Analog Reconversion Discussion Group
General Comments of Electronic Frontier Foundation

November 24, 2003

Seth Schoen, Staff Technologist
454 Shotwell Street
San Francisco, CA 94110
415-436-9333 x107

Executive Summary

- The "analog hole" question raises important public policy issues not discussed within ARDG.
- Technologies proposed to ARDG are likely inadequate to prevent copyright infringement, especially Internet-based copyright infringement.
- Both watermark and VBI-based technologies may be weak and easy to defeat; several technologies can be defeated using techniques that are already public.
- Assessing the efficacy of watermark technologies requires public peer review, but the public lacks enough information to make an informed judgment about watermark vendors' claims presented to ARDG.
- Deploying a VBI technology may create new risks for hearing-impaired viewers and others who rely on closed caption information.

Public Policy Issues

Public policy issues and the larger context surrounding the "analog hole" question were not discussed at ARDG. This practice leaves important issues unaddressed. These issues are of fundamental importance to the Constitutionally-mandated balance between copyright holder's interests and the public interest in U.S. copyright law.

For instance, questions to technology proponents specifically about effects on lawful uses of copyrighted works suggested by Public Knowledge and the Center for Democracy and Technology were removed from the ARDG analysis matrix before it was circulated. As a result, technology vendors were never invited to consider these important questions publicly.

One representative question concerns what happens when someone has the legal right to make a copy of an audiovisual work in a situation where a publisher, broadcaster, or other party has arranged for the work to be marked "Copy Never". Will "compliant" hardware preclude the lawful creation of a copy or excerpt of such work? What recourse will the would-be user have if the technology does forestall a legitimate use?

Such questions are particularly important because copyright holders have, in the past, trumpeted the existence of the "analog hole" as a means of protecting the public's rights to make copies for fair use and other purposes. When the effects of digital rights management technologies and legal anticircumvention measures on the public's traditional access rights were questioned in litigation and before the United States Copyright Office, entertainment interests pointed to the existence of the "analog hole" as a safety valve protecting the public interest. But this argument and these assurances will no longer hold water if the rights of the public to use and benefit from analog are substantially altered.

Should the "analog hole" be utterly stopped up, the public may finally be unable – as a technical matter – to make many lawful uses of audiovisual works by any means whatsoever.

General Inadequacy of Proposed Technologies

It is not at all clear that the proposed technologies will be technologically capable of blocking the analog hole – even if all of them performed as advertised.

It is difficult to see how the public's current level of access to video digitizing technologies can be taken away – "put back in the tube", so to speak. Even redesigning a substantial amount of hardware (with a system which we assume for the moment is very cleverly designed) would have only limited effectiveness. Enormous numbers of video digitizers with unencrypted outputs are already deployed, the market for such hardware is large, and the current cost of these devices is low, at least at standard-definition resolutions.

Some analog-to-digital conversion hardware is not specifically designed for video applications, but it may be possible to repurpose "non-video" devices so that they can usefully process video. The number of analog-to-digital conversion microchips in the world has long exceeded the planet's human population, and the ADC chips are still multiplying faster than we are.

Digitizing hardware can also be fabricated from scratch from other components.[1] While there is some dispute about the cost of creating an analog-to-digital converter with particular specifications, this technology is widely available, widely deployed, and widely understood. We have argued that hobbyists are likely to be able to design and create video digitizers. In addition, consumers will continue to be able to import digitizers from outside the U.S., where billions of dollars of goods sold every year contain analog-to-digital conversion capabilities.

Technological controls may also be ineffective for a different reason. If, as entertainment interests frequency contend, the Internet is an extremely rapid and efficient distribution system for illegal copies, then even a small number of sources could illegally make works widely available, so that restricting the average person's ability to obtain unrestricted copies of a work may have little or no effect on the availability of unlawful copies of works on-line.[2]

Limitations of Technical Approaches: Watermark-Based Technologies

Several proposed technologies submitted to ARDG are based on mandated detection of some digital watermark at the point of analog-to-digital conversion. For the reasons outlined in more detail below, there is a paucity of information about whether these technologies will work as advertised (and good reason to believe that they will not).

Watermarks, like other security systems, can only be evaluated scientifically when they are published for peer review. While no one has found a way to prove conclusively that a particular watermark is secure, it is possible for skilled reviewers to

[1] See, e.g., Andrew Huang, "Myths and Misconceptions About Hardware Hacking", presentation to Analog Reconversion Discussion Group, May 28, 2003, available at http://www.cptwg.org/Assets/Presentations/ARDG/ARDGHardware_hack05-28-03.pdf.

[2] See Peter Biddle, Paul England, Marcus Peinado, and Bryan Willman, "The Darknet and the Future of Content Distribution" (2002), available at http://crypto.stanford.edu/DRM2002/darknet5.doc.

determine whether the state of the art would allow the creation of a straightforward attack against that watermark in a particular application.

Digital watermarking is the subject of serious scientific analysis, and substantial scientific research. Much research about watermarking applications aims at evaluating systems by attempting to discover attacks against them. (Attacks against watermarks and watermark-detection systems may include, without limitation, devising technical means to remove or alter a watermark, or to hide a watermark from a detector so that the detector fails to detect the watermark correctly.)

In the past, many watermarks (even those developed by highly trained scientists and engineers) were quickly defeated by others after they were published. The results of this process fill the proceedings of mathematics and computer security conferences and adorn the pages of peer-reviewed journals. They emphasize just how fundamental public peer review is in assessing security; the creators of security systems are often poor judges of how those systems would hold up in the real world. The iterative nature of the open scientific process frequently leads to major improvements and new discoveries in security.[3]

The range of mathematical tools available to an attacker is significant. Yet dozens of watermarks have been defeated by extremely simple methods such as a tiny distortion of the picture, or a slight change in the speed of playback.

When the SDMI Forum, which was considering audio watermark technologies for copy-control applications, published limited information about some of the proposed technologies, researchers were soon able to point out weaknesses in many of them.[4] Some of the researchers studying the technologies believed that too little information had been published to allow a truly informed assessment.

Unfortunately, most of the techniques presented here have never been published or reviewed. The analysis matrices submitted by proponents generally indicate either that a technology has never been published or disclosed at all, or that it has been disclosed only to a few parties under a confidentiality agreement.

As a result, the public evidence for these technologies' efficacy is typically limited to vendor claims, with no independent analysis or verification. This is obviously a poor way of evaluating and selecting security systems. Even a well-intentioned security technology creator is rarely able to anticipate all of the sources of weakness in his or her own invention.[5] In some cases, various parties have evaluated certain technologies secretly under confidentiality agreements. (For example, the DVD-CCA watermark evaluation process allowed a few companies – but not the general public – to perform evaluations of proposed technologies.) By their nature, these confidentiality agreements

[3] See, e.g., Brief of *Amici Curiae* Dr. Steven Bellovin *et al.*, at 23-30, Universal City Studios v. Eric Corley, 273 F.3d 429 (2nd Cir. 2001) (No. 00-9185) (discussing importance of public peer review to development of cryptography and computer security). Similar points are made by several authors on security engineering; see, for example, Bruce Schneier, *Crypto-Gram*, May 15, 2002, available at http://www.schneier.com/crypto-gram-0205.html. Schneier was also a party to the Bellovin *et al.* brief.

[4] See Scott Craver *et al.*, "Reading Between the Lines: Lessons from the SDMI Challenge", Proceedings of the 10th USENIX Security Symposium (August 13-17, 2001), available at http://www.usenix.org/events/sec01/craver.pdf.

[5] See Bruce Schneier, *Crypto-Gram*, February 15, 1999, available at http://www.schneier.com/crypto-gram-9902.html (discussing inability of inventors to evaluate their own inventions' security properties, and the prevalence of inaccurate claims of security on the part of vendors).

preclude informed assessment by the public of the quality of the evaluation performed. They prevent us from repeating experiments and do not allow us to tell whether particular weaknesses have been considered, or how carefully. All we can tell is that unknown persons performed an unknown amount of unknown research, and later pronounced themselves satisfied. This provides an insufficient basis for public confidence.

By contrast, the NIST AES competition, by which the U.S. government selected the new Federal standard data encryption algorithm, was conducted as a public competition. Submitters had to publish the details of their proposals at the outset; other submitters (and scientists from around the world) then analyzed the submissions for weaknesses. In many cases, significant weaknesses in proposals were uncovered as a result of this process. Several government-sponsored conferences saw the submission of dozens of papers, highlighting many flaws that had not been apparent to the original inventors of AES candidate technologies. Eventually, NIST was able to select an encryption technique (Rijndael) as the AES standard, FIPS 197. Today, AES enjoys a high degree of public confidence because of the transparency of the process and the substantiality of the peer review of candidate technologies. NIST would not have accepted AES proposals from entities that refused to submit to a public review process or kept their technologies secret.[6]

Because few proposed watermark technologies submitted to ARDG have been subject to public analysis, there has been no opportunity to detect and publicize erroneous claims by technology proponents (as routinely took place during the AES selection process).

Princeton University watermark researcher Scott Craver, in a presentation to ARDG, noted that new watermark designs are frequently attacked successfully soon after publication, and that watermarks may have limitations making them unsuitable for use in copy-control applications. He concluded that, for controlling analog reconversion of audiovisual works, the "state of the art favors analysis" (i.e., the attempt to remove or obscure a watermark).[7] Targeted attacks against particular watermarks are often available and effective.

Moreover, regardless of its strength or security, the suitability of any digital watermark for restricting digitization at the point of analog-to-digital conversion in a personal computer is questionable at best. In a computer environment, many attacks try to conceal the presence of the watermark from the detector, rather than removing or

[6] See National Institute of Standards and Technology, "Announcing Request for Candidate Algorithm Nominations for the Advanced Encryption Standard (AES)", 62 Fed. Reg. 48051 (September 12, 1997) (explaining public review process and submission requirements including detailed technical disclosures); "Specification for the Advanced Encryption Standard (AES)", Federal Information Processing Standards Pub. 197 (November 26, 2001) (codifying encryption standard selected as a result of hat process); Schneier (explaining why public review process, including independent third party analysis, yielded a more secure AES with improved public confidence). For the AES conferences, at which flaws or potential flaws in candidate technologies were identified, see http://csrc.nist.gov/CryptoToolkit/aes/round1/conf1/aes1conf.htm, http://csrc.nist.gov/CryptoToolkit/aes/round1/conf2/aes2conf.htm, and http://csrc.nist.gov/CryptoToolkit/aes/round2/conf3/aes3conf.htm.

[7] Scott Craver, "What We Expect from Watermarking", presentation to Analog Reconversion Discussion Group, May 7, 2003, available at http://www.cptwg.org/Assets/Presentations/ARDG/watermarking5-7-03.ppt.

altering the watermark itself. These attacks may not even depend for their success on the technical details of the watermark itself; as the proverb has it, "all cats are gray at night".

This attack can be mounted easily against a watermark detector present at an analog input to a PC. In one version, the watermarked signal is scrambled while still in analog form, before providing it to a video input on the PC. The scrambling is performed using analog components according to a reversible scrambling method whose details are known to the attacker. When the scrambled video is digitized, the watermark detector will be unable to see the watermark hidden within the scrambled video signal. After digitization is complete, and the attacker has a digital copy of the scrambled video saved on the PC's hard drive, the attacker simply reverses the scrambling in software to obtain an unrestricted clear copy.[8]

Broadly speaking, this attack works because watermark designers try to create watermarks that cannot be removed without "destroying the perceptual quality of the signal" or "making the video unwatchable". If a method of removing or hiding the watermark makes the video unwatchable, watermark designers typically assume that nobody would find a reason to apply that method. However, analog encryption schemes deliberately destroy the quality of a signal or deliberately make it unwatchable in a predictable and reversible way. Thus, the signal scrambling can be reversed after digitization has already taken place – making the video watchable again.

Because it need not preserve the appearance or watchability of an image, analog scrambling can alter any feature of a video signal. It can even, like the World War II-era SIGSALY system, incorporate an analog key to control the scrambling process.[9] The output of the scrambling process will then appear to be random noise or static from the point of view of a watermark detector – but someone in possession of the scrambling key can reverse the process, and can reconstitute the original signal. While analog encryption is not widely known today (because of the greater convenience of digital encryption for security applications), it was developed in considerable detail in the pre-transistor era and can likely be implemented for this application at relatively low cost using only analog components.

[8] Many different scrambling techniques are available. One approach mentioned by Ingemar and Linnartz, *infra*, and Craver, *supra*, is inverting some feature of each pixel or group of pixels, such as its luminance. Another method might be the addition of a complicated periodic signal known to the attacker, or even the addition of a nonperiodic and essentially random signal. In the ideal case, the scrambled signal is totally uncorrelated with the original signal, so the watermark is completely undetectable. Some analog scrambling techniques may affect quality if they expand the bandwidth or dynamic range of the signal, but it has not to our knowledge been shown or argued that any significant loss of quality must occur. As we suggest below, the techniques of analog scrambling previously practiced before the digital era could be revived for this purpose; many old techniques would have useful properties for this application. Instead of concealing an analog voice recording's contents against eavesdropping by foreign agents, this system would conceal an analog video signal's watermark against detection by a watermark detector.

[9] SIGSALY was a hybrid analog and digital system. Its keying material was made up of pairs of recordings of random thermal noise on identical phonograph records. SIGSALY and other systems of its era demonstrate that modern digital hardware and digital computers are not necessary in order to scramble an analog signal usefully – and reversibly.

Ingemar Cox and Jean-Paul M. G. Linnartz may have been the first to describe this process, in their 1998 paper "Some general methods for tampering with watermarks".[10] Cox and Linnartz describe the attack as follows:

> [C]opy protection based on watermarking content has a further fundamental weakness [in addition to several discussed earlier]. The watermark detection process is designed to detect the watermark when the video is perceptually meaningful. Thus, a user may apply a weak form of scrambling to copy protected video, e.g. inverting the pixel intensities, prior to recording. The scrambled video is unwatchable and the recorder will fail to detect a watermark and consequently allow a copy to be made. Of course, on playback, the video signal will be scrambled, but the user may then simpl[y] invert or descramble the video in order to watch a perfect and illegal copy of a video. Simple scrambling and descrambling hardware would be very inexpensive […] One way to avoid such circumvention for digital recording is to only allow the recording of content in a recognized file format. Of course this would severely limit the functionality of the storage device.[11]

It does not appear to be feasible to defend against this attack simply by improving the strength of watermark technologies against targeted attacks. This attack is not, strictly speaking, a question of the strength or weakness of particular watermarks. Instead, it is a limitation on the applicability of watermarking technology for particular purposes.

Limitations of Technical Approaches: VBI Signaling Technologies

Several schemes proposed to ARDG use the vertical blanking interval (VBI) of an analog video signal to embed digital labels in a predictable and standardized way. Because the vertical blanking interval does not contain video picture data, it is possible to use it to convey a limited amount of digital data out-of-band with respect to the picture. The best-known application for the VBI data is closed captioning (CC). The VBI data could also include copy-control labels, and, as some presenters observed, some standards already permit (but do not require) the use of portions of the VBI in certain video interfaces to convey copy-control labels. This approach is also vulnerable to straightforward attacks.

Although some VBI signaling proposals are combined with watermarks, VBI signaling is importantly different from watermarking. Watermarking attempts to hide a mark within a signal; VBI signaling does not attempt to hide the label at all. In every digital VBI signaling scheme, the copy label is in a publicly documented format at

[10] Ingemar Cox and Jean-Paul M. G. Linnartz, "Some general methods for tampering with watermarks", 16 IEEE Journal on Selected Areas of Communications 583 (1998), available at http://www.neci.nj.nec.com/homepages/ingemar/papers/jsac98.pdf. See also Craver, "What We Expect from Watermarking".

[11] Cox and Linnartz, Section 6.4. The assumption that the result of this technique is necessarily an "illegal copy" is unwarranted.

a publicly documented location within the VBI. This implies that removing or altering the copy label can, as a technical matter, be done using only public information.

Copy labels in the vertical blanking interval can be removed or altered either accidentally or deliberately. Some existing products strip out the entire VBI portion of certain video signals. Because the copy labels appear at a single known location in the video signal, it is technically easy to remove them using known techniques. (CGMS-A proponents note this: for instance, they observe that "[b]lanking or stripping those lines from the VBI that contain CGMS-A and RCI would be technically the easiest way to attack CGMS-A Plus RC" and suggest that doing so "is not difficult with low-cost specially purposed boxes or circuitry"[12].

Indeed, products available today, both lawfully and unlawfully, can likely be used to remove or alter the contents of CGMS-A labels or other copy-control labels in the vertical blanking interval. For example, devices for the insertion and editing of closed-caption data may allow line-by-line editing of the contents of all VBI lines in a particular video format. Some devices strip the entire VBI, or particular lines, inadvertently or in order to impair copy-control applications.

Accessibility for the Disabled: Consequences of VBI Signaling Technologies

Currently, hearing-impaired people are concerned about the preservation of closed-caption data in video signals. Using the VBI, and especially VBI lines shared with closed-caption data, as a location for signaling in a widely deployed copy-restriction system may create unintended consequences for hearing-impaired and other users of closed captions.

Suppose a general decides to locate a military operation next to a hospital. It is illegal to bombard a hospital, so the general may hope to protect the military operation by placing it nearby the hospital.[13] Even though this tactic may tend to protect the military operation, it simultaneously puts the hospital at risk. Where beforehand there was no incentive for fighting or bombardment in the vicinity of the hospital, the proximity of a military target means that the hospital could well become "collateral damage". The military operation's presence creates risk for the hospital that was previously nonexistent.

In this case, placing valuable data desired by consumers (CC data) adjacent to data they have an incentive to obliterate (copy-control labels) simultaneously diminishes the chance that the latter will be removed and increases the chance that the former will be removed. While ARDG participants suggested that regulations do, or will, protect against alterations to VBI data, it is worth considering the market impact with or without regulation. To return to our analogy, putting military operations near hospitals increases risks even in the presence of a legal standard (such as the Geneva Convention) forbidding the bombardment of hospitals. Today, there is generally no incentive to strip the VBI or

[12] Analysis Matrix submitted to Analog Reconversion Discussion Group by CGMS-A Plus RC proponents, answers to questions 2.8 and 3.3.

[13] Convention for the Amelioration of the Condition of the Wounded on the Field of Battle (Geneva, August 22, 1864), Art. 1. However, the Convention conditions the protection of hospitals on their use for a non-military purpose; their protection "shall cease if the ambulances or hospitals should be held by a military force".

any particular line in a video signal. (Indeed, consumers currently have good reason to prefer devices that preserve the VBI to devices that strip it.) If the VBI is widely used for an application that consumers would prefer be absent, they will for the first time have an incentive to develop, use, or purchase devices that strip some or all of the VBI. That could cause such devices to proliferate in the marketplace, to the detriment of the availability of CC data and everyone who relies upon it. The wisest course in order to ensure that all devices preserve the VBI data would be to avoid creating any new incentive to alter it.

Remarks on Particular Technologies

In addition to our submissions in the ARDG analysis matrices, we make the following brief observations about three particular technologies.

Macrovision Corporation

Macrovision Corporation proposes to use existing Macrovision analog copy-restriction technologies in various combinations as a means of signaling copy-control states.

One consequence of this approach is that older devices that happen to be vulnerable to Macrovision signals will treat all marked materials as Copy Never. Federal law requires certain recording devices to exhibit this vulnerability, so, absent a change in the law, many new devices will continue to have this problem.[14]

At the same time, the effectiveness of using any existing Macrovision analog copy-restriction technology to deter deliberate infringement may be limited, since means of removing Macrovision's video signal degradations are widely known. Indeed, Macrovision has published several such techniques in patents.

Dwight Cavendish Systems

Dwight Cavendish Systems claims to have a technology capable of interfering with the functionality of current digitizer hardware, without requiring such hardware to be redesigned.[15]

To our knowledge, this technology remains unpublished and publicly unproven to date. Dwight Cavendish did not provide relevant technical details to ARDG.

VEIL Interactive Technologies

VEIL's proposal involves a feature called a Visual Rights Assertion Mark, or V-RAM. VEIL's explanation in the course of its presentation at ARDG suggests that the V-RAM has many properties in common with a traditional video watermark and therefore that our concerns about video watermarking schemes generally apply to VEIL's scheme as well.

[14] See 17 USC 1201(k) (requiring analog VCRs to "conform" to Macrovision technologies by refusing to record or "exhibit[ing] a meaningfully distorted or degraded display").

[15] Dwight Cavendish's presentation at the October ARDG meeting says its technology is "Effective on legacy equipment [i]ncluding legacy capture cards". Its Analysis Matrix, in the answer to question 2.5, similarly asserts that the technology "can […] provide some control over legacy devices".