commonplace transactions on the Internet. Without an actual implementation of a copyright filter on a U.S. ISP's network to analyze, it is difficult to predict just how disruptive these delays will be. Given that a copyright filter will likely have to a) stop multiple packets traversing the network b) analyze those packets c) cross-reference the analyzed data against a database of "known" or "protected" copyrighted works d) make a determination as to how to flag those packets (e.g. allow/delay/discard) and e) act on whatever flag has been assigned to the packets in question, it is clear that for each transaction online, the filter will be required to process a great deal of information. As such, we can only assume that there will indeed be a delay and that the delay will be noticeable, at least for certain types of traffic.[89]

A delay of a few seconds or even a few milliseconds might not seem like much, but in realm of Internet use, even such minor delays can have staggering effects. We know for a fact that Internet users with faster connections, "spend more time online, do more things, and do them more often" than users with slower connections.[90] What's more, studies have clearly demonstrated that delays of even a few milliseconds can result in appreciable changes in user behavior. One study found that 100 ms of additional load time translates into a 1 percent drop in sales for Amazon.com,[91] while another study found that 500 ms of delay resulted in a 20 percent drop in both traffic and revenue for the Google search engine.[92] These statistics attest that users are acutely aware of speed differences when using the Internet and will modify their online

[89] Unless the filter does not block, delay or degrade traffic in real-time (see 3.I.D.).
[90] "The Broadband Difference: How online behavior changes with high-speed Internet connections," Pew Internet & American Life Project, June 23, 2002 (http://www.pewinternet.org/Reports/2002/The-Broadband-Difference-How-online-behavior-changes-with-highspeed-Internet-connections.aspx).
[91] "Faster is Better," *Dries Buytaert*, January 6, 2009 (http://buytaert.net/faster-is-better).
[92] "Marissa Mayer at Web 2.0," *Geeking With Greg*, November 9, 2006 (http://glinden.blogspot.com/2006/11/marissa-mayer-at-web-20.html).

behavior accordingly.[93] So it stands to reason that not only will users potentially spend less time on the Internet, do fewer things online and do them less frequently but also that Internet commerce and advertising will be adversely impacted as a result.

Further complicating the matter is the plain fact that regardless of the technology used, any Internet filtering scheme is likely to create greater delays for certain specific types of traffic. As an example, let's consider an Internet filter that is designed with certain built-in exemptions. Such a filter might scrutinize traffic flowing to and from YouTube.com—a site where material suspected of infringing on copyrights is commonly found—but not traffic traveling to and from Hulu.com—a studio-owned streaming video site where infringing content is rarely found.

This, of course, raises a number of network neutrality[94] concerns, many of which were discussed in the context of the Comcast/BitTorrent proceeding.[95] For example, if BitTorrent is more heavily scrutinized than iTunes, thereby creating a greater delay for BitTorrent users, will major label artists be given a competitive advantage over independent artists who distribute their work using BitTorrent? If websites hosting open-source software are treated as suspect while Microsoft's website is given a speed pass, will independent software developers and the open-source community suffer? If copyright filters were instituted on the networks of major American ISPs, would the next YouTube-like service even launch, knowing full well that its traffic would be delayed and that its service would be adversely affected, resulting in a less than optimal user

[93] For additional information on the relationship between speed and user behavior see "Speed Matters for Google Web Search," Google, Inc., June 22, 2009 (http://code.google.com/speed/files/delayexp.pdf) and "The Effect of Network Delay and Media on User Perceptions of Web Resources," *Behaviour & Information Technology*, 2000, Vol. 19, No. 6, 427-439.

[94] For more information on the issue of network neutrality, see http://www.publicknowledge.org/issues/network-neutrality.

[95] See "Formal Complaint of Free Press and Public Knowledge Against Comcast Corporation for Secretly Degrading Peer-to-Peer Applications," November 1, 2007 (http://www.publicknowledge.org/pdf/fp_pk_comcast_complaint.pdf).

experience (i.e. slow, choppy video)? Depending on their design, copyright filters hold the potential to pick winners and losers on the Internet—a fact that could have serious and far-reaching economic consequences.

Even if we ignore the effects of delay, copyright filters will still harm innovators and businesses. In interfering with in-transit Internet traffic, copyright filters will inevitably cause some legal software and protocols to malfunction. A product like Red Lambda's Integrity, for example, would block all traffic that uses a protocol that is associated with P2P use,[96] thereby interfering with legal P2P services like the Octoshape P2P video player used by CNN.com.[97] While developers might be able to modify their software and services to circumvent the filter, they will have to bear the costs of doing so themselves. Of course, in order for a developer to make such modifications effectively, the filter's architects will have to operate transparently, providing legitimate developers with documentation regarding the filtering algorithms and targeting criteria. This, however, presents the architect of the filter with a Catch 22: if this documentation is made available, it will inevitably make its way into the hands of those who aid infringers, resulting in a less effective filter; if the documentation is not made available, legitimate businesses will be harmed. Therefore, it follows that copyright filters will be either ineffective or will stifle innovation. What's more, all of these problems will be compounded if the architects of the filter are involved in an arms race like the one described in the previous section, making iterative tweaks to the filter and subsequently causing recurring, persistent problems for developers.

[96] See Integrity features as listed on Red Lambda website (http://redlambda.com/integrity.php?p=features).
[97] "CNN Video Streaming Tech Raises Questions," *Ars Technica*, February 10, 2009 (http://arstechnica.com/web/news/2009/02/cnn-p2p-video-streaming-tech-raises-questions.ars).

While the long-term effects that copyright filtering would have on the American economy remain unknown, this much is clear: copyright filtering will be a disruptive force on the Internet, one that will be felt by users, online service providers, businesses and investors. The Internet has become a powerful engine for economic growth due to its open architecture, which encourages investment, by providing a platform whereby anyone can build an application or service that can compete on equal footing with incumbent offerings. It is crucial that we allow the Internet to remain a level playing field. Otherwise, we will discourage investment, competition, new entrants and innovation, dealing a blow to our long-term economic health and international competitiveness.

5. Legal Analysis

The laws that currently govern speech and copyright online are the result of a delicate balance that was reached between a number of competing interests and considerations. These laws have allowed ISPs and Online Service Providers (OSPs) to foster a dynamic ecosystem of speech, entertainment, and debate without fear of legal liability for the actions of users online.

If the federal government imposes a mandate requiring ISPs to filter traffic in search of copyright violations or otherwise encourages or condones the act of filtering, this delicate balance will be disrupted. The protections that innovators and citizens have relied on in building the Internet ecosystem would be substantially weakened if not outright eliminated. More troubling, such legislation would undermine current legal and constitutional protections for speech and free expression. Below, we discuss the existing laws with which a copyright filtering mandate would conflict.

I. MANDATORY COPYRIGHT FILTERS WOULD IMPOSE AN UNCONSTITUTIONAL BURDEN ON FREE EXPRESSION, CONTRARY TO THE PRINCIPLES OF COPYRIGHT LAW

Regardless of how it is accomplished, content filtering will interfere with citizens' ability to communicate online. No matter what type of technology is used, no filter will be capable of determining if a communication is authorized, fair use or infringing. As a result, copyright filters will always be overinclusive when blocking online speech. Thus, copyright filters will inevitably interfere with and suppress completely legal forms of speech and expression online. While such interference is worrisome when practiced by a private company, it may well be unconstitutional if imposed by government mandate.

Under the law, a copyright owner is never granted complete control over a copyrighted

work.[98] Limitations on and exceptions from copyright keep copyright law from conflicting with the First Amendment rights of citizens. Fair use and other limitations such as the requirement of originality, the idea/expression dichotomy, and the doctrine of thin copyright[99] allow for free expression including protected forms of speech like parody and criticism. As the Supreme Court has explained, "Copyright … does not impermissibly restrict free speech, for it grants the author an exclusive right only to the specific form of expression … and it allows for 'fair use' even of the expression itself."[100]

A fair use of a copyrighted work is therefore protected free speech. Proponents of copyright filtering suggest that the filtering of copyrighted material would be a straightforward and entirely legal process. However, the nuances of copyright law make distinguishing between a lawful and infringing use of a piece of copyrighted content challenging even for courts. As such, no filtering technology, no matter how advanced, will ever be able to make fair use determinations with 100 percent accuracy. And as courts have held, no law should "allow any copyright owner, through a combination of contractual terms and technological measures, to repeal the fair use doctrine."[101] Because a use labeled by a filter as "unauthorized" is not necessarily an illegal use, no automated system should be allowed to give the desires of a copyright holder priority over the First Amendment. Furthermore, a government mandate requiring copyright filtering might also run afoul of case law governing prior restraint, which states that speech cannot be preemptively censored, except in extenuating circumstances (i.e.

[98] Exclusive rights "do[] not give a copyright holder control over all uses." *Fortnightly Corp. v. United Artists*, 392 U.S. 390, 393 (1963).
[99] "Thin copyright" refers to the lower level of copyright protection for compilations of fact. The doctrine was first explained by the Supreme Court in *Feist Publications, Inc. v. Rural Telephone Service Company, Inc.* 499 U.S. 340 (1991).
[100] *Eldred v. Ashcroft*, 537 U.S. 186, 197 (2003).
[101] *Chamberlain Group, Inc. v. Skylink Technologies, Inc.,* 381 F.3d 1178, 1202 (Fed. Cir. 2004).

issues of national security).[102]

As discussed in the technological analysis section of this paper, copyright filters are not only overinclusive—they will also be underinclusive. As we have seen, today's filters have proven ineffective at stopping even moderately sophisticated infringers. And there is no indication that future filters will be able to overcome this technological shortcoming. As a result, a government mandate that ISPs filter for copyright infringement would substantially interfere with free speech and non-copyrighted content while providing little or no benefit to copyright holders to justify this constriction of First Amendment rights.

By virtue of their technological limitations, copyright filters will inevitably block some protected forms of speech while allowing some infringement. The simultaneously overinclusive and underinclusive nature of filters will ultimately result in unconstitutional restrictions on free speech, just as other government attempts to block access to objectionable web sites have in the past.[103] Therefore, we can conclude that a government copyright filtering scheme that fails to specifically accommodate fair use is unlikely to be constitutional.

II. COPYRIGHT FILTERING COULD UNDERMINE THE SAFE HARBOR PROVISIONS GRANTED TO ISPS UNDER THE DIGITAL MILLENNIUM COPYRIGHT ACT (DMCA)

Currently, the Digital Millennium Copyright Act offers ISPs certain safe harbors from copyright infringement liability for activity which occurs over their networks.[104] It is not hard to understand the value of this protection to an ISP, and to service providers in general. If ISPs could be held liable for every infringement perpetrated by their customers, they would be exposed to a flood of lawsuits from the sound recording, motion picture and software industries.

[102] *Near v. Minnesota*, 283 U.S. 697 (1931).
[103] *CDT v. Pappert*, 337 F. Supp. 2d 606 (E.D. Pa. 2004).
[104] 17 U.S.C. § 512(a).

And unlike with individual infringers, ISPs are easy to find and likely to have the resources to pay a judgment and are therefore, attractive targets for lawsuits (suing individuals has proven to be a time- and cost-intensive activity that is unlikely to result in a significant financial recovery[105]).

These ISP safe harbors under the DMCA are not the result of a happy accident. During the process of drafting of the DMCA, Congress was convinced of two points—first, that effectively monitoring every bit that travels through an ISP's network was not technologically feasible[106] and second, that if ISPs were forced to monitor their networks, that requirement would effectively cripple the nascent Internet.[107]

Congress ultimately determined that limiting the liability of ISPs would ensure "that the efficiency of the Internet [would] continue to improve and that the variety and quality of services on the Internet [would] continue to expand."[108] Essentially, Congress decided to allow ISPs to develop viable business models, fearful that the specter of liability could, in the words of one commentator, "virtually halt American participation in the emerging information society."[109] Today's Internet is a testament to the wisdom of that decision. It is hard to imagine any common activity on the Internet today—be it making a comment on a news site, posting a photograph on a social networking site, or perusing indexed links on a search engine site—that has not benefited from the DMCA safe harbors.

In order to qualify for the safe harbor provisions, an ISP has to meet a certain set of

---

[105] See*, e.g.*, *Capitol Records Inc. v. Thomas*, 579 F. Supp. 2d 1210. (D. Minn. 2008).
[106] See *Copyright Protection on the Internet*: *Hearings on H.R. 2441 – The NII Copyright Protection Act of 1995 Before The Subcommittee on Courts and Intellectual Property of The House Committee on the Judiciary*, 104th Cong. (1996) (statements of Stephen M. Heaton, General Counsel and Secretary, CompuServe Incorporated and Scott Purcell, Commercial Internet eXchange Association).
[107] *Ibid.*
[108] S. Rep. No. 105-190, at 8 (1998)
[109] *Copyright Protection on the Internet*: *Hearings on H.R. 2441 – The NII Copyright Protection Act of 1995 Before The Subcommittee on Courts and Intellectual Property of The House Committee on the Judiciary*, 104th Cong. (1996) (statements of Stephen M. Heaton, General Counsel and Secretary, CompuServe Incorporated).

requirements, which are outlined in the DMCA. The first of these requirements is that all material that travels over the network must be "initiated by or at the direction of a person other than the service provider."[110] In layman's terms, this means that the ISP must act as a simple conduit through which data travels—any transaction that takes place on the network must be initiated by someone at the edge of the network—either a client or a server—but may not be initiated by someone who sits in the middle of the network. If an ISP implemented a copyright filter, that ISP could, arguably, become an active participant in the chain of transmission, making decisions about what data to transmit and what data to discard. Instead of merely passing a bit of data along, the ISP would inspect, categorize, and possibly interrupt, delay or discard that bit of data. In so doing, the ISP could potentially be disqualified from the DMCA's safe harbor protections and therefore, would be exposed to liability for any infringement that takes place over its network.

Filtering similarly jeopardizes the second requirement that an ISP must meet in order to qualify for DMCA safe harbor protection. This requirement states that the transmission of data must occur "through an automatic technical process without selection of the material by the service provider."[111] This use of the word "selection" is not further clarified in the statute, creating an open question as to what degree of filtering would qualify as "selection". Depending on the level of sophistication of the prioritization process, certain packet management techniques could be interpreted as constituting a "selection" of material. If an ISP could be described as actively selecting what material is allowed to travel over its network, its safe harbor protection would be jeopardized. Furthermore, this selection process could quickly rise to the level of an "editorial function" (*i.e.* choosing to prioritize data from a preferred source over a non-preferred

[110] 17 U.S.C. § 512(a)(1).
[111] 17 U.S.C. § 512(a)(2).

source), which is likely to disqualify an ISP from DMCA safe harbor protection.[112]

In establishing the safe harbor provisions in the DMCA, Congress recognized that in order for the Internet to thrive, ISPs would need to be shielded from liability for their users' actions. Given that any filtering technology that actively inspects the content that flows over an ISP's network may eliminate that ISP's ability to claim protection under the DMCA's safe harbors, it should be clear that any government mandate that requires or condones copyright filtering by ISPs would undermine the safe harbors on which today's Internet was built.[113]

III. ISPS THAT ENGAGE IN PACKET INSPECTION RISK VIOLATING THE ELECTRONIC COMMUNICATIONS PRIVACY ACT (ECPA)

The Electronic Communications Privacy Act (ECPA) extended traditional wiretap privacy protections to include communications between computers.[114] In passing this Act, Congress recognized that private communications that take place on a computer network deserve the same sorts of protections traditionally granted to communications that take place over voice networks. Congress also extended these protections to other types of private communications, including email.[115]

Specifically, ECPA prohibits the interception of electronic communications by those not party to the communication. Intercepting the packets of Internet users, whether for copyright filtering or other purposes, would certainly seem to implicate the statute. Even when law enforcement agencies in the course of their duties wish to intercept communications, the government must justify to a neutral third party why a specific intrusion of privacy is

[112] S. Rep. No. 105-190, at 42 (1998).
[113] "Has AT&T Lost its Mind?," *Slate*, January 16, 2008 (http://www.slate.com/id/2182152).
[114] See P. L. 99-508, 100 Stat.1848.
[115] See 18 U.S.C. §2701-11.

necessary.[116] While the standards of proof that the government must meet in order to intercept private communications vary, under no circumstance is the government authorized to examine all communications in the hope of stumbling upon evidence of illegal conduct.

A commercial enterprise possesses even less justification for warrantless and unwarranted surveillance of private communications.[117] Nor does it seem likely that any of the exceptions provided within ECPA allow an ISP to inspect all of its customers in order to pass judgment on the desirability or legality of their communications. While a provider may engage in activity "necessary incident to the provision of service, or to the protection of the rights or property of the provider of that service," this exception would not apply to an ISP seeking to enforce an outside party's rights. The exception exists to preserve the ability of a network to operate, not to give free reign to a carrier to act as arbiter of its users' content. In order to make use of this exception, a provider must show a "necessary incident," or at the very least, a "substantial nexus," between the monitoring and the protection of the provider's rights and property.[118] Nor would fine print or a clickthrough EULA seem to suffice for the purposes of consent to wiretapping. The necessary amount of notice for consent to wiretapping may be far higher than for consent in ordinary contract terms.[119]

Since such exceptions would likely not apply to ISP filtering, we are left with the fact that a filtering ISP would be unlawfully intercepting electronic communications between citizens, a

---

[116] See 18 U.S.C. § 2710 (15) and § 2711(2).
[117] Controversy surrounding ISP use of DPI for advertising purposes raised the question of ECPA violations by ISPs and their ad network partners. *See, e.g.*, Center for Democracy and Technology, *An Overview of the Federal Wiretap Act, Electronic Communications Privacy Act, and State Two-Party Consent Laws of Relevance to the NebuAd System and Other Uses of Internet Traffic Content from ISPs for Behavioral Advertising,* July 8, 2008 (http://www.cdt.org/privacy/20080708ISPtraffic.pdf) (commenting on the likelihood of ECPA violations by NebuAd and other DPI-based advertising systems).
[118] *Id.* at 5; Paul Ohm, *The Rise and Fall of Invasive ISP Surveillance*, U. Ill. L.R. (forthcoming 2009) (manuscript at 69-71), *available at* http://ssrn.com/abstract=1261344 (detailing the high bar providers must meet to avail themselves of the "protection of rights and property" exception).
[119] *Id.* at 72-73.

violation of the intricately constructed wiretap laws.

Allowing ISPs to filter all communications for potential copyright infringement would turn this finely calibrated system on its head. Congress would be permitting private parties to engage in privacy-violating activities that it elsewhere explicitly prevents private actors and government agencies alike from conducting. If ISPs were authorized to engage in copyright filtering, the end result would be a massive invasion of the privacy of all Internet users, carried out by those few, privileged gatekeepers through whom U.S. web traffic passes.

6. Conclusion

As we have seen, copyright filtering is, at its core, a technology that is ill-suited for use on U.S. broadband networks. Due to their inherently underinclusive and overinclusive nature, copyright filters will never be an effective solution to the problem of online copyright infringement and furthermore, will harm free speech, online privacy and the speed and affordability of broadband services. If implemented, a copyright filtering mandate would force ISPs to inspect all data transmitted by all Internet users, invading the privacy of hundreds of millions of citizens at the behest of one industry. This would likely have far-reaching consequences for the Internet ecosystem as a whole, disincentivizing investment, innovation and creativity and undermining the goals of the National Broadband Plan.

If we are serious about combating copyright infringement online, we should use the tools already at our disposal for identifying, trying and prosecuting infringers. The Digital Millennium Copyright Act (DMCA), for example, already contains robust provisions for copyright owners looking to have infringing content removed from the web and copyright law allows copyright holders to sue for statutory damages that are more than adequate—if not excessive—in many cases.

While existing law already provides copyright holders with a number of options for combating infringement in all its forms, the best solution for content companies will ultimately be one that has little to do with copyright enforcement. If entertainment companies and other content providers adapt to meet the needs of the digital economy, by providing consumers with access to the content that they want in the ways that they want, it is likely that copyright infringement will become a secondary concern for many of these companies. By exploring new,

innovative business models, the content industry can encourage consumers to purchase entertainment goods even when that same content is available for illegal download online. This fact is evidenced by increases in online music sales,[120] decreases in unlawful music file sharing[121] and strong sales of movie tickets[122] and high-definition physical video products like Blu-Ray discs,[123]—all despite the widespread availability of albums and full-length films on P2P file-sharing networks. Additionally, some have suggested that even if unlawful file sharing continues, artists and content companies could still be compensated if they are willing to adopt a more innovative licensing model, such as voluntary collective licensing.[124]

Ultimately, the content industry will have to work closely with technologists, innovators and policymakers to find solutions that make sense in a digital economy. Until then, however, we should not rush blindly to implement ineffective, dangerous solutions like copyright filtering. To do so would be to endanger free speech, to imperil user privacy and to undermine our efforts to deliver the promise of broadband connectivity to all Americans.

---

[120] According to *the New York Times*, online sales of music rose by 27 percent between 2007 and 2008. See "Music Sales Fell in 2008, but Climbed on the Web," *The New York Times*, December 31, 2008 (http://www.nytimes.com/2009/01/01/arts/music/01indu.html).
[121] "Report: UK Filesharing Drops, Even Among Teens," *Ars Technica*, July 13, 2009 (http://arstechnica.com/media/news/2009/07/report-more-uk-users-going-the-legal-route-for-music.ars).
[122] "What Piracy? Movie Biz Sees Record Box Office in 2008," *Ars Technica*, January 5, 2008 (http://arstechnica.com/media/news/2009/01/what-piracy-movie-biz-sees-record-box-office-in-2008.ars).
[123] According to the Digital Entertainment Group, sales of Blu-Ray products tripled between 2007 and 2008. See "DVD Sales Down 5.5% in '08," *Variety*, January 7, 2009 (http://www.variety.com/article/VR1117998174.html).
[124] "A Better Way Forward: Voluntary Collective Licensing of Music File Sharing," The Electronic Frontier Foundation, April 2008 (http://www.eff.org/wp/better-way-forward-voluntary-collective-licensing-music-file-sharing).