The Limits of Filtering:

A Look at the Functionality & Shortcomings of Content Detection Tools

Evan Engstrom and Nick Feamster
March 2017

Engine

Table of Contents

Executive Summary

Nearly twenty years ago, Congress passed the Digital Millennium Copyright Act (DMCA), establishing copyright rules tailored for the emerging internet ecosystem. In order to provide legal certainty for internet startups and investors, as well as to encourage websites and copyright owners to cooperate in addressing online infringements, the DMCA created "safe harbors" for online service providers (OSPs) that limit exposure to secondary infringement claims for the actions of their users. At its core, the DMCA exempts OSPs from monetary liability for user infringements so long as the OSP removes access to infringing material when it becomes aware of it. But, because OSPs cannot readily identify potentially copyrighted material and determine whether or not its use is authorized, the DMCA places the burden of policing infringements squarely on copyright owners.

The DMCA's limited liability regime prompted an unprecedented boom in internet activity to the benefit of users, OSPs, and creators alike. Although the DMCA has succeeded admirably in fostering the growth of the internet, some policymakers and copyright industry lobbyists have advocated for drastic changes to the copyright system to force OSPs to implement content filtering technologies in order to obtain the protections of the safe harbor and potentially punish them if and when those filters fail. Ultimately, these proposals hinge on a misunderstanding of the technical capabilities and likely effect of filtering technologies. In this paper we examine the functionality and inherent limitations of the most common filtering technologies to demonstrate why a mandatory filtering regime would pose grave dangers to the viability of the internet ecosystem in exchange for a minimal effects on online infringement:

- Metadata Filtering: Digital files are often labeled with information describing the contents of a given file (aka "metadata"). In the case of a song file, metadata could cover the song's title, performer, and length. Automated programs can search a group of files for particular metadata matching a target copyrighted work and mark matching files for removal requests. These tools are inaccurate and easily circumvented, as a file's metadata is often inaccurate and can be easily manipulated or encrypted to avoid detection. Similarly, metadata searches often misidentify non-infringing work, since two pieces of different content can share the same metadata (two songs with the same title, for example).

- Hash-Based Filtering: Using a file as the input in a specialized mathematical algorithm called a hash function produces a unique alphanumeric code (a "hash") that can be used to identify specific files. Automated programs can compare the hash of a reference file containing a copyrighted work against a database of hashes for different files hosted on

an OSP to identify copies of the target work. But, because any alteration in a file's data (such as changing the encoding format, shortening a song by a fraction of a second, or compressing the file) will produce an entirely different hash, hash-based filtering can only identify exact matches of specific files but cannot identify modified files containing the same underlying copyrighted work. Like metadata searches, hash-based filtering tools are also easy to circumvent through file manipulation and encryption.

- Content Fingerprinting: Unlike metadata search tools (which look at the metadata of particular files) and hash-based filtering tools (which turn the bits comprising a file into a unique identifier), fingerprinting tools examine characteristics of the underlying media content—such as the frequency values in a song file—to make identifications. While these tools are more robust to alterations in the file itself than hash or metadata-based techniques, they are similarly limited in their capacity to identify content. Fingerprinting tools are narrowly tailored to particular media types (an audio fingerprinting tool cannot be used to match copyrighted text files), and such tools only exist for a small subset of the many types of copyrighted content available online. Because fingerprinting tools require access to the underlying media content, they cannot be used to process encrypted files or torrent files, and they are not consistent with search platform functionality.

Critically, all content filtering technologies are at best capable of simply identifying the contents of a file, not making the often complex determination as to whether the use of a particular file constitutes an infringement. Furthermore, no filtering technology can process encrypted files.

For the great majority of OSPs that do not host large volumes of infringing content, the cost of filtering technologies far exceeds their benefit in limiting infringement. Such tools can require significant human and financial capital to maintain, and most of the websites trafficking in infringing material are based overseas or are otherwise beyond the reach of U.S. copyright law. Ultimately, conditioning access to the safe harbor on an OSP implementing a filtering technology would undermine the certainty that the DMCA has created without a commensurate impact on infringement. Policymakers should resist calls to upend the DMCA on the illusory hopes that filtering technologies can replace the commonsense copyright regime that has allowed the internet to flourish.

I. Introduction

In the span of a generation, the internet has gone from relative obscurity to near universality. Almost 90 percent of Americans directly use the internet, and essentially all aspects of the modern economy depend on internet access.[1] Given this fast pace of innovation, it is something of a surprise that many of the laws governing the internet that were created when the Web was emerging are still effective today.

One of the most critical of these laws is the Digital Millennium Copyright Act (DMCA). Passed with almost unanimous, bipartisan support in 1998, the DMCA provides rules clarifying the circumstances under which online companies may be held liable for the copyright infringements of their users, encouraging cooperation between websites and rightsholders to address online infringement and driving investment in the sector. Congress determined that these goals were best met by placing the burden of identifying infringing material on copyright holders, not online service providers (OSPs). At its core, the DMCA establishes that if a website removes access to infringing content when a copyright owner provides notice of the infringement, it cannot be held monetarily liable for the infringement.

The DMCA's limitation on liability and division of responsibility between copyright owners and websites spurred the internet's boundless growth. But, in recent years, copyright owners have argued that the internet's rapid expansion has led to an equivalent increase in online infringement, such that the DMCA's division of responsibilities between OSPs and copyright owners no longer makes sense. According to this view, websites should be required not only to remove access to a particular infringing file upon request but also to (1) proactively identify and delete other copies of the copyrighted work at issue and (2) guarantee that no future copies of that work ever appear their sites again. Arguments for this shift from a "notice-and-takedown" standard to "notice-and-staydown" depend on the notion that new technologies make it easy for OSPs to automatically and accurately

[1] Monica Anderson and Andrew Perrin, *13% of Americans Don't Use The Internet. Who Are They?*, PEW RESEARCH (Sept. 7, 2016), http://www.pewresearch.org/fact-tank/2016/09/07/some-americans-dont-use-the-internet-who-are-they/.

locate and remove infringing content. Proponents of "notice-and-staydown" are advocating for laws that require websites hosting user-posted content to implement such "filtering" technologies, rather than treating these automated technologies as voluntary measures, as they currently are under the DMCA.

Although filtering technologies have some role to play in the online ecosystem to identify and remove infringing material, their many inherent limitations make them incapable of fully addressing copyright infringement. For many websites, the costs and problems associated with filtering technologies far outweigh their benefits for reducing infringement. All of the filtering technologies available today, from hash-based filtering to fingerprinting, are limited in their capacity to accurately identify content. Simple systems such as hash-based filtering and metadata searches rely on crude approximations of the contents of a given file, making them easy to circumvent and prone to false identifications. More complex "fingerprinting" tools that examine the underlying contents of a file are not much more effective, as they also suffer from regular false positives and can at best be used to identify only a small subset of the many different types of copyrightable media that OSPs host. Although there are fingerprinting tools available to scan and compare audio, video, and image files, no such tools exist to process other forms of copyrightable content, such as software. And, since these filtering tools require access to the complete, raw, unencrypted content of files, they cannot process encrypted files or torrents. As a result, the range of infringing activity that filtering tools can effectively address is rather narrow. And, even for media types for which filtering tools exist, such tools are only capable of matching content, not determining whether the use of a particular work constitutes an infringement.

With increasing calls to modify the DMCA to require OSPs to proactively filter content, it is useful to consider how filtering tools actually work to shed light on whether they are likely to cause more harm than good for most OSPs, users, and creators. This paper examines the technical functioning of common filtering tools and their effects on infringement and the internet ecosystem to highlight how the DMCA's approach to combatting online infringement is, despite its age, still well-suited to the technological realities of the internet and the goal of promoting the growth of both creative production and internet services at large.

II. Legal Background

A. The Origins of the DMCA

At the dawn of the internet era, it became clear that one of the internet's central virtues—its capacity to facilitate the almost instantaneous distribution of information of all kinds, including copyrighted material, to any connected point on the globe—presented complications for existing copyright law. Congress recognized that the decentralized nature of the internet meant that it would be essentially impossible for OSPs to know about, much less review and control, all infringing activities occurring on their systems. Existing principles of copyright infringement liability—both direct and secondary—seemed an awkward fit for this new medium. Early decisions on copyright liability for OSPs applied different tests and came to often conflicting conclusions.[2] Should OSPs be held directly liable for the infringing activities of their users? Does the operator of an intangible, unmoderated service like an online bulletin board exercise the same ability to control infringing activity as the owner of a dance hall[3] or swap meet?[4] Is the control an OSP exercises over its users more like the control a landlord exercises over her tenants or like a radio station over its disc jockeys?[5] Congress understood that allowing courts to make these determinations on an ad hoc basis would preclude an OSP from knowing in advance whether or not it could be held liable for copyright infringement, severely discouraging investment in the sector.

To help mitigate this uncertainty while at the same time encouraging OSPs and rightsholders to cooperate in limiting online copyright infringement, Congress passed the DMCA, which at its core creates a legal "safe harbor" that allows compliant OSPs to avoid monetary liability for secondary copyright infringement

[2] *Compare Sega Enterprises v. MAPHIA*, 948 F. Supp. 923 (N.D. Cal. 1996) *with Playboy Enterprises v. Frena*, 839 F. Supp. 1552 (M.D. Fla. 1993).
[3] *Dreamland Ball Room v. Shapiro, Bernstein & Co.*, 36 F. 2d 354 (7th Cir. 1929).
[4] *Fonovisa v. Cherry Auction*, 76 F. 3d 259 (9th Cir. 1996).
[5] *Religious Technology Center v. Netcom On-Line Communications Services, Inc.*, 907 F. Supp. 1361 (N.D. Cal. 1995).

claims.[6] Essentially, as long as an OSP removes access to allegedly infringing material that it becomes aware of and terminates users who repeatedly infringe copyrights,[7] it cannot be held monetarily liable for the copyright infringements of its users.

Since the obligations and protections of the DMCA's safe harbor depend on whether or not an OSP has knowledge of specific infringing activity, defining the circumstances when an OSP should be held to have such knowledge is central to the operation of the statute. Critically, Congress determined that an OSP should not be assumed to have knowledge of all user activities and established that eligibility for the safe harbor does not require an OSP to "monitor[] its service or affirmatively seek[] facts indicating infringing activity."[8] Rather, the "DMCA notification procedures place the burden of policing copyright infringement—identifying the potentially infringing material and adequately documenting infringement—squarely on the owners of the copyright."[9] As such, to qualify for the safe harbor, an OSP must prevent access to infringing material if it independently becomes aware of such material or if a copyright owner or its agent sends a statutorily compliant notice of claimed infringement to the OSP, but it need not take proactive steps to seek out infringing material on its site.

B. Policy Rationale for Putting Burden of Policing on Owners

Congress decided to refrain from automatically imputing knowledge of user infringements on OSPs and to place the burden of policing infringement on copyright owners in large part because "[c]opyright holders know precisely what materials they own, and are thus better able to efficiently identify infringing copies than service providers…who cannot readily ascertain what material is copyrighted and what is not."[10] Given the informational asymmetry between OSPs and copyright owners, requiring OSPs to identify potential infringements would be

[6] *See, e.g.* 144 Cong. Rec. S11,889 (daily ed. Oct. 2, 1998) (statement of Sen. Hatch) (DMCA meant to provide OSPs with "more certainty … in order to attract the substantial investments necessary to continue the expansion and upgrading of the Internet.").
[7] An OSP that wants to obtain the protection of the DMCA's safe harbor must also meet several other requirements, such as designating an agent to receive takedown notices. 17 U.S.C. § 512.
[8] 17 U.S.C. 512(m).
[9] *Perfect 10, Inc. v. CCBill LLC*, 488 F.3d 1102, 1113 (9th Cir. 2007).
[10] *UMG Recordings, Inc. v. Shelter Capital Partners LLC (UMG II)*, 718 F.3d 1006, 1021-22 (9th Cir. 2013).

inefficient, impractical, and prohibitively expensive for many OSPs. As the Senate noted in its report accompanying the DMCA, OSPs—often fledgling startups—should not be expected "to determine whether [a] photograph was still protected by copyright or was in the public domain; if the photograph was still protected by copyright, whether the use was licensed; and if the use was not licensed, whether it was permitted under the fair use doctrine."[11] It is difficult to imagine any small OSP being able to thoroughly monitor all of the content on its site and make accurate infringement determinations on a case-by-case basis.

Because most startups are ill-equipped to perform this level of analysis and in any event typically do not have enough infringement occurring on their platforms to warrant the imposition of a stringent content monitoring obligation, Congress recognized that placing the burden of policing infringement on OSPs would simply result in some combination of severely curtailed investment in online startups and the systematic removal of non-infringing material. Because copyright holders can seek statutory damages awards well in excess of the actual harm caused by an act of infringement,[12] OSPs already have strong incentives to process takedown notices for material that is arguably non-infringing, as the cost of a statutory damages award usually far exceeds the benefit of resisting a questionable removal request. Had Congress determined that OSPs were obliged to proactively monitor for potentially infringing material, and that they could be held liable for infringements for which they had no knowledge, it is unlikely that any Web 2.0 companies would have found investors in light of the constant threat of unforeseen copyright liability. As a result, the many OSPs that allow users to distribute, share, and store content online would not exist today.

C. Filtering Technologies Under the DMCA

Underlying Congress's decision to relieve OSPs from the burden of monitoring user activities and searching for infringing material was the recognition that such a task was technologically infeasible. As then-Senator John Ashcroft noted in the lead-up to the DMCA's passage:

[11] S. Rep. No. 105- 190, at 48.
[12] 17 U.S.C. § 504(c).

> It would be impossible for any carrier to review all of the material; and we cannot create a legal obligation that is technologically impossible to satisfy. Clearly, the potential for copyright infringement is real—as real as the impossibility of requiring a service provider to monitor every communication, including e-mail, homepages, and chat rooms [for infringing activity].[13]

Although Congress thoroughly considered how best to encourage the cooperative development and deployment of technologies that would help minimize copyright infringement on the internet, it ultimately decided to exempt OSPs from an obligation to affirmatively monitor their services for possible infringements, whether through manual review or automated technologies.[14]

As Congress anticipated, despite the lack of an affirmative mandate to deploy technologies to identify potentially infringing content, many such technologies emerged in the years after the DMCA's passage, and courts were tasked with determining how the use of such programs might affect an OSP's obligations under the DMCA. Perhaps the most direct analysis of the interplay between filtering technologies and the DMCA's requirements for OSPs came from the Central District of California in *UMG Recordings, Inc. v. Veoh Networks, Inc.*[15] In *Veoh*, the court granted summary judgment against UMG on its secondary infringement claims against Veoh, a user-generated content (UGC) video platform, holding, *inter alia*, that Veoh's decision to adopt two different filtering technologies did not affect its ability to claim the protections of the DMCA's safe harbor. UMG argued that even though Veoh implemented a hash-based filtering[16] system early in its existence, its delay in adopting an allegedly more robust filtering system amounted to willful blindness of the infringing activity occurring on its system. UMG claimed that this delay should deprive Veoh of the safe harbor, as it was intentionally avoiding gaining knowledge of infringements.[17] The court rejected this argument on the grounds that the DMCA does not impose "an obligation on a service provider to implement filtering technology at all, let alone technology from the copyright

[13] 144 Cong. Rec. S8729 (daily ed. Sept. 3, 1997) (statement of Sen. Ashcroft).
[14] *See, e.g.* Salil K. Mehra & Marketa Trimble, *Secondary Liability, ISP Immunity, and Incumbent Entrenchment*, 62 Am J. Comp. L. 685, 690-691 (2014)(describing discussion at congressional hearing regarding deployment of technologies to combat copyright infringement).
[15] *UMG Recordings, Inc. v. Veoh Networks Inc. (UMG I)*, 665 F. Supp. 2d 1099, 1108 (C.D. Cal. 2009).
[16] *See* section III.A.2., *infra*.
[17] *UMG (I)*, 665 F. Supp. 2d. at 1111.

holder's preferred vendor or on the copyright holder's desired timeline."[18] Rather, the court noted that Veoh's voluntary adoption of filtering technologies demonstrated its good faith efforts to avoid or limit storage of infringing content."[19] Similarly, the court held that "Veoh's 'right and ability' to implement filtering software, standing alone or even along with Veoh's ability to control users' access" did not mean that Veoh had the "right and ability to control [infringing] activity" on the site such that it should be excluded from the safe harbor pursuant to 17 U.S.C. 512(c)(1)(B):[20]

> If courts were to find that the availability of superior filtering systems or the ability to search for potentially infringing files establishes – without more – that a service provider has "the right and ability to control" infringement, that would effectively require service providers to adopt specific filtering technology and perform regular searches. That, in turn, would impermissibly condition the applicability of section 512(c) on "a service provider monitoring its service or affirmatively seeking facts indicating infringing activity.[21]

As other courts have noted,[22] conditioning eligibility for the safe harbor on an OSP's decision to implement content filtering tools would directly violate Congress's intent:

> [The DMCA] is not intended to discourage the service provider from monitoring its service for infringing material. Courts should not conclude that the service provider loses eligibility for limitations on liability under section 512 solely because it engaged in a monitoring program.[23]

Although it is well-established that the DMCA neither requires an OSP to implement a filtering system to obtain the protection of the safe harbor nor punishes an OSP that voluntarily adopts such a system by imparting on the OSP disqualifying knowledge of infringing activity, an OSP's decision to implement or

[18] *Id.*
[19] *Id.* at 1112.
[20] *Id.* at 1113.
[21] *Id.*
[22] *See, e.g. Hendrickson v. eBay*, 165 F. Supp. 2d 1082 (C.D. Cal. 2001).
[23] H. Rep. 105-796, at 73 (Oct. 8, 1998).

forego filtering technologies has on occasion impacted a court's consideration of an OSP's intent for purposes of a secondary infringement claim. Though it was not strictly relevant to the outcome in *UMG Recordings, Inc. v. Veoh Networks, Inc.*, the court noted that Veoh's adoption of two different filtering systems was evidence of Veoh's "good faith efforts to avoid or limit storage of infringing content."[24] In *Metro-Goldwyn-Mayer, Inc. v. Grokster,* the Supreme Court cited the defendant OSP's failure to "develop filtering tools or other mechanisms to diminish the infringing activity using their software" as additional evidence of their "intentional facilitation of their users infringement," though it cabined this dicta by noting that "in the absence of other evidence of intent, a court would be unable to find contributory infringement liability merely based on a failure to take affirmative steps to prevent infringement."[25]

D. Cooperation Between OSPs and Copyright Owners

Though there has been relatively little judicial analysis of the role that filtering technologies play under the DMCA, the congressional record, statutory text, and legal interpretations all clearly establish that the use of filtering technologies is purely voluntary for OSPs and does not impose any barrier to accessing the safe harbor.[26] Despite some concerns at the time of the DMCA's passage that making the use of filtering technologies optional would not adequately incentivize OSPs to develop and adopt tools to address online infringements,[27] many such tools have been developed and implemented on web platforms, at least on larger OSPs where the greater volume of user activity makes these tools more useful. For all but the largest OSPs, the small scope of infringing material on their portals renders

[24] *UMG (I),* 665 F. Supp. 2d at 1112.

[25] *Metro-Goldwyn-Mayer, Inc. v. Grokster,* 545 US 913, 939; *but see Capitol Records, LLC v. Vimeo, LLC*, 972 F.Supp.2d 500, 534 (S.D.N.Y. 2013) ("Plaintiffs further contend that evidence of inducement may be found in Vimeo's failure to implement filtering technologies that could be used to locate infringing content. But...just because Vimeo can exercise control does not mean that it must.").

[26] This rule is constrained in theory by language in the statute establishing that an OSP seeking the safe harbor may be required to proactively monitor for user infringements "to the extent consistent with a standard technical measure." *See* 17 USC §512(m). This limitation is only theoretical, since there are no such "standard technical measures" in existence. The statutory definition of "standard technical measures"—limited to technologies "used *by copyright owners* to identify or protect copyrighted works" that have been "developed pursuant to a broad consensus of copyright owners and service providers in an open, fair, voluntary, multi-industry standards process" and "do not impose substantial costs on service providers"—is very narrow and appears incompatible with all past and present filtering technologies. 17 U.S.C. § 512(i)(2) (emphasis added); *See also* L. Gallo, *The Impossibility of "Standard Technical Measures" for UGC Websites,* 34 COLUM. J.L. & ARTS 283 (2011). Considering the severe negative impacts that a mandatory content policing regime would have on the online ecosystem, it is no surprise that Congress limited "standard technical measures" to those technologies that are so uncontroversial as to have near unanimous approval.

[27] *See* Mehra & Trimble, *supra* note 14, at 690-691.

filtering tools impractical and unnecessary.[28] As Congress hoped, content owners and OSPs have entered into a variety of cooperative arrangements to help mitigate online infringement (and even allow copyright owners to monetize infringing user uploads rather than simply removing them, as with YouTube's Content ID). While early examples of such cooperation (such as the Principles for User Generated Content Services ("UGC Principles") from 2007)[29] were predicated in part by OSPs' concerns regarding potential liability, many of the arrangements in place today often derive from cooperative business partnerships between OSPs and copyright owners.

E. The Success of Congress's Voluntary Approach to Content Filtering

Congress's decision to encourage voluntary cooperation between content owners and OSPs to combat online infringement rather than placing the burden of searching out infringements on OSPs has succeeded in diminishing legal uncertainty regarding OSP liability for user infringements, allowing rightsholders to efficiently address copyright infringement on the many OSP platforms that do not experience significant infringing user activity, and fostering collaboration between OSPs and rightsholders on sites with enough user traffic such that more proactive measures are sensible. More critically, it has furthered the fundamental purpose of copyright law by facilitating the growth of a vibrant ecosystem of online portals that have helped spur a boom in creative production since the DMCA's inception. The economic value of the worldwide entertainment industry grew almost 66 percent from 1998 to 2010, and U.S. households spent almost 15 percent more on entertainment as a percentage of income in 2008 than in 2000.[30] In terms of actual content creation, the years since the DMCA's passage have been amongst the most prolific in history.[31] OSPs have seen similar growth, as the internet sector

[28] Gallo, *supra* note 26, at 295 ("[A]ll major right holders and all major UGCs...employ some form of the fingerprinting technology.")

[29] The UGC Principles are a voluntary, non-binding agreement between content owners and OSPs providing that content owners will not sue signatory OSPs that, among other things, implement "effective content identification technology." *Principles for User Generated Content Services*, http://www.ugcprinciples.com (last visited March 20, 2017).

[30] Michael Masnick and Michael Ho, *The Sky is Rising: a Detailed Look at the State of the Entertainment Industry*, at 2, http://www.techdirt.com/skyisrising/ (last visited March 20, 2017); *see also*, Joshua P. Friedlander, *News and Notes on 2015 RIAA Shipment and Revenue Statistics,* https://www.riaa.com/wp-content/uploads/2016/03/RIAA-2015-Year-End-shipments-memo.pdf (last visited March 20, 2017) (U.S. music industry revenues increased to $7.0 billion in 2015 and have increased at wholesale value for five consecutive years).

[31] Between 2001 and 2010, the number of songs identified in the Gracenote database jumped from 11 million to over 100 million; between 2002 and 2010, the number of new books available on the market increased more than twelvefold. *See* Masnick *supra* note 30, at 3.

contributed nearly $1 trillion to the U.S. economy—nearly 6 percent of real GDP—in 2014.[32]

F. New Calls for Mandatory Filtering

Despite the success of the DMCA, some copyright holders and policymakers have pushed to rewrite the law, arguing that the internet has grown in a way that was not anticipated by the DMCA's drafters such that its provisions no longer make sense in the current online environment.[33] They claim that the incredible growth of the internet and the exponential increase in content available online renders the DMCA's "notice and takedown" process ineffective at addressing online infringements, as content owners cannot identify and request removal of infringing content as quickly as new infringing content is posted.[34] Content owners complain that because the DMCA only requires OSPs to disable access to specifically identified infringing files and imposes no obligation to independently locate and remove other copies of works identified in takedown requests, they must send a takedown notice for each copy of a particular work on any given website—a Sisyphean task, according to copyright holders.[35]

To address the DMCA's perceived shortcomings, many content owners have argued that OSPs should be obligated to implement content filtering technologies to pre-screen all user uploads as a requirement to qualify for the safe harbor.[36] Under this proposal, "websites will be immune from copyright liability—both direct and indirect—as long as they can show that they employed the best filtering technology available on the market at the time the alleged infringement occurred."[37] These proposals are often conjoined with so-called "notice-and-staydown" obligations under which a standard takedown notice from a copyright

[32] Stephen E. Siwek, *Measuring the U.S. Internet Sector*, at 5, http://internetassociation.org/wp-content/uploads/2015/12/Internet-Association-Measuring-the-US-Internet-Sector-12-10-15.pdf (last visited March 20, 2017).
[33] *See, e.g., Comments of Universal Music Group*, U.S. Copyright Office, Section 512 Study, at 9 (2016) ("Whatever its original intent, Congress simply could not have contemplated either the massive scope or the diverse methods of online infringement when it passed Section 512 in 1998").
[34] *See, e.g.,* Ernesto, *RIAA Wants Google to End Piracy "Whack-a-Mole,"* TorrentFreak (March 14, 2014), https://torrentfreak.com/riaa-wants-google-end-piracy-whack-mole-140314/.
[35] *See Comments of the Motion Picture Association of America*, U.S. Copyright Office, Section 512 Study, at 5 (2016) ("Because infringing content is uploaded on a mass scale, requiring copyright owners to provide individualized URLs in order to trigger service providers' obligation to act further ensures that the process will be repetitive, mechanical and ineffective.").
[36] *See UMG Comments*, *supra* note 32 ("UMG respectfully submits that a service provider wishing to take advantage of the Section 512 safe harbors must, at a minimum, implement effective content identification technologies.").
[37] Lital Helman & Gideon Parchomovsky, *The Best Available Technology Standard*, 111 Colum. L. Rev. 1194, 1197 (2011).

owner for a particular infringing file would trigger an obligation for the OSP to locate and remove every other instance of the identified work from the website and prevent future uploads of that work.[38] Since manually comparing all uploads against a database of prior takedown requests is logistically impractical, notice-and-staydown proposals tacitly include a mandatory content filtering obligation and would punish OSPs when those filtering tools fail to remove all infringing content.

Although the DMCA's voluntary approach to OSP content filtering remains in place domestically, policymakers in Europe recently proposed an update to the European Commission's copyright regime that includes a requirement for OSPs to implement content filtering technologies.[39] The European Commission's (EC) proposal is a significant departure from existing EU copyright law, which largely mirrored the DMCA's notice-and-takedown approach.[40] A leaked document outlining the EC's policy rationale for proposing a mandatory filtering obligation reflects a misunderstanding of the technological and economic realities of content filtering, implying that filtering technologies are capable of accurately identifying a wide range of content and are available to startups for negligible costs.[41]

Proposals from the EC and copyright industries that would require OSPs to implement filtering technology misunderstand the inherent technical limitations of these technologies, as well as the costs and consequences of mandating their use. While using certain filtering tools may be an effective way for some OSPs to address online infringements, their cost and limited utility militate against mandating universal adoption or conditioning qualification for safe harbor protection on their use. Given the technical shortcomings of content filtering tools, requiring OSPs to implement such technologies will likely have minimal effect on copyright infringement but seriously harm the OSPs that have helped make the internet the dominant cultural and economic medium of our time.

[38] *See Comments of Authors Guild, Inc.*, U.S. Copyright Office, Section 512 Study, at 12-14 (2016).
[39] Proposal for a Directive of The European Parliament and of the Council on Copyright in the Digital Single Market, COM(2016) 593, https://ec.europa.eu/transparency/regdoc/rep/1/2016/EN/1-2016-593-EN-F1-1.PDF.
[40] *See* Directive 2000/31/EC of the European Parliament and of the Council of 8 June 2000 on Certain Legal Aspects of Information Society Services, in Particular Electronic Commerce in the Internal Market.
[41] Commission Staff Working Document Impact Assessment on the Modernisation of EU Copyright Rules (draft), at 138, (2016) (available at http://statewatch.org/news/2016/aug/eu-com-copyright-draft.pdf).

III. Technical Analysis of Common Filtering Tools

As with all areas of computing, there have been significant advances in content filtering technologies since the adoption of the DMCA. The technologies available today can more quickly scan files to identify the underlying content with more adaptability to potential alterations in the file format or encoded media. However, all of the tools currently available to locate potentially infringing material are subject to severe limitations with respect to their accuracy and adaptability. Additionally, given the wide range of copyrightable content, there are many types of media content for which no filtering tools currently exist. To help explain why filtering technologies cannot realistically be expected to accurately identify all infringing content or otherwise eliminate online copyright infringement, it is helpful to examine the basic functionality of the most commonly employed filtering technologies and to describe their uses and limitations.

A. Overview of Existing Filtering Mechanisms

1. Content Metadata Search

One common technique for identifying content online, in use since at least the early 1990s with the advent of Gopher,[42] involves searching files—either manually or through the use of automated programs—according to data that accompany the actual content, or "metadata." The term "metadata" refers to the fact that the content is annotated with additional data as opposed to simply being described by the content itself; in other words, it is structured information about the associated media resource. Metadata may contain various information about a given file; in the case of an audio or video media file, for example, metadata may include the content's title, data, file size, length, encoding rate, and so forth. Metadata can take on many formats, but it is generally structured in a way that makes it easy to search. Given such data about associated media, a simple takedown strategy entails searching an OSP for offending content based on its metadata and removing any matches. For example, an automated script could search for content

[42] Farhad Anklesaria, et al. *The Internet Gopher protocol (a distributed document search and retrieval protocol).* No. RFC 1436 (1993).

based on metadata associated with the content, such as the document's title (e.g., a song title), author, or artist. In the case of YouTube, for example, the metadata include information such as the content's author (or owner), a title, a text description, a set of text-based keywords, the time at which the video was uploaded, the duration of the file, the name of the raw file, and various viewing statistics (e.g., the number of times the video has been viewed). In principle, one could perform a search on *any* of this metadata to retrieve files with associated matching metadata. Note that it is possible to fetch this metadata programmatically (e.g., with an automated script) without ever examining the contents of the file itself. Similar metadata exist for content that is hosted on other OSPs. Searching for content based on its metadata is simple and efficient, as it does not involve direct analysis of the content itself.

Metadata searches are efficient in that they allow for the quick analysis of a large volume of files without needing to actually download any files and can be used across a variety of media.[43] However, they are imprecise and often inaccurate. Clearly, metadata searches will be unable to identify content if the metadata of a particular file do not accurately describe its content. Content is often mislabeled, either accidentally or intentionally. Similarly, two different pieces of media content can have the same metadata—a book and a movie with the same title, for example.[44] As a result, a search that aims to identify offending content for removal simply based on metadata may be subject to both false positives (in cases where portions of the metadata, such as a title, match the title of infringing content, but the content itself does not match the infringing content) and false negatives (in cases where content that corresponds to infringing content is labeled with metadata that do not match the metadata for the content). For example, audio, video, or other content might be re-titled, re-encoded, shortened, or otherwise modified in ways that affect the content metadata and thus the ability to perform a metadata search to identify corresponding infringing content. In particular, converting a media file from one format to another can often alter or eradicate

[43] As described in more detail *infra*, some fingerprinting tools are limited to a particular media type. A tool designed to identify the contents of video files will not work on text, for example. And no such fingerprinting tools even exist for many types of copyrightable media.

[44] See, e.g., Ernesto, *"Warner Bros. Admits Sending Hotfile False Takedown Notices,"* TorrentFreak, (Nov. 10, 2009) https://torrentfreak.com/warner-bros-admits-sending-hotfile-false-takedown-requests-111109/ ("Hotfile pointed out that [Warner's] automated process resulted in the removal of many files that do not belong to Warner. The movie studio admits this and confirms that while searching for 'The Box (2009)' many unrelated titles were removed. 'Warner admits that its records indicate that URLs containing the phrases 'The Box That Changed Britain' and 'Cancer Step Outsider [sic] of the Box' were requested for takedown through use of the SRA tool.'")

metadata, making searches for metadata inaccurate or otherwise impractical. In short, the accuracy of metadata searches is limited by the accuracy and precision of the metadata associated with every file at issue. Because it is easy to alter the metadata of a file, and because a file's metadata are not necessarily unique markers for particular content, metadata searches are not reliable for accurately identifying potentially infringing content.

2. Hash-Based Identification

Another approach to automated identification of content involves representing a piece of content (e.g., an audio or video file) by a content *hash*. A *hash* is a numerical representation of a file that is significantly smaller than the original file. The hash value of a file is computed using a cryptographic hash function that takes the file as input. A file's hash nonetheless uniquely represents the contents of the file, so that files with different contents will generally have different associated hash values. Commonly used cryptographic hash functions have a property called collision-resistance that aims to guarantee that two distinct inputs (i.e., files) will not produce the same hash value. Moreover, even slight alterations to a file will generally result in completely different hash values. As such, except in very rare cases of hash collisions, hashes are effectively perfect identifiers of a particular file.

Additionally, because the hash of a file is generally much smaller than the file itself (a hash might be, for example, 256 bits—a small fraction of the overall size of the file, since video files can easily exceed 100 megabytes), it is much more convenient to analyze hashes than the full files themselves. This smaller size lends itself to convenience for automated search and identification as well, since it is much easier to maintain a reference database of hashes than full files. Rather than relying on the accuracy of the metadata or requiring close examination of the entire file, an automated takedown mechanism could compute the hash value of a piece of content and compare that hash value against a database of hash functions that correspond to copyrighted content. If the same mechanism were required to compare full files against a database of full files, the process would require considerably more time, storage, and computing power. Consider, for example, that an an audio file that is encoded at 128 kilobits per second requires approximately one megabyte of storage per minute of audio; therefore, a typical audio file will require approximately 5 megabytes of storage—more than 100 times the size of

the corresponding hash value. In addition to the storage costs associated with storing and comparing files, the computational cost would be similarly higher, as comparing the contents of a larger file requires proportionally more time to read and process the file.

As with searching for metadata, searching for infringing content by hash values is more efficient than techniques that require analyzing full files; yet, a simple hash-based comparison also has multiple drawbacks. Notably, altering the original file in any way—shortening or excerpting it, re-encoding it, and so forth—will invariably alter the hash value. For example, if the same song is stored in two different file formats, one as a WAV and another as an MP3, each file will have a unique hash and the database of hash values against which a search is referenced will need to contain the hash values of both files to accurately identify both of them. A search for the hash of the MP3 version will not match the WAV file, even though the original data in each file were the same. As a result, just as altering a file's metadata will undermine the accuracy of metadata-based content filtering techniques, infringing content may escape automated detection simply by virtue of being a slightly altered version of the original content.

3. Audio and Video Fingerprinting

Newer, more sophisticated tools for content identification, such as Audible Magic, PhotoDNS, Zefr, and Echoprint, involve a technique known as *fingerprinting*. Unlike hash-based filtering, these programs typically do not examine the bits in the media file (which may change with alteration and editing) to make identifications, but rather characteristics of the media itself—the notes in a song, for example, rather than the bits in the file that encode the song. By looking at the characteristics of the media itself rather than the bits comprising the entire file (as with hashing), a fingerprint computes a representation of the content that is more robust, even under various previously described alterations, edits, and modifications that would change the raw content of the media. More specifically, each note in a song could be represented by the presence or absence of specific frequency values; the volume of a particular note is roughly represented by some amplitude at the frequency corresponding to that musical note. Specific fingerprints may include information such as the amplitude of different frequency values over a sequence of time intervals. For example, an audio file may have more "energy" in higher or

lower frequencies (a high-pitched tone has more energy in higher frequencies). The amount of energy in each part of the frequency spectrum will differ across files, and that representation may be both unique to the file *and* robust to various transformations of the original files. One piece of content can often have multiple fingerprints: algorithms can produce fingerprints based on different excerpts of a file, for example, and different algorithms for producing fingerprints will naturally produce different fingerprints.

Because fingerprinting technologies rely on algorithms that process the underlying media content of a given file, they are naturally constrained to a small subset of copyrightable content. For example, because an audio fingerprinting tool's central algorithm examines, say, the frequency values in a song file, it cannot be used to identify copyrighted photographs or software programs, which contain no audio frequency values. As such, to filter all files on a given site, an OSP would need to obtain a different fingerprinting tool for each type of media that is (or could be) hosted across the entire site. Considering the incredibly wide scope of copyrightable content, there are many types of content for which no fingerprinting tool exists, such as architectural designs or handmade items sold on Etsy.

For types of media that are amenable to fingerprinting, there are various algorithms that can produce such fingerprints; early versions of the tools to generate fingerprints for audio files, for example, generated fingerprints using the chroma vectors[45] of successive smaller segments of the original file. In short, chroma vectors measure the relative intensities of the twelve pitches in the music scale within a given time interval (for example, how often the tone of "C#" is heard during a time interval and at what intensity). These vectors tend to be robust to noise such as ambient sounds, and they are also independent of volume and dynamics. Although relying on chroma vectors to produce an audio file fingerprint represents an improvement over previous hash-based matching mechanisms, the original fingerprinting techniques were not robust to distortions that may be introduced in over-the-air recordings; subsequent improvements to the original algorithm have enabled the encodings to be robust to these distortions.

[45] Meinard Müller, Frank Kurth, and Michael Clausen, *Audio Matching via Chroma-Based Statistical Features*, Proceedings of the International Conference on Music Information Retrieval (2005), at 288–295.

Although audio fingerprinting tools overcome some of the limitations of tools based on hash or metadata searches—such as an improved capacity to accurately identify files that may have been altered in simple ways—these techniques remain at best an imperfect mechanism for automatically identifying copyrighted content. For example, fingerprinting techniques can be susceptible to false positives or false negatives when scanning files for which the encoded content has been altered in some way, and they are usually narrowly tailored to operate only on particular types of media; for example, a fingerprinting tool created to identify audio files cannot be used to identify software programs.

Most fingerprinting tools are proprietary, making it difficult to evaluate their technical functionality. Nevertheless, an examination of the mechanisms that underpin an open-source audio fingerprinting tool, Echoprint, is a representative illustration of how other fingerprinting programs operate. Although many such tools exist, they operate similarly to Echoprint and are thus subject to similar limitations.

B. Case Study: Echoprint

Echoprint is an open-source audio fingerprinting tool that is used by Spotify, among others. Because the tool is open-source, it is possible to analyze the tool in more detail than other contemporary fingerprinting tools, but other programs most likely rely on similar mechanisms. The main fingerprinting mechanisms described in this section are viewable in the Echoprint source code on GitHub.[46] The designers of Echoprint have also released a white paper that describes the fingerprinting algorithm in detail,[47] and many of the techniques are also summarized in a patent.[48]

Echoprint has several components: (1) a tool to generate the "codes" or fingerprints for a particular media file; (2) a query server that stores the database of codes to match against; and (3) the codes themselves that are used to match against the codes that are generated from any particular media file. Given any media file or excerpt, the code generator can produce a code that is subsequently matched against codes stored at the query server.

[46] https://github.com/spotify/echoprint-codegen/blob/master/src/Fingerprint.cxx.
[47] http://ismir2011.ismir.net/latebreaking/LB-7.pdf.
[48] Daniel Ellis and Brian Whitman, *Musical fingerprinting based on onset intervals*. U.S. Patent No. 8,586,847. (Nov. 19, 2013).

The process Echoprint uses to generate a fingerprint relies on the relative timing of successive onsets (i.e., beat-like events) in the audio sample. Recently, these fingerprinting algorithms have also added "whitening" filters to remove resonance (i.e., echo) that may result from various artifacts that live recordings or over-the-air broadcasts may introduce. Although such a whitening filter makes fingerprinting more effective for accurately matching both live and over-the-air recordings with an original recording, this filtering process does not necessarily correct for more "natural" distortions that may result from a live performance (e.g., changes in beat onsets that may simply be an aspect of the performance itself, as opposed to signal distortion).

After applying a sequence of filters to the audio signal, the codes that Echoprint generates are based on analysis from eight separate frequency bands. One configuration uses frequency bands spanning from 0 to 5,500 Hz (i.e., samples per second), but other frequency bands may be used; the method does not depend on a precise specification of frequency bands. To identify matching content, the fingerprinting mechanism compares the relative timing between beat onsets in the audio that are present in each of these different frequency bands. An "onset" is the period of increased magnitude of the music sample, such as the start of a musical note or a beat; given onsets in each frequency band, the algorithm may associate a timestamp within each onset. The resulting fingerprint or "code" is based on the relative location of onsets within the audio sample.

Given this fingerprint, the algorithm identifies matching documents by comparing the fingerprint against a database of known tracks. During the fingerprinting process, each audio file is split into 60-second segments, with each adjacent segment overlapping by 30 seconds to produce more matching query segments for an audio file of a given length. The codes for any 60-second segment are represented as terms in an *inverted index*. An inverted index takes an audio sample and returns the corresponding document that contains that sample, as opposed to a conventional index, which typically takes a track or document ID and returns the corresponding file. The concept is analogous to a search index, whereby sets of keywords map to corresponding documents that contain those keywords. The track index and the segment number is the document ID; the underlying database facilitates a fast lookup of the document ID, given a code that is being

queried. Any 30-second query contains about 800 samples that are used to perform the inverted lookup, and the query server returns the documents with the most matches for each code term in the query. In practice, many documents may contain a set of matching hash values; each document has a corresponding score based on the number of matching hash keys across the segments in the file. The algorithm only returns a matching document (i.e., media file) if the score of a single document is significantly higher than all other candidates. Otherwise, the algorithm does not return a match.

1. Accuracy of Echoprint in Controlled Experiments

Echoprint's designers performed an evaluation based on a database containing 30 million audio files. Depending on the length of the file and the type of encoding, the actual error rate for the algorithm varies between 1–2 percent, where an error could either be a false positive (the wrong file was identified), a false negative (the file was contained in the database but not identified), or a false accept (the file was not contained in the database at all, but a false match was detected).[49] Given the reported error rates, one could thus expect the state of the art fingerprinting algorithm to misidentify about one or two in every 100 pieces of audio content. As a point of comparison for acceptable false positive rates, consider that mail service providers consider any false positive rate higher than about 0.1 percent too high to be used in practice for spam filters, due to the potential limitations on speech that could arise as a result of legitimate email messages being misidentified as spam.[50] Accordingly, a 1–2 percent false positive rate for an automated filtering procedure is problematic for the same reasons, as such a technique would result in filtering legitimate content at rates that would frequently obstruct speech.

2. Echoprint's Limitations

Even the most sophisticated audio fingerprinting technologies such as Echoprint have several fundamental limitations. Because Echoprint is based on the inter-onset intervals in the audio file, any distortion to these intervals will disrupt Echoprint's ability to identify the corresponding audio file. Such distortions might be

[49] Daniel PW Ellis, Brian Whitman, and Alastair Porter. "Echoprint: An open music identification service." *Proc. ISMIR*. 2011.
[50] Shuang Hao, et al., *Detecting Spammers with SNARE: Spatio-temporal Network-level Automatic Reputation Engine*, USENIX Security Symposium. Vol. 9. (2009).

commonplace in live performances of copyrighted content, but other transformations may also distort these intervals. One transformation that could distort these intervals might be excerpting certain tracks of a recording, or altering the playback rate of the audio.

Also, Echoprint, like other fingerprinting tools, is designed to operate on a single media type: audio. As such, it is limited in its efficacy, as an OSP cannot use Echoprint or other audio filtering technologies to accurately identify video, text, or binary executable files like software programs. Fingerprinting a binary executable, in particular, is not currently feasible, because binaries are not amenable to the signal processing techniques that Echoprint and other media fingerprinting technologies rely on. Any OSP that allows users to upload a variety of different media would likely have to implement a different fingerprinting tool for each different type of content; even then, it is unlikely that the OSP would be able to filter all of the content on the site.

In addition to its inability to identify files with audio distortions and its applicability to audio files only, Echoprint, like other fingerprinting tools, is subject to many other fundamental limitations that *any* filtering technology faces.

C. Fundamental Problems with Filtering Technology

Filtering technologies such as those described above have a variety of limitations, many of which make it fundamentally difficult to perform automated identification of content. We briefly outline these limitations below and explain why relying on automated filtering technologies for copyright enforcement may result in false positives and other unintended outcomes.

First, because all filtering technologies rely on examining and analyzing aspects of target files (whether the bits in a file as with hash-based filtering, identifying information describing a file as with metadata searches, or unique aspects of the media content of a file as with fingerprinting), all filtering tools can be evaded through basic manipulation of the file. In the case of identification based on metadata, even simple manipulation of the metadata (such as altering the title of the content) is sufficient to disrupt automated filtering. In the case of file content

hashes, the value of the resulting hash is based on a computation over the contents of the file; thus, even small changes to a file's contents will result in a completely different hash computation. Thus, filters that rely on matches against hash values that correspond to infringing content can be undermined simply from small manipulations (such as shortening an audio file by a fraction of a second). Transformations and encryption will also undermine more sophisticated fingerprinting technology. Fingerprinting mechanisms that are based on aspects of the file's contents (e.g., chroma vectors in the case of audio fingerprinting and techniques such as those used by Echoprint) are typically more resilient to minor manipulation because these types of features are preserved. Encryption, however, stymies even these more sophisticated fingerprinting techniques. When a file is encrypted, the entire contents of the message are completely obscured, making it impossible to analyze the contents of the file for specific patterns. Although it is possible to encrypt a file's contents without encrypting the metadata, such a method would be subject to the previously discussed limitations of metadata-based fingerprinting—specifically, the file's contents may not match the associated metadata, and a file could thus be identified based on incorrect metadata. If the file's contents were encrypted, it would also be difficult to determine whether the contents matched the metadata, since the content itself could not be analyzed or inspected. Because all of these alteration techniques aimed at evading filtering technologies are easy to perform and in most cases do not require changes to a file that would render the encoded media less functional, filtering technologies are unlikely to be effective for mitigating copyright infringement.

Second, even if it were possible to consistently identify content accurately with filtering technologies, such technologies are not sufficient to consistently identify *infringements* with accuracy, as they can only indicate whether a file's contents match protected content, not whether a particular use of an identified file is an infringement in light of the *context* within which the media was being used. It is often permissible to excerpt or otherwise refer to copyrighted content in contexts that are permitted by fair use (e.g., for educational purposes, in art). Although an automated algorithm could determine whether the content (or excerpt) matched known copyrighted content, such an algorithm would not be able to determine whether the particular use of a given file is infringing or not. Since making nuanced determinations as to the legality of a file's use is appropriately left to legal

practitioners and juries, automated filtering technologies will never be sufficient to accurately police for infringing content.

Third, filtering technologies are only effective if an OSP has access to the file in question, making it difficult or impossible to use filtering tools on search platforms. Unlike OSPs that store material at the direction of a user (such as UGC sites), search engine platforms link to content that is hosted elsewhere often without downloading or otherwise processing the linked content. Although search engines build indexes based on content that they "crawl" from the web, ultimately these search engines map keywords and search phrases to links that are elsewhere on the internet, making it difficult, if not impossible, to control links to infringing content if the contents of a link were to change after the index was built. Similarly, search platforms cannot employ tools like hash-based filtering or fingerprinting, as they do not generally have access to linked files as would be necessary to compute a file's hash value or fingerprint. If search platforms were required to actually download linked content and run the associated files through filtering technologies, the cost of indexing the content on the web would increase significantly, potentially crippling the efficiency of online search.

In some cases, it may be extremely difficult for the search engine to obtain the content in the first place. For example, protocols that facilitate file sharing such as BitTorrent rely on application-specific clients to locate, download, and assemble file segments (i.e., "chunks" in BitTorrent parlance) based on information contained in an index file (in BitTorrent, a so-called "torrent" file).[51] Although the torrent file may be accessible via the web (and thus indexed by the search engine), the *contents* of the file that the torrent file refers to are impossible to retrieve without participating in the file-sharing protocol. As a result, search engines typically do not ever retrieve or store the contents of files shared using file-sharing software such as BitTorrent, even though the search engine may index the torrent files. In other cases where media might be streamed, the search engine would be unable to download the media in the first place, requiring any fingerprinting process to operate in real-time over a streaming audio file, which is impractical.

Similarly, the contents of any individual webpage are dynamic (i.e., generated programmatically with scripting languages such as Javascript), and are often even

[51] Bram Cohen, *Incentives Build Robustness in BitTorrent*, Workshop on Economics of Peer-to-Peer Systems, Vol. 6. 2003.

personalized based on the user or device that is retrieving the content. Website owners have at times attempted to conceal copyright infringing activity, malware hosting, or other malfeasance by exploiting such dynamic behavior to conceal the infringing or offending content from the search engine itself.[52] Ultimately, these limitations make filtering technologies ineffective for use in linking or other search-based applications.

Fourth, all automated filtering technologies depend on having a reference database for copyrighted content (such as matching fingerprints or hashes), but such a database is only possible if copyright owners submit accurate fingerprints/hashes of content to the database in the first place. Reported problems with the accuracy of takedown notices raise questions as to whether the databases on which filtering tools depend are or will soon become suffused with false identifications where the reference hash or fingerprint does not actually correspond to the content it purportedly identifies. One study found that almost 5 percent of takedown notices in the chosen sample did not match the infringed work.[53] In the context of the study, this amounted to around 4.5 million takedown requests.[54] Similarly, many copyright owners simply refuse to provide any fingerprint data to some OSPs, rendering futile their efforts to use filtering technologies to mitigate infringement. One file-hosting site, 4shared, was purportedly unable to obtain fingerprint data from large rightsholder organizations for use in its implementation of the Echoprint system.[55]

Finally, audio and video standards continue to evolve rapidly. For example, the H.264/MPEG-4 video coding standard has had 22 different versions since its inception in 2003;[56] many of these version updates involve substantial updates to the coding function and format that enable more scalable video coding algorithms and faster processing rates. Similarly, the JPEG image compression format specification has had six different updates.[57]

---

[52] David Y Wang, Stefan Savage, and Geoffrey M. Voelker, *Cloak and dagger: dynamics of web search cloaking*, Proceedings of the 18th ACM Conference on Computer and Communications ACM, (2011).
[53] Jennifer Urban, et al., *Notice and Takedown in Everyday Practice*, at 88 (available at http://ssrn.com/abstract=2755628).
[54] *Id.*
[55] *See, e.g.,* Ernesto, *4shared's Piracy 'Fingerprint' Tool Helps to Reduce Takedown Notices*, TorrentFreak (Nov. 23, 2016), https://torrentfreak.com/4shared-copyright-holders-abuse-googles-dmca-takedown-system-161123/ ("While 4shared has been using the content recognition software for quite a while already, not all copyright holders are eager to use it. Several large industry groups such as IFPI refuse to provide 4shared with fingerprint data.").
[56] https://en.wikipedia.org/wiki/H.264/MPEG-4_AVC
[57] Gregory K. Wallace, *The JPEG Still Picture Compression Standard*, IEEE Transactions on Consumer 38.1 (1992), at xviii-xxxiv.

Audio and video streams can also be transcoded from one format to another, or from one bitrate to another. This process often distorts aspects of the video encoding that might be used for fingerprinting. For example, in the case of video compression, a common technique is to perform a process akin to static image compression on an "anchor frame," with subsequent frames in the video encoded as differences from the anchor frame with motion vectors. This encoding process has slight differences between codecs; for example, the MPEG-4 standard defines forward motion vectors based on frames that are subsequent to an anchor frame; H.263, on the other hand, encodes motion vectors based on anchor frames that both precede and follow the frame being coded.

Transcoding between these formats fundamentally changes the nature of the encoding, given that the encoded motion vectors themselves will be different in each case.[58] For example, because an MPEG-encoded video will have only forward motion vectors and an H.263 video will have motion vectors that refer both forwards and backwards in time in the video stream, converting between the two formats will fundamentally change the nature of how the underlying compressed video stream is represented. Similarly, audio codecs use different methods to encode the underlying waveform. One such method, called *waveform encoding*, tries to represent an analog waveform as accurately as possible with its digital equivalent; pulse code modulation (PCM) is an example of such an encoding. Other audio codecs use a method known as *differential encoding* (also known as *predictive coding*), which predicts a subsequent audio sample based on a previous one and stores only the difference between the predicted audio sample and the actual sample. Differential coding significantly reduces the bandwidth required for transmitting an audio signal; for example, G.729, a differential encoding, uses about half of the number of bits as G.711, a waveform encoding. Naturally, each of these encoding mechanisms produces very different digital representations of the same audio sample.

In general, the evolution and proliferation of standards increases the likelihood the underlying representation of the compressed video stream will not be preserved across formats. If the database containing a fingerprint is based on characteristics that are present in one format but absent in another, then the fingerprint will fail to

[58] Susie J. Wee, John G. Apostolopoulos, and Nick Feamster, *Field-to-frame transcoding with spatial and temporal downsampling*, IMAGE PROCESSING (1999). *ICIP 99. Proceedings. 1999 International Conference on*. Vol. 4. IEEE, 1999.

produce a match on the content. As standards and formats proliferate, so do the mechanisms for transcoding between them, each of which may manipulate the underlying format in ways that disrupt the attributes of the file that are used for fingerprinting.

As encoding and compression standards continue to evolve, the fingerprinting technologies themselves must also develop to keep pace with standards. In the case of an audio fingerprinting mechanism such as Echoprint or other audio fingerprinting technologies, the inter-onset interval features may be preserved across encoding formats because they are based on properties of the audio. On the other hand, if the technique is based on properties of the *compressed* encoded video stream, as discussed in the video compression example above, the fingerprinting mechanisms will fail as encoding mechanisms improve and evolve. At the same time, standards must continue to evolve to support the proliferation of digital content on the Internet; given that video is a significant fraction of all Internet traffic at peak, innovation in coding standards will improve the efficiency of compression as well as the resilience to errors, packet loss, and (ever increasing) congestion. Keeping standards fixed solely to facilitate fingerprinting will impede innovation and is not feasible or sensible.

IV. Negative Implications of Filtering Requirement

The technical limitations of filtering tools are reason enough to question the reasonableness of policy proposals that would require OSPs to implement such technologies. Given the limited types of content these tools can scan and the ease with which they can be circumvented, it is unrealistic to believe that use of filtering systems will have the type of impact on infringement that proponents would like. Of course, the merits of requiring OSPs to use content filtering tools must be evaluated in terms of their cost as well as their efficacy. The negative impact a proactive filtering mandate would have on OSPs—and the concomitant decrease in creative output from those who rely on the internet as a medium of production and distribution—significantly outweighs its benefits.

A. Filtering Tools Are Prohibitively Expensive for Many Small OSPs

Considering the cost of filtering technologies in the most literal sense—the economic cost to OSPs—reveals problems with injecting a filtering obligation into a law meant to promote investment. Contrary to claims made by proponents of mandatory filtering, such tools are quite expensive—particularly for the startups that have historically driven the growth of the internet sector. To support its proposed filtering mandate, the European Commission argued that the cost of filtering tools would be negligible for startups: "it is estimated that a small scale online service provider can obtain such services for less than 900 euros a month."[59] This estimate derives from comments Audible Magic submitted in a Copyright Office study of the effectiveness of the DMCA's safe harbor provisions.[60] But this estimate is only accurate for an incredibly small number of OSPs. Audible Magic's website indicates that this price only applies to tools that filter audio files and is available only for OSPs hosting less than 5,000 song files per month—an incredibly low volume for an OSP.[61] To put this in perspective, when Soundcloud was only five years old, users were uploading twelve hours of audio content every minute.[62] The cost to most OSPs to source filtering technology from third parties is likely to be far more expensive in absolute terms than the EC's projection (assuming such technologies even exist to scan the types of content such OSPs host). A recent survey of OSPs reported that medium-sized companies engaged in file-hosting services paid between $10,000 and $25,000 a month in licensing fees alone for Audible Magic's filtering tool.[63] It is worth noting that the licensing fees for the software amount to only a portion of the total costs associated with using fingerprinting software. Any OSP's hosting platform must be altered or augmented to perform the fingerprint lookups and comparisons against a fingerprint database, a substantial software integration task.

[59] *Draft EC Report*, *Supra*, note 41 at 138.
[60] Comments of Audible Magic, U.S. Copyright Office, Section 512 Study, at 2 (2016).
[61] Audible Magic, *Copyright Compliance Service*, https://www.audiblemagic.com/compliance-service/#pricing (last visited March 20, 2017).
[62] Janko Roettgers, *SoundCloud Turns 5, Creators Now Upload 12 Hours of Audio Every Minute*, GIGAOM (Nov. 13, 2013), https://gigaom.com/2013/11/13/soundcloud-turns-5-creators-now-upload-12-hours-of-audio-every-minute.
[63] Urban, et al., *supra* note 53, at 64.

And, because not all OSPs that host "large amounts of content" (the undefined qualifying criteria for the EC's filtering mandate)[64] host large amounts of *infringing* content, requiring OSPs of a certain size to proactively filter user uploads will create massive costs for many OPSs without a commensurate benefit. For example, Kickstarter—a crowdfunding platform through which users can raise funds for creative projects—has hosted 338,368 projects on its site since its inception,[65] most of which contain a range of image, audio, and video content; but in 2015, it removed only 78 projects in response to copyright infringement notices targeting 215 distinct projects.[66] Because copyright infringements are exceedingly rare on its platform, the cost to Kickstarter of implementing filtering technologies that could process image, audio, and video content—likely separate tools for each media type—would far exceed the benefit of removing less than 100 infringements per year. And, since so many of the infringement notifications Kickstarter receives each year are false or deficient (around two-thirds of all infringement notices in 2015), it would likely still have to devote human resources to review the content that filtering technologies identified to prevent inappropriate removals.

This is not to say that filtering technologies are unaffordable for all OSPs. Indeed, many larger OSPs that allow users to share content already use filtering tools. YouTube famously spent $60 million developing its Content ID tool.[67] Soundcloud spent more than €5 million building its own filtering technology and still must dedicate seven full-time employees to maintain the technology.[68] But, for smaller OSPs, the cost of these systems would make it significantly harder to attract investors and compete with dominant incumbents. In a survey of investors in the U.S. and EU, a majority of respondents said they would be "uncomfortable investing in businesses that would be required by law to run a technological filter on user-uploaded content."[69] Since startups are responsible for virtually all new net job

[64] *Draft EC Report*, *Supra*, note 41, at 132.
[65] *Stats*, KICKSTARTER, http://www.kickstarter.com/help/stats (last visited March 20, 2017).
[66] *Kickstarter Transparency Report 2015*, KICKSTARTER (Apr. 25, 2016), https://www.kickstarter.com/blog/kickstarter-transparency-report-2015.
[67] Hearing on Section 512 of Title 17 before the H. Judiciary Subcomm. on Courts, Intellectual Prop., & the Internet, 113th Cong. 47 (2014), at 49 (testimony of Katherine Oyama).
[68] https://ec.europa.eu/eusurvey/pdf/answer/6acf2b21-865a-402c-876a-e2b67c0ceef9.
[69] Evan Engstrom et al., *The Impact of Internet Regulation on Early Stage Investment*, (2014) , https://static1.squarespace.com/static/571681753c44d835a440c8b5/t/572a35e0b6aa60fe011dec28/1462384101881/EngineFifthEraCopyrightReport.pdf.

growth in the U.S.,[70] the overall economic cost of a mandatory filtering regime for OSPs would likely be significant.

B. A Filtering Obligation Would Undermine the Certainty of the Safe Harbor

Even if the monetary and labor expense of employing content filtering technology alone had a minimal direct impact on OSP viability, conditioning the safe harbor on implementing a satisfactory content filtering tool would likely curtail investment in OSPs by creating uncertainty as to whether an OSP's chosen filtering technology would pass muster under the statute. A safe harbor is only useful for mitigating uncertainty if the prerequisites for obtaining its protections are clear, but it is decidedly unclear how policymakers could precisely define what sorts of filtering technologies would be deemed adequate to claim the safe harbor.[71] Any filtering requirement would either have to endorse a particular technology (and quickly become outdated) or establish a "reasonableness" standard that would require clarification from courts before OSPs could have any confidence in their protection under the safe harbor.

Given the fast pace of development, attempting to identify in statute which particular filtering technologies would qualify for the safe harbor would be a fool's errand. Since 2000, at least twelve new audio encoding formats have emerged,[72] each of which has a plethora of encoding and sampling rates that can introduce distortion or loss. Given the rate of innovation in encoding and compression technologies, as well as the rate at which fingerprinting technologies themselves continue to evolve, mandating the use of a specific filtering technology is likely to be ineffective. Audio and video compression technologies typically are subject to significant updates every couple of years; requiring the use of a specific technology would most likely result in the use of technologies that quickly become obsolete and increasingly inaccurate as the standards and technologies themselves evolve.

[70] Tim Kane, *The Importance of Startups in Job Creation and Job Destruction*, KAUFFMAN FOUNDATION, (2010), http://www.kauffman.org/~/media/kauffman_org/research%20reports%20and%20covers/2010/07/firm_formation_importance_of_startups.pdf.
[71] Some commenters have proposed an even more stringent standard, arguing that OSPs should be required to implement the "best available" filtering technology to qualify for the safe harbor. *See*, Helman & Parchomovsky, *supra* note 37.
[72] Cornell University Library, *Digital Preservation and Technology Timeline*, DIGITAL PRESERVATION MANAGEMENT, (2003).

Inevitably, a mandatory filtering regime would require a "reasonableness" standard for determining which filtering technologies are sufficient for an OSP to claim the safe harbor. Although "reasonableness" standards are common throughout the law—the DMCA itself conditions the safe harbor on an OSP "reasonably implement[ing]" a repeat infringer policy[73]—determining what filtering technologies might be considered "reasonable" is a particularly uncertain inquiry, considering the wide range of technologies that can be employed to filter content and the fast pace at which new technologies emerge and replace prior systems. While the "reasonableness" of an OSP's repeat infringer policy is likely to remain consistent over time, a filtering technology that is reasonable for an OSP one year may be legally inadequate the next.

Even amongst existing filtering technologies, there is little agreement about which systems are sufficiently effective. For example, while some copyright owners have commended YouTube for developing its Content ID fingerprinting system, others argue that even Content ID—a technology that cost YouTube $60 million to develop—is inadequate and should be supplemented by some other, more effective system.[74] Such complaints about the reasonableness and efficacy of filtering systems seem motivated in part by unrealistic expectations about what they can actually accomplish. According to Sony Music Entertainment, "if even a single copy of a particular work slips by a filter, in practical terms it is tantamount to a failure by the filter to lessen the availability of that work on the service at all."[75] Similarly, in the litigation over the infamous Napster peer-to-peer file sharing service, a judge held that Napster was in violation of an injunction because, despite filtering out 99.4 percent of infringing content, the court demanded the site use a filter that could remove *all* infringing material.[76] Of course, since no such filtering technology exists, Napster promptly shut down.

Even determining what criteria courts should use in evaluating the reasonableness of a particular filtering system will likely produce significant uncertainty and

[73] 17 U.S.C. § 512(i)(1)(A).
[74] *Compare*, *Comments of the Motion Picture Association of America*, *supra* note 35 ("voluntary cooperation between content owners and service providers could go a long way to solving this problem…[YouTube] uses a Content ID system that is very similar to those described in the *UGC Principles*.") *with Comments of Universal Music Group*, *supra* note 33 ("while YouTube's content management tools were essential, and have improved to a certain extent today, they still are not sufficient.").
[75] *Comments of Sony Music Group*, U.S. Copyright Office, Section 512 Study, at 12-13 (2016).
[76] Evan Hansen, *Court: Napster Filters Must Be Foolproof*, CNET (January 2, 2002) https://www.cnet.com/news/court-napster-filters-must-be-foolproof/.

potential conflicting judicial opinions: should the law favor systems that are capable of identifying a wide range of potentially infringing content at the risk of deleting a non-negligible amount of non-infringing content (such as metadata searches) or should it prioritize accuracy at the risk of underinclusiveness, as with hash-based filtering? Is a system that is technically state-of-the-art "reasonable" for purposes of a filtering mandate if copyright owners refuse to submit fingerprints for that system?[77]

Ultimately, since what is "reasonable" is generally a question of fact for a jury to decide,[78] an OSP would likely not know the answer to these questions—and consequently whether it could claim the safe harbor—until after a full jury trial, rendering the safe harbor of little value. For many OSPs, simply engaging in litigation over the adequacy of a filtering system would lead to bankruptcy, even with a favorable judgment on the merits.[79] The threat of protracted litigation would likely deter investment in startups seeking to compete with incumbent content distribution OSPs, since venture investors would have to consider whether a potential portfolio company could afford to implement a content filtering technology *and* defend a secondary infringement lawsuit. The likelihood of decreased venture activity is not merely speculative, as the negative investment consequences associated with increased litigation risk are well-documented.[80] In short, if the DMCA's safe harbor protections were predicated on an OSP implementing a "reasonable" content filtering system, the uncertainty surrounding any OSP's eligibility for the safe harbor would undermine one of the central purposes of the DMCA: providing "more certainty … in order to [allow OSPs to] attract the substantial investments necessary to continue the expansion and upgrading of the Internet."[81]

---

[77] *See, e.g.,* Ernesto, *4shared's Piracy 'Fingerprint' Tool Helps to Reduce Takedown Notices*, TORRENTFREAK (Nov. 23, 2016), https://torrentfreak.com/4shared-copyright-holders-abuse-googles-dmca-takedown-system-161123/ ("While 4shared has been using the content recognition software for quite a while already, not all copyright holders are eager to use it. Several large industry groups such as IFPI refuse to provide 4shared with fingerprint data.").

[78] *See Perfect 10, Inc. v. CCBill LLC,* 488 F.3d 1102 (9th Cir. 2007) (reasonableness of repeat infringer policy a question of fact); *see also Tran v. State Farm Mut. Auto. Ins. Co.*, 999 F. Supp. 1369, 1372 (D. Haw. 1998) ("An analysis of what is reasonable is almost always de facto a question for the jury.").

[79] For example, Veoh, a video hosting OSP, was forced to declare bankruptcy after its protracted and ultimately successful defense of a secondary infringement lawsuit brought by Universal Music Group. Eliot van Buskirk, *Veoh Files For Bankruptcy After Fending Off Infringement Charges*, WIRED (February 12, 2010) https://www.wired.com/2010/02/veoh-files-for-bankruptcy-after-fending-off-infringement-charges/.

[80] *See, e.g.* Michael A. Carrier, *Copyright and Innovation: The Untold Story,* 2012 WIS. L. REV. 891, 905–917 (2012) (describing decreased venture investment as a result of litigation uncertainty after Napster lawsuit); *see also* Stephen Kiebzak et al., *The Effect of Patent Litigation and Patent Assertion Entities on Entrepreneurial Activity*, 45 RES. POL'Y. 218 (2016).

[81] Hatch, *supra* note 6.

C. A Filtering Obligation Would Likely Have Minimal Impact on Infringement

Though the DMCA does not require OSPs to employ content filtering technologies, many of the larger OSPs that can afford to do so voluntarily use some kind of content filtering,[82] often because they are concerned that failing to do so could result in secondary liability (though as described in more detail *supra*, some OSPs have implemented filtering tools pursuant to cooperative agreements to help copyright owners monetize user posts).[83] As such, passing a law requiring OSPs to filter content would not result in any decreased infringing activity on many of the larger OSPs that already use filtering tools, such as YouTube, Soundcloud, and Spotify. And, even assuming that there are large numbers of OSPs that would implement filtering tools for the first time should such a requirement become law, and assuming that these tools were effective at identifying and removing infringing content, requiring U.S.-based OSPs to implement such tools to obtain the safe harbor would likely not have a major impact on copyright infringement, since most infringing activity appears to involve so-called "pirate sites" that are based overseas or are otherwise ignoring U.S. copyright laws. One recent study analyzing the interplay between online advertising and copyright infringement identified the top 589 "pirate" sites from 2014, virtually all of which are foreign domains.[84] Similarly, all of the sites identified in a report from *TorrentFreak* on the ten most popular torrent[85] sites of 2017 appear to be located outside the U.S.[86] While it is impossible to know with certainty where some websites are located, the Motion Picture Association of America's 2016 public commentary submission in an inquiry from the Office of the U.S. Trade Representative regarding "notorious markets" purports to identify the locations of many such sites, all of which are based outside

[82] There is some debate about the prevalence of content filtering technologies amongst OSPs. One survey of OSPs concluded that "most OSPs do not filter," while other commentators have suggested that certain forms of filtering are common, at least amongst larger OSPs. *Compare* Urban, et al., *supra* note 53 at 58 *with* Gallo, *supra* note 26 at 296 ("[F]ingerprinting technology is widespread among UGCs").

[83] *See* Urban, supra note 53 at 58 ("OSPs reported that filtering is likely to be adopted under considerable pressure and concern for liability.").

[84] Digital Citizens Alliance, *Good Money Still Going Bad: Digital Thieves and the Hijacking of the Online Ad Business* (2015) http://illusionofmore.com/wp-content/uploads/2015/05/latest-DigitalCitizensAlliance5.pdf. Moreover, of the top 596 "pirate" sites identified in the 2013 report, less than 60 percent made the list in 2014, indicating that these sites, regardless of domicile, are too transient to be impacted by anti-copyright infringement measures.

[85] Torrent technology is content-neutral and can be used to distribute non-infringing and infringing content alike.

[86] Ernesto, *Top 10 Most Popular Torrent Sites of 2017*, TORRENTFREAK (Jan. 7, 2017), https://torrentfreak.com/top-10-most-popular-torrent-sites-of-2017-170107/.

the U.S.[87] Requiring OSPs to filter user uploads would do nothing to stop infringement on these websites.

V. Conclusion

In the two decades since the DMCA was passed, internet activity and total creative output have both increased exponentially. Considering the DMCA was designed to mitigate legal uncertainty in order to drive investment to the internet sector and to encourage copyright owners and OSPs to cooperate to combat online infringement in order to promote creative production, its success seems unquestionable. But, despite this strong record, the DMCA has come under fire in recent years from copyright industries for allegedly not doing enough to stop copyright infringement. These calls to modify the DMCA to require OSPs to implement technical measures to police infringement overestimate the technical capacity of these tools and underestimate the damage that their obligatory use would have on the internet ecosystem. Before considering dangerous mandatory content filtering rules, policymakers should understand the inherent limitations of filtering technologies. Reversing two decades of sensible copyright policy to require OSPs to deploy tools that are costly, easily circumvented, and limited in scope would deeply harm startups, users, and content creators alike.

[87] Motion Picture Association of America Submission in Request for Public Comment on 2016 Special 301 Out of Cycle Review of Notorious Markets, Docket No. USTR-2016-2013, available at http://www.mpaa.org/wp-content/uploads/2016/10/MPAA-Notorious-Markets-2016-Final.pdf.

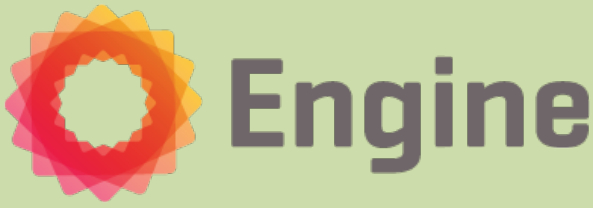

Engine (www.engine.is) is a non-profit organization that supports the growth of technology entrepreneurship through economic research, policy analysis, and advocacy on local and national issues.