Oral presentation at the Security Issues in Multimedia Systems Workshop, IEEE Multimedia Systems '98, Austin, TX, 1998.

Published the Multimedia & Security Workshop, ACM Multimedia '98, Bristol, UK, Sept. 12-16, 1998. GMD Report 41, Ed. J. Dittmann, P. Wohlmacher, P. Horster and R. Steinmetz, 71-76, (1998)

Watermarking in the Real World: An Application to DVD

Matt L. Miller
Signafy, Inc.
4 Independence Way
Princeton, NJ 08540
(609) 734-7620
mlm@signafy.com

Ingemar J. Cox
NEC Research Institute
4 Independence Way
Princeton, NJ 08540
(609) 951-2722
ingemar@research.nj.nec.com

Jeffrey A Bloom
Signafy, Inc.
4 Independence Way
Princeton, NJ 08540
(609) 734-7620
bloom@signafy.com

1. ABSTRACT

The prospect of consumer DVD recorders highlights the challenge of protecting copyrighted video content from piracy. Digital watermarking can be used as part of a copy protection. We describe the copy protection system currently under consideration for DVD. We will also highlight some implementation issues that are being addressed.

1.1 Keywords

Watermarking, DVD, copy protection

2. INTRODUCTION

Digital multimedia watermarking is a field that has received an increasing degree of interest from researchers in both academic and practical settings. The fundamental challenge is to hide a piece of information into a digital image file or a video or audio stream (cover) such that the information is not perceived and cannot be removed without causing significant perceptual degradation to the cover [1]. Since the watermark is embedded into the media, it has the property that it will undergo the same transformations as the media and can thus be used as an indicator of what those transformations may have been.

Some potential applications include the use of a watermark as a signature identifying the copyright owner, as a fingerprint identifying the customer of the cover media, as an authentication key describing some feature of the media which would likely change if the cover were manipulated, or as a copy control mechanism indicating copy permission. Most of these applications rely on the property that watermarks are not easily separated from the content or cover media and, consequently, research into watermarking has focused on the problem of making watermarks difficult to remove without making them perceptible[2].

Since the middle of 1996, we have been working on a copy control application in which watermarks will be one part of a system for protecting video on digital versatile disks (DVD). While the difficulty of removing watermarks is an important problem in this application, we have been confronted with a wide collection of other problems that have been given much less attention in the literature. In this paper we will briefly describe the DVD copy protection framework in which watermarking technology is to be applied and present some of the technical challenges which have not yet been adequately addressed.

3. APPLICATION FRAMEWORK – DVD COPY PROTECTION SYSTEM

In 1996, the Motion Picture Association of America (MPAA), the Consumer Electronics Manufacturers Association (CEMA), and members of the computer industry put together an ad hoc group to discuss the technical problem of protecting digital video from piracy, particularly in the domain of DVD [3]. This group, the Copy Protection Technical Working Group (CPTWG), is open to anyone who wishes to participate, and has no official decision-making power. However, over the past year and a half, it has succeeded in designing the major part of a copy protection system that is bound to become the defacto standard for DVD.

Two major principals have guided the CPTWG's work. The first principal is that the copy protection system should not be mandatory. This immediately divides devices into two categories: "compliant" devices, which implement the protection system, and "non-compliant" devices, which do not. The media to be protected must be scrambled in such

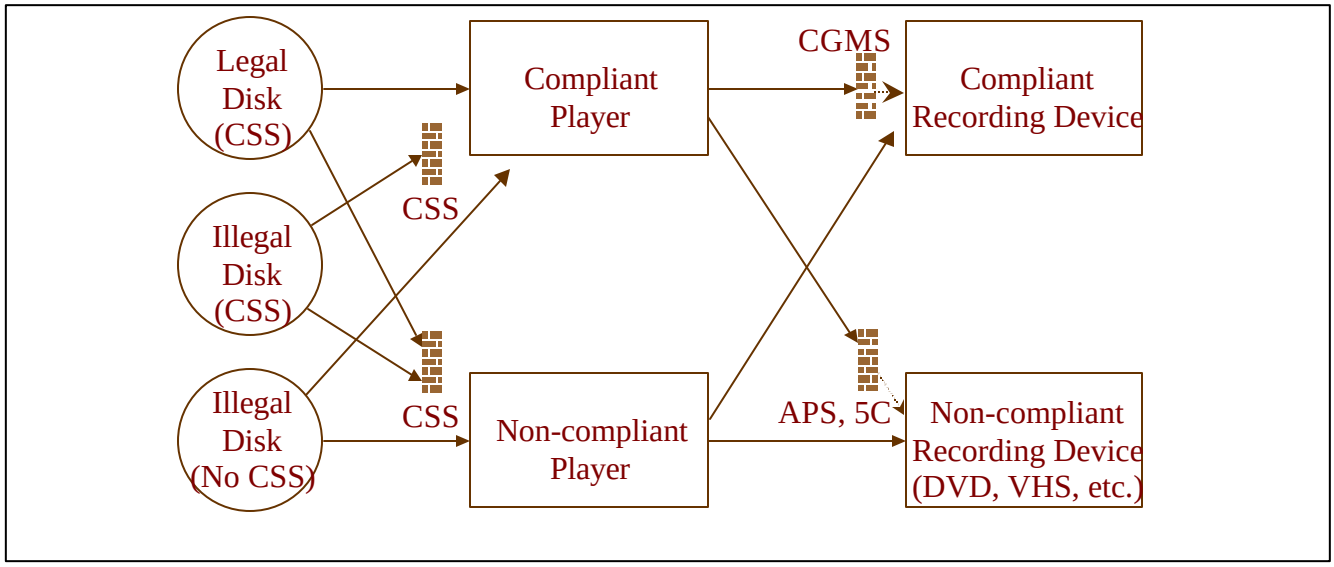

Figure 1. DVD copy protection system *without* watermarking

way that it cannot play on non-compliant devices, or else there will be no protection at all.

The second principal is that the system must be cost-effective. This means it is unlikely to be secure against determined hackers, since that level of security would require more computing power than is reasonable in low-cost consumer devices. Rather, the aim is to come up with a system that is cheap, and good enough to prevent the kind of mass, casual copying that has become prevalent in audio. The design mantra is "keeping honest people honest."

The system designed by the CPTWG is still a work in progress. At present, there are three components that are already being built in to consumer devices. These are the Content Scrambling System (CSS), the Analog Protection System (APS), and the Copy Generation Management System (CGMS). Two additional components are being seriously considered: a system for secure communications across a PC bus (designed by a coalition of 5 companies, and hence referred to as 5C), and watermarking. The watermarking component, of course, is the topic of this paper. The other four components are described below.

- CSS is a low-cost method of scrambling MPEG-2 video, developed by Matsushita. To descramble the video, a device requires a pair of keys. One of the keys is unique to the disk, while the other is unique to the MPEG file being descrambled. The keys are stored on the lead-in area of the disk, which is generally only read by compliant drives. Keys can be passed from a DVD drive to a descrambler over a PC bus using a secure handshake protocol (different from 5C).

  The purpose of CSS is twofold. First and foremost, it prevents byte-for-byte copies of an MPEG stream from being playable, since such copies won't include the keys. Second, it provides a reason for manufacturers to make compliant devices, since CSS scrambled disks won't play on non-compliant devices. Anyone wishing to build compliant devices must obtain a license, which contains the requirement that the rest of the copy protection system be implemented.

- The APS system, developed by Macrovision, is a method of modifying NTSC signals so that they can be displayed on televisions, but cannot be recorded on VCR's. It works by confusing the automatic gain control in VCR's, and this usually leads to unwatchable recordings. Before being adopted for DVD, it has been widely used on videocassettes and in set top boxes.

  Of course, the data on a disk is not NTSC encoded, so APS has to be applied by the NTSC encoder in a DVD player. The information of whether a given video stream should have APS applied, and details about how it should be applied, is stored in the MPEG stream header.

- CGMS is simply a pair of bits in the header of an MPEG stream that encode one of three possible rules for copying: "copy-always" (the video may be freely copied), "copy-never" (the video may never be copied), or "copy-once" (a first generation copy may be made, but no copies may be made of that copy). The copy-once case is included to support such uses as time shifting, where a copy of broadcast media is made for later viewing. Copy-once is unlikely to appear on

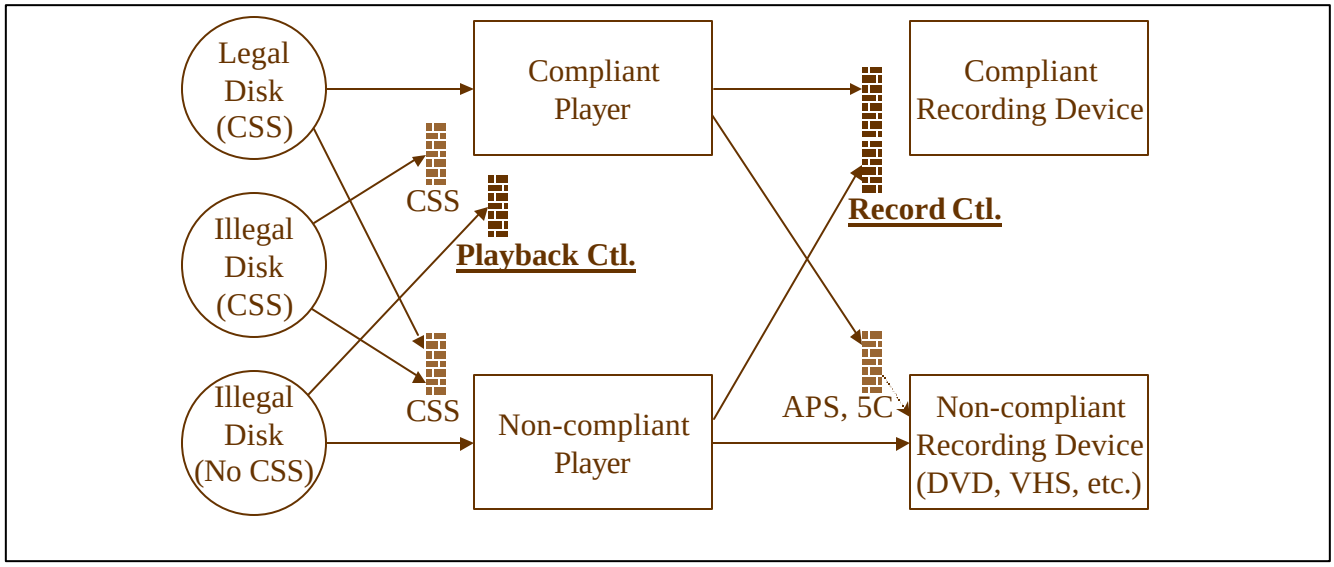

Figure 2. DVD copy protection system *with* watermarking

recorded disks, but it is important for DVD recorders to support it.

- The proposed secure transmission system, 5C, provides a mechanism for pairs of compliant devices on a computer bus to exchange keys, so they can send encrypted data to one another that no other devices can decrypt. The system is more secure than the handshake used for CSS.

  Development of 5C was prompted by the advent of high-speed computer busses such as 1394, which can potentially carry uncompressed digital video from a player or set-top-box to a monitor. The fear is that a pirate could tap into the bus and record any unencrypted video being transmitted.

The role of these copy protection devices is illustrated in Figures 1 and 2. Figure 1 shows the system without watermarking and demonstrates the need for watermarking. In this illustration we assume that available in the marketplace will be both compliant and non-compliant players and recording devices. Three possible types of disks are considered: factory-pressed, legal disks containing copy protected video, bit-for-bit illegal copies of these disks, and illegal copies made of the video after descrambling.

Legal disks will be scrambled with CSS and can be played only on compliant devices. Bit-for-bit copies of these disks won't be playable on any devices, because they won't contain the descrambling keys. This in ensured by storing the keys on the lead-in area of the legal disk, which is only read by compliant drives. The compliant drives take precautions to prevent the keys from being copied.

CGMS is intended to prevent illegal copies, however a non-compliant player may strip out these copy control bits from the header, leaving the video in the clear, or unprotected. At this point there is nothing left to indicate copy restrictions to the compliant recording device and DVD RAM disks without CSS or CGMS can be generated.

Another potential weak point in the system is in the protection against copies being made on non-compliant recorders. APS works only on VCR's, and 5C works only when the display device is a compliant, digital monitor. If the output of the player is, for example, analog RGB, a pirate can simply route it into an appropriate non-compliant recorder and make an unencrypted copy. Of course, such a copy would not contain the CGMS bits.

Because of these two weaknesses, it can be expected that many unprotected, illegal copies will be made. These can be widely distributed, since they will play in either compliant or non-compliant devices. The purpose of introducing watermarking into this system is twofold: first, to improve the protection provided by CGMS by making the copy-control information harder to remove, and, second, to reduce the value of illegal, unencrypted copies when they are made, by making them unplayable on compliant devices.

Figure 2 shows the same scenario except that now watermarking is included. The two functions of the watermark mentioned above are referred to as "record control" and "playback control", respectively. Record control takes over the job of CGMS. It works regardless of how the video reaches the compliant recorder, since the watermark that contains the CGMS data is never removed by normal video processing.

Copy-once control can also be implemented in the compliant recording device. Recording of source data containing this copy-once watermark is allowed, however

some modification is made to indicate a third state called copy-no-more which can be treated the same as copy-never.

Playback control introduces a new point of protection in the system. Should a pirate be successful in generating a DVD RAM copy of a protected video without CSS, this copy will still contain the watermark. The compliant players can now recognize as illegal a video marked with copy-never that is being read from an unscrambled DVD RAM and refuse playback. This playback control limits the potential market for pirated DVD to those consumers who own non-compliant players, which will not play legal disks.

In the summer of 1997, after receiving presentations on watermarking technologies from several companies, the CPTWG set up the Data Hiding SubGroup (DHSG) to evaluate these systems and determine whether the technology is mature enough for inclusion in the copy protection system. The CPTWG issued a call for proposals [4] in July 1997. Eleven companies responded with proposals. After the initial round of testing, seven proposals remain under consideration.

The remainder of this paper describes some of the challenges that are faced by the companies that submitted proposals to the DHSG.

4. CHALLENGES

As the copy protection system described above and illustrated in Figure 2 is implemented an array of challenges related to the watermarking technology have arisen. The issue of watermark removal is often addressed in watermarking literature and remains an important concern [5]. There are a number of other issues, some technical and some non-technical, which have also come to play an important role. In the remainder of this section we briefly introduce and discuss the following issues: enforcement, system tampering, detector placement within the system, computational cost of the detector, effects of geometric distortion, interaction between the watermarking and compression systems, false positive rates and analysis, and copy generation control.

Enforcement - One interpretation of Figure 2 is that the DVD world may be split in two, one compliant and one non-compliant. The copy protection system, specifically the watermarking technology and the CSS, will prevent legal copies from being played on non-compliant players and illegal copies from being played on compliant players. This doesn't stop consumers from owning two players, one compliant and one non-compliant, and does not prevent the sale of a "dual" player containing both compliant and non-compliant drives. The approach taken to discourage the manufacture of "dual" players is to note that both the CSS and watermarking technologies are protected by patents and may only be used in a DVD player with the proper licenses. These licenses will specify that the player must not possess the capability of playing non-compliant DVD sources. We will then rely on the expense of owning two DVD players and the fact that non-compliant DVD copy protected source is illegal as a violation of the content provider's legal copyright, to help "keep honest people honest."

System Tampering - The illegal copy without CSS of the Figure 1 scenario was rendered unplayable by the watermarking technology in Figure 2. This suggests that the pirate has an interest in being able to remove the watermark. Watermarks that are image independent can easily be reconstructed by frame averaging and, once found, can be subtracted from the watermarked video source. Another documented “attack” on watermarks is called sensitivity analysis in which a detector is used to reconstruct the watermark in a frame by a systematic degradation of the image. Again, once found, the watermark can be subtracted from the video source. The field of watermark removal is very active and the robustness of watermarking techniques is constantly being challenged. While possession, sales, and distribution of illegal copies are prohibited by law, there are no such constraints on the sales of watermark removal hardware or software.

There are two common approaches to this problem. The most obvious approach is to invent a watermark that is truly tamper resistant. The other, perhaps more realistic approach may seem at first to be counter intuitive. A company that relies on the tamper resistance of a watermarking technology may wish to actively seek out, invent, and patent any reasonable technique for removing that watermark. Any watermark removal software or hardware using these techniques would then represent a patent infringement.

Beyond watermark removal there are other ways to circumvent the copy protection system. These include hardware modification to disable watermark detection and source scrambling such that the watermark detector does not recognize the source as watermarked video. In this latter case the video must be descrambled after it passes by the watermark detector. Neither of these two approaches can be used to generate an illegal copy that will play on compliant players.

Detector Placement – An issue of significant debate within the DHSG involves the physical placement of the watermark detector in the system. This is of particular interest for DVD drives installed in personal computers. Two reasonable approaches are shown in Figure 3. In the scenario of Figure 3a the watermark detector is located inside the MEPG codec and in Figure 3b it is in the DVD drive. Each of these solutions has its advantages and its disadvantages. Having the detector in the MPEG codec is an efficient solution since both the codec and the detector can share many of the same elements (tables, buffers, etc.)

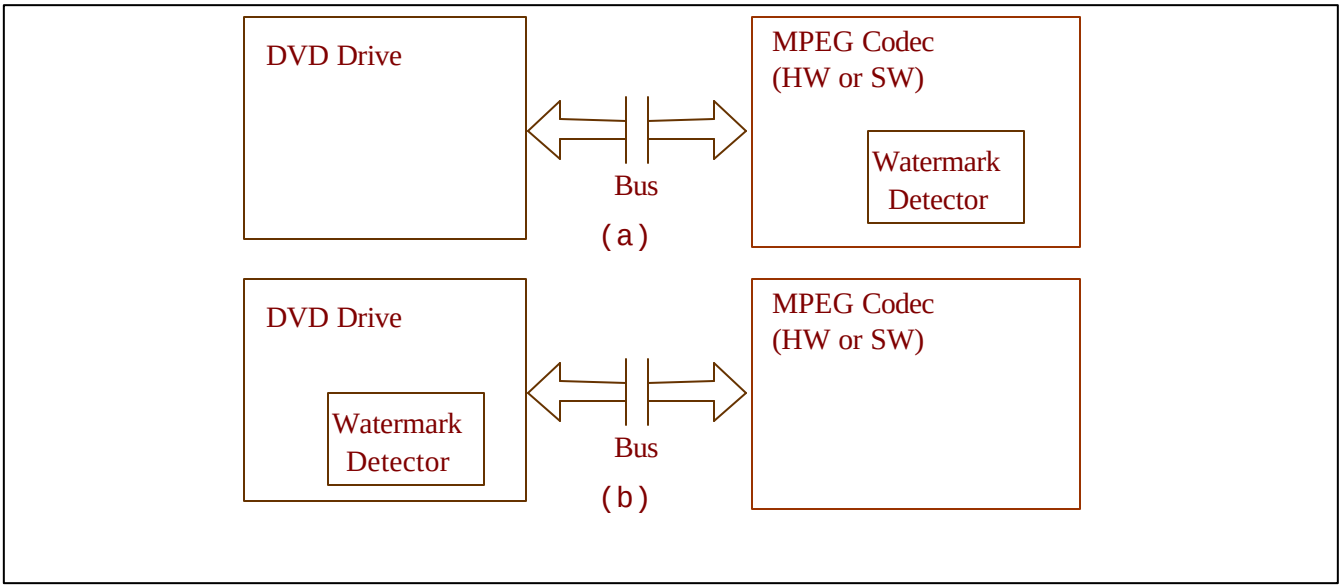

Figure 3. Watermark Detector Placement

However, this solution also allows easy creation of a "dual" system in a computer for example, since most MPEG decoding applications will use non-compliant MPEG decoders.

The second scenario, which is currently leading in the debate, places the watermark detector in the DVD drive. This has the advantage that it is more tamper resistant. Record control will prevent watermarked, non-compliant MPEG bitstreams from being recorded. The DVD player also has knowledge of the disk type (ROM or RAM) from which the video is being read and can check for an allowed combination of disk type and watermark (e.g. copy-never and copy-once should not be found on a RAM disk).

Detector Computational Cost – Adding a watermark detector to a DVD RAM drive will require some degree of redesign. In order to minimize that cost, drive manufacturers have indicated that the detector must fit onto unused silicon that already exists in the drives. This restriction on the cost of the watermark detector in the DVD application is means that the detector must be implemented in about 30k gates. A significant implication is that the detector may not use a frame buffer and must process the video in real time without reference to previous frames. This shows the asymmetry between the watermark embedder and decoder since the motion picture industry is likely to accept an embedder with very high computational cost and physical cost on the order of $100,000.

Geometric Distortion – DVD players have the facility to geometrically alter the video in two important ways. Letterbox is a technique which changes the aspect ratio from 4:3 to 16:9. Panscan represents a cropping of the larger image. The watermark must survive these geometric distortions as well as more arbitrary scaling and cropping which a pirate may use to avoid watermark detection. While these issues are generally addressed in watermarking literature, this special case where a frame buffer may not be available is particularly difficult.

Watermark/Compression Interaction – It can be argued that a goal of video compression, to remove all visually imperceptible information, makes the challenge of imbedding a visually imperceptible watermark much more difficult. If the watermark is placed in perceptually significant component, the source may be more difficult to compress.

In the DVD application, MPEG-2 compression is used and it is required that the watermark be detectable in both the compressed data stream and the reconstructed video. The former case requires detection in the block-based DCT domain (without frame buffers as previously mentioned) and both cases require that the watermark survive MPEG quantization. Another requirement is that the watermarks be modifiable in the compressed data stream without complete decompression and that the modifications not affect the bit-rate or position of I-frames.[6] The scalability features of MPEG-2 further complicate watermark detection and modification in the bitstream.

False Positive Rate – Watermark detection can generally be expressed as a binary decision and there are penalties associated with incorrect decisions. In the DVD application, when the detector decides that a watermark is present in video that does not contain a watermark, the result will be that a user cannot do some action that should be allowed. A couple might never be able to watch their wedding video. A football fan might not be able to record the Super Bowl for time shifting. The latter example is particularly catastrophic; if a piece of the Super Bowl

triggers a false positive, *no one* will be able to record it on DVD. Our estimates of the required false positive rate are about one in 10^{11} or 10^{12} distinct frames. A recent model for predicting the false positive rate can be found in [7].

Copy Generation Control – There are a number of proposed methods for using watermarks in a copy generation control system. The goal is to detect a copy-once state and change it to a copy-no-more state as the video is being recorded. One approach is to use a watermark that can actually be changed. Recall that this will need to be done in the MPEG stream without changing the bitrate. This approach is likely to be more susceptible to tampering since the ability to change a watermark implies the ability to remove it.

Another approach involves the addition of a separate watermark. Thus the copy-once state will be indicated by the presence of one watermark and copy-no-more by the presence of both. To do this, the DVD recorder, with it's limited computational complexity and cost, must be able to insert the copy-no-more watermark. As with the other watermarks, this copy control watermark must be unobtrusive, indelible, and robust.

The opposite approach can also be taken where the presence of two watermarks, one of which is fragile, represents the copy-once state. The recorder then has the task of removing the fragile watermark. An interesting example of this can be found in the Macrovision/Digimark proposal [4] in which the fragile watermark is a visible pattern (placed in the overscan area of the frame so that it will be hidden by the edge of a television screen). This pattern is designed in such a way that it cannot be recorded on a VCR. Thus, a copy on a VCR removes the fragile mark, and automatically converts copy-once into copy-no-more. Of course, a digital recorder will still have to remove the fragile watermark explicitly.

A completely different approach to generation control is to use information that is not embedded in the watermark, but must be available to the recorder if copy-once video is to be recorded. Such information, often referred to as a "tag" or "ticket", might be stored in MPEG headers or in the vertical blanking interval of analog video. In such a system, the copy-once state is represented by the presence of a copy-once watermark *and* an appropriate tag. The watermark without the tag would indicate copy-no-more. Since no mechanism would be provided for copying the tag, any copy would necessarily be labeled with copy-no-more. A weakness of this method is that all devices that do not copy the video, such as set top boxes, must preserve the tag. Current set top boxes would have to be modified for this purpose.

5. CONCLUSION

We have described here several of the difficult problems encountered in designing a real-world application of watermarking for copy control in DVD. While the problems of fidelity and robustness have received ample attention in the literature, several of the problems encountered here are less studied. The most notable of them are

- Interaction with compression algorithms
- Overall system design to avoid circumvention
- False positive rates
- Issues of computational costs, such as designing detectors without using frame buffers

The details and relative importance of these problems change with different applications. But they all pose fundamental challenges that must be met before watermarking can fulfill its promise as a tool for copyright protection.

6. REFERENCES

[1] Cox, I.J.; Kilian, J.; Leighton, F.T.; Shamoon, T., Secure spread spectrum watermarking for multimedia, *IEEE Transactions on Image Processing*, vol.6, no.12, p. 1673-87, 1997.

[2] Cox, I.J. and Miller, M.L., Review of watermarking and the importance of perceptual modeling, Proc. SPIE, vol.3016, p. 92-9, 1997.

[3] Bell, A., Personal communication, 15 May, 1998.

[4] DHSG Call for Proposals, http://www.dvcc.com/dhsg.

[5] Cox, I.J. and Linnartz J-P., Some General Methods for Tampering with Watermarks, *IEEE Journal on Selected Areas in Communications*, Vol. 16, pp. 587-93, 1998.

[6] Hartung, F. and Girod, B., Digital watermarking of MPEG-2 coded video in the bitstream domain, *IEEE International Conference on Acoustics, Speech, and Signal Processing*, pp. 2621-4, 1997.

[7] Hernández, J. R., et.al., Performance analysis of a 2-D-multipulse amplitude modulation scheme for data hiding and watermarking of still images, *IEEE Journal on Selected Areas in Communications*, Vol. 16, pp. 510-24, 1998.