Copy Protection for DVD Video

JEFFREY A. BLOOM, INGEMAR J. COX, SENIOR MEMBER, IEEE, TON KALKER, MEMBER, IEEE, JEAN-PAUL M. G. LINNARTZ, MEMBER, IEEE, MATTHEW L. MILLER, AND C. BRENDAN S. TRAW

The prospect of consumer digital versatile disk (DVD) recorders highlights the challenge of protecting copyrighted video content from piracy. We describe the copy-protection system currently under consideration for DVD. The copy-protection system broadly tries to prevent illicit copies from being made from either the analog or digital I/O channels of DVD recorders. An analog copy-protection system is utilized to protect the NTSC/PAL output channel by preventing copies to VHS. The digital transmission of content is protected by a robust encryption protocol between two communicating devices. Watermarking is used to encode copy-control information retrievable from both digital and analog signals. Hence, such embedded signals avoid the need for metadata to be carried in either the digital or analog domains. Finally, the copy-protection system provides the capability for one-generation copying. We discuss some proposed solutions and some of the implementation issues that are being addressed.

Keywords*—*Copy protection, digital versatile disk (DVD), digital video, encryption, MPEG, tickets, watermarking.

I. INTRODUCTION

Since the middle of 1996, the authors of this paper have been working on a variety of issues related to the copy control of video stored on digital versatile disks (DVD's). DVD players were introduced into the consumer market segment in late 1996 and have received an enthusiastic response. Current DVD players provide a very high-quality video signal that is encoded on a read-only disk with the same form factor as conventional audio CD's. The capacity of these disks is significantly larger than audio CD's, 4.7 Gbytes per layer per side versus 650 Mbytes. Prerecorded movies are MPEG-2 compressed and subsequently encrypted prior to being stored on DVD.

The advantages of digital video come at the price of an increased risk of illegal copying. Hollywood studios are very familiar with piracy issues, however DVD disks, and digital video recording in general, raise even more concern since each copy is a perfect reproduction. This is in contrast to traditional VHS tape copying in which the video fidelity is degraded with each generation of copying, e.g., a copy of a VHS tape looks inferior to the original tape. Thus, as second-generation DVD players with digital video recording capabilities are likely to be introduced in the 1999 timeframe, there is a pressing need to provide several levels of copy protection.

Traditionally, protection of digital data has been provided by a variety of encryption methods and DVD is no exception. However, encryption alone does not provide an adequate solution as it only provides for robust delivery of the content. Once the content is decrypted, it is no longer protected and the decrypted content is accessible through the analog input/output (I/O) channels, e.g. NTSC, PAL, and RGB ports. Content providers have therefore insisted on additional protection. An overview of the current DVD copy-protection system and its weaknesses is provided in Section II.

Section III then provides a more detailed description of the robust digital transmission system that will allow digital video content to be transferred between devices.

The U.S. Supreme Court has held that personal, home videotape recording of a television broadcast for time-shifting purposes is a fair use and therefore does not constitute copyright infringement.[1] However, consumers are not entitled to make a copy of this copy, i.e., a second-generation copy. The DVD copy-protection system is designed to support a copy generation management system. This requires at least two bits of information to be associated with a piece of video indicating one of the following copy states: "copy_never"; "copy_once"; "copy_no_more"; or "copy_freely." These metadata are problematic in two respects. First, the computer industry does not want to be held responsible for propagating these metadata, as they may be easily stripped from a video stream. Second, it is desirable that this copy-control information (CCI) survive both digital-to-analog and analog-to-digital conversions. It was therefore decided that CCI would also be encoded using watermark technology. This will be the first widespread adoption of video watermarking technology and will require that all DVD players and

Manuscript received Macrh 2, 1999; revised April 25, 1999.

J. A. Bloom and M. L. Miller are with Signafy, Inc., Princeton, NJ 08540 USA.

I. J. Cox is with NEC Research Institute, Princeton, NJ 08540 USA.

T. Kalker and J.-P. M. G. Linnartz are with Philips Research, NL-5656-AA, Eindhoven, The Netherlands.

C. B. S. Traw is with Intel Corporation, Hillsboro, OR, 97124 USA.

Publisher Item Identifier S 0018-9219(99)04956-7.

[1] U.S. 417 (1984) Sony Corp. of America versus Universal City Studios, Inc.

recorders possess watermark detection circuitry. At the time of this writing, two proposals are under consideration for the DVD watermark technology. These proposals are sponsored by two coalitions of companies: IBM, NEC, Sony, Hitachi, Pioneer, and Signafy (a subsidiary of NEC USA) on the one hand and Philips, Macrovision, and Digimarc on the other. Section IV describes some of the design constraints that a watermark must meet.

Perhaps the most difficult or awkward issue in the design of the copy-control system has been generational copy control. Two classes of solution have been proposed in order to represent the change from "copy_once" to "copy_no_more." These two solutions are discussed in Section V. Finally, Section VI gives brief conclusions.

As a final remark of this section, we would like to note that some of the authors of this paper are affiliated with companies which have directly competing DVD watermarking technologies. In this way we, the authors, would like to make a case for scientific cooperation, despite conflicting business interests.

II. 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 problems of protecting digital video from piracy, particularly in the domain of DVD [1]. 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 likely to become the *de facto* specification for DVD copy protection.

Two major principles 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 "noncompliant" devices, which do not. The medium to be protected must be scrambled in such way that it cannot play on noncompliant devices.

The second principle is that the system must be cost-effective. This means that 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 system must be cheap and robust 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 into consumer devices. These are the Content Scrambling System (CSS), the Analog Protection System (APS), and the Copy Generation Management System (CGMS). Three additional components are being seriously considered: a system for robust exchange of content across digital interconnect (designed by a coalition of five companies, and hence referred to as 5C), media identifiers, and watermarking. The watermarking and robust communications components are discussed in detail in this paper. A brief description of each of the six components is given below.

- CSS is a low-cost method of scrambling MPEG-2 video, developed by Matsushita. Descrambling 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 will not include the keys. Second, it provides a reason for manufacturers to make compliant devices, since CSS scrambled disks will not play on noncompliant 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/PAL 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 a severe degradation of the content quality. Before being adopted for DVD, it has been widely used on videocassettes and in set-top boxes (STB's).

  Of course, the data on a disk are not NTSC/PAL encoded, so APS has to be applied by the NTSC/PAL 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 are stored in the MPEG stream header.
- CGMS is a pair of bits in the header of an MPEG stream that encode one of three possible rules for copying: "copy_freely" (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 recorded disks, but it is important for DVD recorders to support it.
- The proposed content protection transmission system, 5C, provides a mechanism for compliant devices on a bus to exchange keys in an authenticated manner, so they can send encrypted data to one another that no other devices can decrypt. The system is more robust than the handshake used for CSS.

  Development of 5C was prompted by the advent of high-speed interconnects such as 1394, which can potentially carry digital video between devices such as PC's, players, STB's, and displays. The fear is

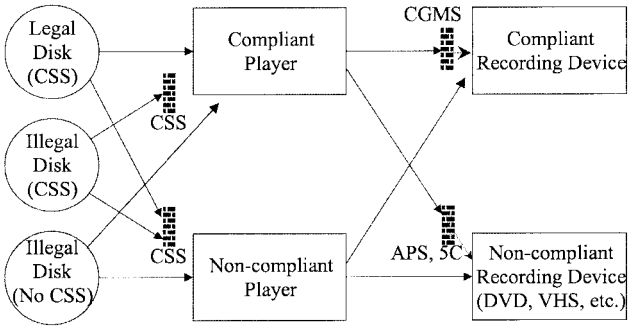

Fig. 1. DVD copy-protection system without watermarking.

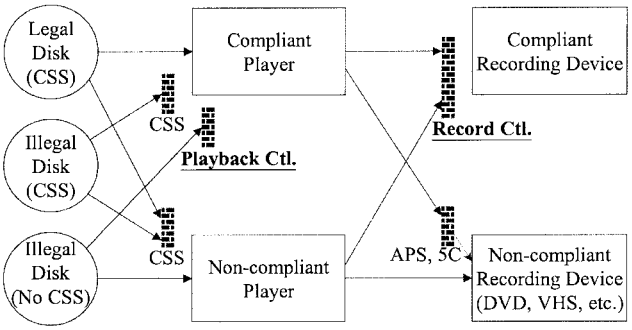

Fig. 2. DVD copy-protection system with watermarking.

that a pirate could tap into the bus and record any unencrypted content being transmitted. Digital transmission content protection is discussed in more detail in Section III.

- Secure (physical) media identifiers can distinguish between original media and copies. Several initiatives have been launched to find technical mechanisms by which a player can identify recordable media, identify whether a compliant recorder has produced the disk, and distinguish between original ROM discs and piracy stamped copies.
- Watermarking is a technique for hiding information directly in video by making small, unnoticeable distortions in the frames. In DVD, it is intended primarily as a more secure form of CGMS. The CGMS bits do not survive digital to analog conversion, and can be trivially stripped from an MPEG stream. Watermarks encoding the same information will not be so easily stripped in normal video processing.

  A secondary purpose of watermarking is to encode the bits used for controlling APS, which have the same weaknesses as the CGMS bits.

The role of these copy-protection devices is illustrated in Figs. 1 and 2. Fig. 1 shows the system without watermarking and demonstrates the role of watermarking. In this illustration we assume that both compliant and noncompliant players and recording devices will be available in the marketplace. Three possible types of disks are considered in this figure: factory-pressed, legal disks containing copy-protected video; bit-for-bit illegal copies of the video sectors on these disks; and illegal copies made of the video after descrambling. Other types that are relevant in the discussion are legal home recordings without CSS and legal prerecorded disks without CSS. For the sake of simplicity, the set of disk types has been limited in Figs. 1 and 2. Particularly, legitimate home recording is interesting, because a hacker or pirate may attempt to disguise an illegal copy to appear as a legal recording of copy_freely content.

Most legal disks will be scrambled with CSS and can be played only on compliant devices. Bit-for-bit copies of these disks will not be playable on any devices because they will not contain the descrambling keys. This is 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, in particular of unscrambled content. However, a noncompliant 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 noncompliant recorders. APS works only on VCR's and 5C is designed for digital connections and will not provide viewing capabilities to analog and noncompliant monitors. If the output of the player is, for example, analog RGB, a pirate can simply route it into an appropriate noncompliant 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 would be made. These could be widely distributed, since they would play in either compliant or noncompliant devices. The purpose of introducing watermarking into this system is twofold: first, to improve the protection provided by CGMS by making the CCI harder to remove, and second, to reduce the value of illegal, unencrypted copies when they are made by making them unplayable on compliant devices.

Fig. 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 fourth 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 noncompliant players (which will not play legal disks). As playback control basically prevents against illegal distribution of copies, its potential exceeds that of only countering simple casual copying to DVD RAM. Playback could potentially be controlled by the relationship between the watermark and physical disk properties or the presence of CSS. Reliance on the presence of CSS alone, however, would base the security of the playback control system on the impossibility of reading from and writing in the lead-in area. This has potential weaknesses and may preclude certain future enhancements of the standards for recordable media.

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 DHSG issued a call for proposals [2] in July 1997. Eleven companies responded with proposals. After the initial round of testing, seven proposals remained under consideration. There has since been some consolidation such that two joint proposals are currently being evaluated.

The technical solution is only part of the solution to the complicated copy-protection problem. The solution will work only if the majority of the recording devices in the marketplace are compliant. One interpretation of Fig. 2 is that the DVD world may be split in two, one compliant and one noncompliant. The copy-protection system, specifically the watermarking technology and the CSS, will prevent legal copies from being played on noncompliant players and illegal copies from being played on compliant players. This does not stop consumers from owning two players, one compliant and one noncompliant, and it does not prevent the sale of a "dual" player containing both compliant and noncompliant 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 noncompliant DVD sources. We will then rely on the expense of owning two DVD players, the fact that copy-protected DVD source will not play on noncompliant players, and the fact that noncompliant DVD copy protected source is illegal as a violation of the content provider's legal copyright, to help "keep honest people honest." This is of course impossible to enforce in a PC environment.

III. Content Protection During Digital Transmission

The proliferation of digital entertainment content has led to the introduction of interconnect technologies to enable its exchange between devices within a home entertainment environment. The IEEE 1394 serial bus [3] has been widely adopted by the consumer electronics and personal computer industry to connect digital content handling devices. Content being exchanged via a digital interconnect is vulnerable to unauthorized copying since it is traveling between devices across well-defined interfaces typically with publicly available specifications. Without technical means to protect the content, it could be copied by any device that can be connected to the bus.

Because of the risk of unauthorized coping, content owners require content protection. These requirements have been legally expressed in the DVD CSS license agreement. Specifically, this license agreement requires that all digital outputs on a device that can access CSS protected content must include an approved content protection technology.

A. Content Owner Protection Requirements

The content owners have established the following basic content protection requirements for digital interconnects [4]:

- robust, mature, cryptographic authentication and protection for content and CCI; the content owner uses CCI to specify under what conditions the content can be copied;
- system renewability to maintain the integrity of the protection system in the event that reverse engineering and/or development of circumvention devices occurs;
- licensable components to enable use of intellectual property law to prosecute manufacturers of circumvention devices.

B. Digital Transmission Discussion Group

The Digital Transmission Discussion Group (DTDG) is a subgroup of the CPTWG chartered to define a content protection system. This system should be capable of preventing the unauthorized use of commercial entertainment content by ordinary consumers when that content is transferred in digital form between interfaces compliant with the IEEE 1394-1995 High Performance Serial Bus standard. The DTDG released a call for proposals (CFP) [5] requesting solutions addressing the following three technical elements of a system.

- Robust exchange of CCI.
- Encryption of content. By encrypting the content, complaint devices are able to transfer the content in a manner which prevents illegal usable copies of the protected content from being made.
- Authentication and key exchange (AKE). Authentication provides a technical means for compliant devices to establish the authenticity of another device. Key exchange enables authenticated devices to establish the keys necessary to exchange encrypted content.

The DTDG received 11 responses to the CFP. After several months of technical discussions, three proposals remained under consideration after several of the proposals were withdrawn, transferred to the DHSG, or merged together. Extensive discussions were stimulated by the

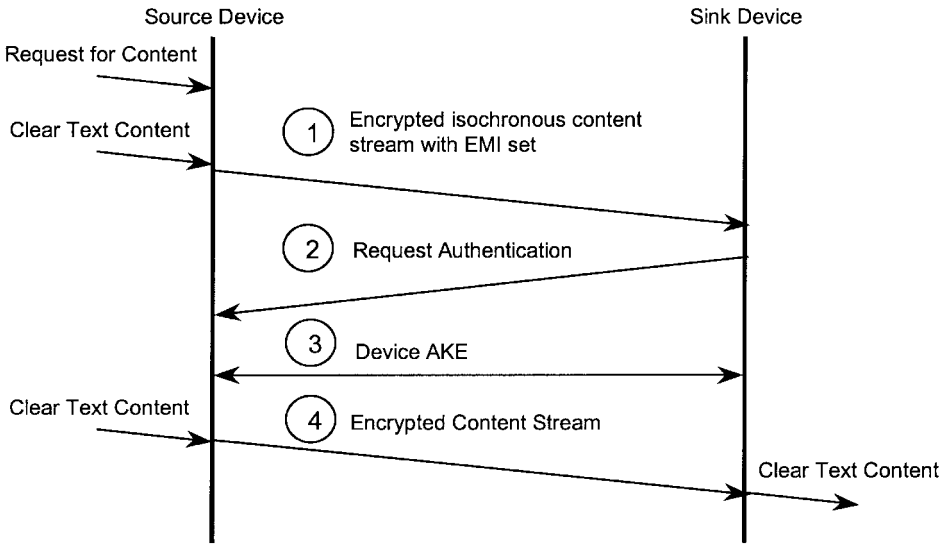

Fig. 3. Operation of 5C DTCP.

engineering tension between the content industries' need for robust protection, the consumer electronics industries' need for a low-cost hardware implementation, and the information technology industries' need for an efficient software implementation.

According to the discussion group's final report [6] issued prior to disbanding, all of the surviving proposals addressed the three requested system elements and were deemed to meet the MPAA standard of "keeping honest people honest."

C. 5C Digital Transmission Content Protection (DTCP)

The 5C DTCP technology [7] developed by Hitachi, Intel, Matsushita, Sony, and Toshiba is an example of the digital transmission solutions considered by the DTDG. In response to MPAA requirements, the 5C DTCP provides system renewability as well as an encrypted exchange of content and CCI between authenticated devices.

CCI is carried embedded in the content stream according to the content format (e.g., MPEG). In addition, it is mapped into an encryption mode indicator (EMI) that provides protected, yet easily accessible, access to the CCI.

Content is encrypted using the M6 block cipher that is used in converted cipher block chaining (CBC) mode with 56-bit keys.

Two AKE procedures based on challenge/response procedures are defined to enable manufacturers to trade off implementation complexity verses value of content to be handled:

- full authentication (for all content) is based on Digital Signatures and Diffie–Hellman Key Exchange using a 160-bit elliptic curve public key cryptosystem compatible with IEEE P1363 [8];
- restricted authentication (acceptable for "copy_once" and "copy_no_more" content only) is based on shared secret techniques.

System renewability is provided through device certificate revocation. The license administrator can, under a rigorously specified set of conditions, exclude individual, compromised devices from participating in the protection system with devices supporting full authentication. Revocation lists are carried in system renewability messages that are distributed with content and between compliant devices.

The application of this system is not limited to IEEE 1394 serial busses. It can be used with any interconnect which provides, at a minimum, low bandwidth, bidirectional communications.

Fig. 3 shows an overview of the operation of the content protection system. The device that is the source of protected content has been instructed to transmit the content via the IEEE 1394 serial bus' isochronous transport.

Step 1) The source device is requested to initiate the transmission of a stream of protected content. The embedded CCI of the content is examined to determine the appropriate EMI value (e.g., "copy_once," "copy_never," or "copy_no_more") to associate with the encrypted content stream. The source device may choose to transmit an empty content stream until at least one device has completed the appropriate authentication procedure.

Step 2) Upon receiving the content stream, the sink device inspects the EMI to determine the copy-protection status of the content. If the content is marked "copy_never," the sink device requests that the source device initiate full AKE. If the content is marked "copy_once" or "copy_no_more" the sink device can request restricted AKE if full authentication is not available. If the sink device has previously performed the appropriate authentication, it can immediately proceed to Step 4).

Step 3) When the source device receives the authentication request, it proceeds with the type of authentication requested by the sink device, ensuring that Full AKE is performed if the content is marked "copy_never".

Step 4) Once the devices have completed AKE, the keys required to access the encrypted content stream are exchanged between the devices.

D. Future Work

While several suitable content protection technologies for IEEE 1394 serial busses are available, additional work is still needed to address the specific content-protection requirements of other busses and interconnects including the Universal Serial Bus, conditional access smartcard interfaces, and extremely high-bandwidth baseband digital video interfaces under development.

IV. Watermark System Design Issues

As the copy-protection system described in Section II and illustrated in Fig. 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 [9]. There are a number of other issues, some technical and some nontechnical, which have also come to play an important role. In this section we briefly introduce and discuss the following issues: computational cost of the detector and embedder; false positive rates; detector placement within the system; interaction between the watermarking and video compression systems; and robustness of the watermark to common signal processing and intentional tampering.

A. Economic Costs

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 means that the detector must be implemented in about 30 000 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. In the standardization of DVD audio, 90 000 gates are often mentioned as a target for the complexity of the watermark detector.

Watermark insertion is expected to be performed during postproduction of the movie and prior to MPEG compression. Thus, the number of watermark inserters will be small, especially compared with the number of watermark detectors. Consequently, the cost of watermark insertion can be considerably higher. 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 and physical cost.

During the design of the watermark, there has been a conscious effort to standardize the watermark detector but leave the encoder undefined, in a fashion similar to MPEG. In so doing, it is hoped that watermark-insertion technology can continue to improve even though the detection circuitry is fixed.

It should be noted that some proposals support copy generation control by embedding a secondary watermark (see Section V for a detailed discussion). This secondary watermark embedder, must also be very low cost, i.e., implemented within the gate count restrictions for the detector.

B. False Positives 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. Recent models for predicting the false positive rate can be found in [10] and [11].

C. Interaction with MPEG Compression

Currently the storage capacity of DVD disks is limited to 4.7 Gbytes for read-only disks (single side, single layer) and 5.2 Gbytes for DVD RAM (two sided). If a typical feature-length film is to fit on one side of a read-only disk, it must be compressed to an average of 3.8 Mbits/s. Obtaining this compression rate without seriously damaging the quality of the video requires several passes with high-end equipment. Watermarks can sometimes make this compression more difficult by introducing details into the video that the MPEG encoder tries to preserve, thus reducing the bits available for the underlying content. Clearly, it is desirable that the watermarking scheme adopted for DVD minimizes this effect.

A second desirable feature is that the watermark be detectable in both the compressed data stream and the reconstructed baseband video. The former case requires detection in the block-based discrete cosine transform (DCT) domain (without frame buffers, as previously mentioned) and both cases require that the watermark survive MPEG quantization.

D. 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 Fig. 4. In the scenario of Fig. 4(a), the watermark detector is located inside the MPEG codec, and in Fig. 4(b) it is in the DVD drive. Each of these solutions has its advantages and its disadvantages.

1) Watermark Detection in the Drive: The first scenario places the watermark detector in the DVD drive. This has the obvious advantage that as long as the watermark is not compromised, pirated content will never leave the drive (in playback mode) or will never get copied onto a disk (in recording mode). The assumption is that the drive has sufficient intelligence to detect the presence of MPEG

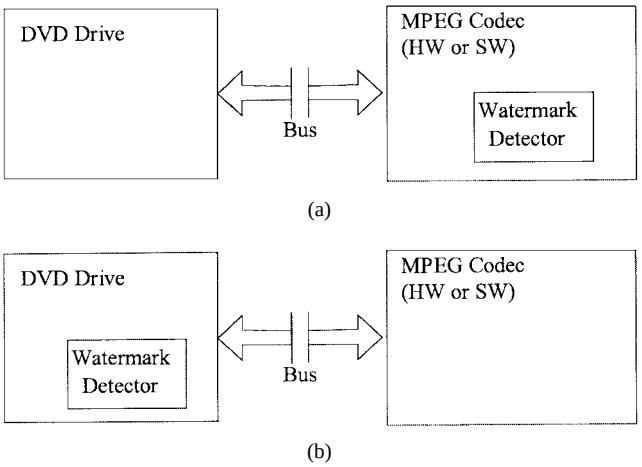

Fig. 4. Watermark detector placement.

streams. Given that assumption, the location of the watermark detector in the drive is secure and tamper resistant. Record control will prevent watermarked, noncompliant MPEG bit streams 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).

Watermark detection in the drive has a number of consequences. First, it implies additional complexity compared to a watermark detector near the application. Minimally, an MPEG parser and dequantizer have to be present because resources cannot be shared with an MPEG decoder. Second, the watermark detector needs logic to recognize the MPEG data type. This task may become challenging since there are a number of possible circumvention strategies [9], including PC device drivers that intentionally read out disk sectors in random order. Increases in DVD drive data rates will place higher demands on such MPEG detection and watermark detection circuitry.

2) Detector Within the Application: An MPEG decoder that implements the watermark detection would be a compliant decoder. Having the detector in the MPEG decoder is an efficient solution since both the decoder and the detector can share many of the same elements (tables, buffers, etc.). However, without special measures, this solution allows easy creation of a "dual" system in a computer. For example, since most MPEG decoding applications will use noncompliant MPEG decoders, a consumer might install both compliant and noncompliant decoders. Then, depending on whether the disk is a legal or illegal copy, one of the two MPEG decoders is invoked. A possible solution to this problem is the use of authenticated links between compliant devices, e.g., a compliant drive and a compliant application program.

A second problem with watermark detection within the application is the threat of illegal taps on MPEG in the clear on links. If a compliant source and sink device have authenticated each other, the video data need to be sent in a robust fashion. If not, a pirate may break into the link and make a copy of the unencrypted content. Furthermore, precautions need to be taken to ensure the integrity of the data. Otherwise, a compliant drive may be fooled into believing that it is outputting legal data when it is not. It follows that a watermark detector within an application needs the following:

- a protocol which enables the recognition of compliant devices;
- a bidirectional link which implements authentication, encryption, and data integrity;
- a protocol between source and sink which informs the drive about whether or not to stop giving out data.

It is not a trivial task to design an efficient bidirectional link as above with a limited amount complexity. One option might be to build upon the existing CSS, and possibly 5C, infrastructure.

However, assuming that a robust digital link can be realized, there are a number of advantages to detection within the application. First, the scheme is extendible to other data types. The method can therefore be upgraded and would support other copy protected data types on the same physical disk. Second, the burden on the complexity of the DVD basic engine (hard drive) is minimized. Third, as the watermark decoder can share many resources with the application, a more powerful watermark detector may be possible.

E. Robustness

1) Common Signal Processing: DVD players have the facility to geometrically alter the video in two important ways. Letterbox is a technique that changes the aspect ratio from 4:3 to 16:9, and 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.

The effect of scaling is usually not that the watermark is lost, but that the watermark is difficult to find. In most practical cases the retrieval of the watermark entails: 1) finding the proper scale and 2) subsequent retrieval of the payload. Retrieval of the scale parameters usually involves a large frame buffer. Moreover, latency is introduced due to the fact that the retrieval of the scale parameters is not immediate. As soon as the watermark detector is locked onto the proper scale, the payload of the watermark may be extracted.

An aspect of scale invariance that is not very often highlighted involves the false positive rate. Unmarked video will trigger the watermark detector to search for watermarks in a large number of scales. This means that the search space is effectively enlarged, and therefore the false positive rate is increased.

2) Intentional Tampering: The illegal copy without CSS of the Fig. 1 scenario was rendered unplayable by the watermarking technology in Fig. 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 [9], [12]. 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. This approach has been successfully implemented by Macrovision for their analog copy-protection system.

Beyond watermark removal, there are other ways to circumvent the copy-protection system. These include hardware modification to disable watermark detection, modification of the geometry of the source and source scrambling, the intention of the last two methods being to hide the presence of a watermark. A good example of the first case is the introduction of slowly varying horizontal and vertical offsets. In fact, jitter in horizontal offsets may be introduced unintentionally when video is played back from a low-quality VCR. This circumvention method is so cheap and easily available that a watermark system which is not shift invariant is unsuitable for copy-protection purposes. A more complicated geometric attack relies on scaling the video source. For source scrambling, the video must be descrambled after it passes by the watermark detector. However, inexpensive source scramblers and descramblers may be difficult to outlaw since users may argue that they serve a legitimate privacy purpose, e.g., preventing children from watching inappropriate content.

V. Copy-Generation Management

Copy-generation management requires that the "copy_once" state be detected and changed to a "copy_no_more" state as the video is being recorded. Two approaches have been proposed. In the first approach, copy-generation management is completely implemented in the watermark domain. In the second approach, the state change is effectuated by removal of additional information known as tickets.

A. Secondary Watermarks

A straightforward method of changing a content's state from "copy_once" to "copy_no_more" is to replace the "copy_once" watermark with a "copy_no_more" watermark. However, such a direct approach is not practical. First, the primary watermark that encodes the "copy_once" state is computationally expensive to insert, and it is therefore not economically feasible to include such an inserter in every DVD recorder. Second, the widespread availability of watermark encoders is a significant risk to its long-term integrity.

Alternatively, the primary watermark denoting "copy_once" can be left unchanged while a secondary watermark is added. The presence of both the primary and secondary watermarks denotes the "copy_no_more" state. Some of the requirements of the secondary watermark are quite different from the primary watermark. In particular, it is imperative that the secondary watermark inserter be computationally inexpensive. Insertion must be possible in both the baseband and compressed video domains. And when the MPEG stream is modified, it must be done so without changing the bit rate. Like the primary watermark, the secondary watermark should be unobtrusive and robust. There is disagreement between the authors as to whether these design constraints can all be satisfied.

A secondary watermark is likely to be more susceptible to tampering than the primary mark since the constraint on computation requires a less sophisticated insertion strategy. In addition, possession of both a watermark encoder and a watermark detector can simplify the problem of rendering the watermark undetectable. However, this approach has the advantage of not requiring any associated metadata. Thus, the solution is independent of the transmission channel, e.g., analog, digital, cable, satellite. This is especially important given the installed base of cable STB's and satellite receivers.

B. Tickets

Tickets represent a second solution to the tampering problem. A ticket T is a cryptographic counter, which is implemented as a multibit random number. The counter value $\langle T \rangle$ depends on the presence of a watermark W with a multibit payload $\langle W \rangle$ and is defined as the number n such that $F^n(T) = \langle W \rangle$. Here, $F(\cdot)$ denotes a fixed and secure cryptographic one-way function. The ticket value $\langle T \rangle$ represents the number of playback and recording generations allowed. The predominant requirement of the one-way function $F(\cdot)$ is that it is computationally unfeasible to compute the inverse, i.e., given the watermark, it is not practical to calculate a valid value for the ticket T. Such one-way hash functions are well known [13].

To provide generation management, the recorder/player modifies the ticket by feeding it through $F(\cdot)$, thereby effectively decrementing the ticket value $\langle T \rangle$. The ticket concept can be viewed as a cryptographically secured CGMS. Tickets can also be viewed as some sort of fragile watermarks: it is up to compliant devices to process tickets properly. Noncompliant devices, being ignorant about tickets, will in general loose the ticket, prohibiting re-entry of the content into the compliant world. As illustrated in Fig. 5, the ticket changes state during every passage through

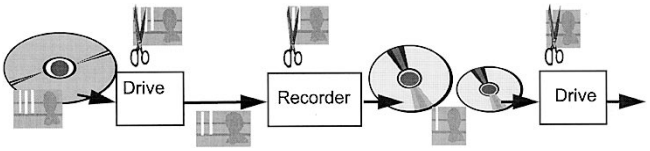

Fig. 5. The ticket is clipped (cryptographically modified) during each playback or recorder passage.

a playback and recording device. In other words, that ticket behaves as a counter that gets decremented every time it goes through a player or recorder and permits operation of this device as long as this counter is greater than zero. Thus, it becomes (computationally) impossible for attackers to increment the counter again. The ticket T in the stream is replaced by $T' = F(T)$ during each recording or playback operation, whereby $F(\cdot)$ is a publicly known cryptographic one-way function. Neither the player nor the recorder pass T transparently.

For a copy-never signal, T specifies that only playback is allowed. When in transit from a player or an STB to a recorder, such video signal carries a ticket for one passage.

The implementation of a ticket depends on the type of link or medium considered. For a copy-never signal, T specifies that only playback is allowed. For DVD-ROM disks it is proposed to implement the ticket by means of a wobble in the lead in. The wobble is a radial deviation of the position of pits and lands relative to the ideal spiral, which can be detected from the optical laser pickup or from signals in the servo feedback tracking loop. The wobble format is currently discussed in the DVD forum for formal approval. Wobbled disks are backward compatible with respect to current drives because the deviation is only very slight and the wobble frequency is sufficiently high not to disturb the positioning of the optical. The payload of the wobble is only a modest 64 bits, but this is sufficient to implement a ticket. A coarse description of playback control with a wobble is as follows. Upon insertion of a disk in a compliant drive, the drive will look for the presence of a wobble, and if present, read out the 64 bits of payload. If the (compliant) MPEG decoder informs the drive that a copy-never watermark is read, the drive: 1) feeds the 64 wobble bits through the one-way function F and 2) requests the MPEG decoder to read out the additional payload of the watermark. Only if the additional watermark payload and transformed wobble bits match is playback allowed. Because it is mechanically impossible to write a wobble track on recordable media, the wobble is a also powerful method for distinguishing ROM disks from recordable and rewritable disks.

For copy-once signals in transit from a player or an STB to a recorder, such video carries a ticket for two passages, i.e., one for recording and one for playback, so $F2(T) = W$. Several options exist to store the transformed ticket bits $F(T)$ on the recordable media to allow play control. We will mention two options. The first option, put forward by Hewlett-Packard, proposes to exploit redundancy in the DVD channel modulation code (i.e., EFMP). By slightly changing the run-length statistics of EFMP, a few bits per sector can be embedded. The second option proposes to induce intentional errors in the ECC blocks. Both of these methods create a secure and fragile channel because, similar to the wobble channel, it is inaccessible to the user. Moreover, both methods can easily be implemented in a DVD basic engine.

Tickets values also have to be passed from source devices to sink devices, and therefore tickets have to live on connecting links as well. As digital links between compliant devices are assumed to be secure, on digital links there are many options for secure transmission of ticket bits. The most problematic case is the transmission of tickets over analog links, as, for example, from an STB to a compliant DVD recorder. Several options have been put forward, such as transmission in the vertical blanking interval (VBI) transmission in the overscan area of active video or in the extended data services (XDS) of line 21, but no agreement has yet been reached by all involved parties on the best option.

The ticket concept is a powerful idea. However, the ticket itself is metadata that must be broadcast along with the video content. This is likely to mean that some existing broadcast channels will need modification or will not be able to support copy-generation management. This also means that every future broadcasting standard and all equipment for converting between standards must maintain the ticket data. The seriousness of this problem is another source of disagreement between the authors.

VI. Conclusion

The risks associated with the digital distribution of content have proven to be a strong incentive to develop a broad array of technologies to deter illegal copying. This has proven to be particularly true with the introduction of DVD video players and recorders. However, the design challenges are compounded by the need to establish a consensus between the often conflicting goals of the motion picture, consumer electronics, and computer industries.

The formation of the CPTWG provided a forum within which the three industry groups, the MPAA, CEMA, and IT, agreed to technical solutions to a variety of problems. Currently, DVD video disks are encrypted and an analog protection system is included to inhibit illegal VHS copies. A robust digital transmission protocol has been designed for the safe transfer of digital video between compliant consumer and computer devices and is well on the way to being adopted as a *de facto* standard. The CPTWG has made substantial progress toward the adoption of watermarking technology to prevent unauthorized copying of in-the-clear video, as well as playback of pirated disks. Once adopted, this will be the first large-scale deployment of video watermarking technology for copy protection.

Many issues still remain to be addressed, and the CPTWG has begun discussions on several of them. These include a system for encrypting legal copies on RAM disks, a method of preventing RGB to NTSC encoders from being used to make illegal VHS recordings from RGB output, and the development of a complete system for protecting audio content. In addition, discussions are beginning on the application of all the above system to new video formats such as HD/DTV.

REFERENCES

[1] A. Bell, private communication, May 15, 1998.
[2] Data Hiding Subgroup (DHSG). Call for proposals. [Online]. Available WWW: http://www.dvcc.com/dhsg.
[3] *Standard for a High Performance Serial Bus*, IEEE Standard 1394-1995, Dec. 12, 1995.
[4] C. Cookson and B. Lambert, "MPAA comments on the digital transmission findings document," presented at the January 1998 CPTWG Meeting, Burbank, CA.
[5] "Official call for proposals, digital transmission discussion group," presented at the March 1997 CPTWG Meeting, Burbank, CA. [Online]. Available WWW: http://www.dvcc.com/cptwg/dtcfp10.pdf.
[6] Digital Transmission Discussion Group Copy Protection Technical Working Group. (1997, Nov. 11). Review and findings of submitted proposals, version 1.0. [Online]. Available WWW: http://www.dvcc.com/ cptwg/dtdgr10.pdf.
[7] 5C digital transmission content protection white paper, version 1.0. [Online]. Available WWW: http://www.dtcp.com.
[8] *Editorial Contribution to Standard for Public Key Cryptography, Preliminary Draft, P1363/D3*, IEEE P1363, May 11, 1998.
[9] I. J. Cox and J.-P. Linnartz, "Some general methods for tampering with watermarks," *IEEE J. Select. Areas Commun.*, vol. 16, pp. 587–593, May 1998.
[10] J. R. Hernández, F. Pérez-González, J. M. Rodríguez, and G. Nieto, "Performance analysis of a 2-D-multipulse amplitude modulation scheme for data hiding and watermarking of still images," *IEEE J. Select. Areas Commun.*, vol. 16, pp. 510–524, May 1998.
[11] G. F. G. Depovere, A. C. C. Kalker, and J. P. M. G. Linnartz, "Improved watermark detection reliability using filtering before correlation," in *Proc. Int. Conf. Image Processing (ICIP)*, Chicago, IL, Oct. 1998, pp. 430–434.
[12] T. Kalker, "Watermark detection through detector analysis," submitted for publication.
[13] B. Schneier, *Applied Cryptography: Protocols, Algorithms, and Source Code in C*. New York: Wiley, 1996.

Jeffrey A. Bloom received the B.S. and M.S. degrees in electrical engineering from Worcester Polytechnic Institute, Worcester, MA, in 1987 and 1990, respectively. He is currently pursuing the Ph.D. degree in the Electrical Engineering Department, University of California, Davis.

From 1985 to 1986, he worked for Eastman Kodak Company as an Engineer in their Quality Assurance Testing Group. In 1998, he joined Signafy, Inc., as a Member of the Technical Staff researching issues in the watermarking of image, video, and audio data.

Ingemar J. Cox (Senior Member, IEEE), for a photograph and biography, see this issue, p. 1141.

Ton Kalker (Member, IEEE) was born in The Netherlands in 1956. He received the M.S. degree in mathematics in 1979 from the University of Leiden, The Netherlands, and the Ph.D. degree in mathematics in 1996 from the Technical University of Delft, The Netherlands.

From 1979 to 1983, he worked as a Research Assistant at the University of Leiden. From 1983 to December 1985, he worked as a Lecturer at the Computer Science Department of the Technical University of Delft. In December 1985, he joined Philips Research Laboratories, Eindhovenn, The Netherlands. Until January 1990, he worked in the field of computer-aided design. He specialized in (semi-) automatic tools for system verification. Currently, he is a member of the Digital Signal Processing Group of Philips Research. His research interests include wavelets, multirate signal processing, motion estimation, psycho physics, digital video compression, digital watermarking, and multimedia security.

Jean-Paul M. G. Linnartz (Member, IEEE) received the Ir. (M.Sc.E.E.) degree (cum laude) in electrical engineering from Eindhoven University of Technology, Eindhoven, The Netherlands, in 1986 and the Ph.D. degree (cum laude) from Delft University of Technology, Delft, The Netherlands, in 1991 for his work on multi-user mobile radio nets.

He is with Philips Natuurkundig Laboratorium (Nat.Lab.), Eindhoven, The Netherlands, where he leads a team of researchers developing multimedia conditional access and smart card security system. In 1994, he was with Delft University of Technology as an Associate Professor. In 1992–1994, he was an Assistant Professor with the Department of E.E.C.S., the University of California, Berkeley, where he was Assistant Adjunct Professor from 1994 to 1998. From 1988 to 1991, he was Assistant Professor at Delft University of Technology. During 1987–1988, he worked with the Physics and Electronics Laboratory (F.E.L-T.N.O., The Hague) of The Netherlands Organization for Applied Scientific Research on frequency planning and UHF propagation. He is the author of the book *Narrowband Land-Mobile Radio Networks*. His main research interests are in conditional access and information security, electronic watermarks, (wireless) multimedia communications, and multicarrier CDMA (combining OFDM and CDMA).

Dr. Linnartz received the Dutch "Veder Prize" in 1991 for his research on teletraffic aspects in mobile radio networks. He is Editor-in-Chief of *Wireless Communications*, the interactive multimedia CD-ROM.

Matthew L. Miller, for a photograph and biography, see this issue, p. 1141.

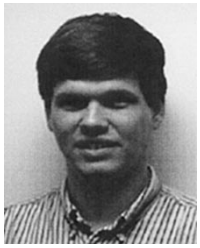

C. Brendan S. Traw received the Ph.D. degree in computer and information science from the University of Pennsylvania, Philadelphia.

He is currently a Staff System Architect with Intel's Platform Architecture Laboratory, Hillsboro, OR. His research focus is on the architecture and performance of I/O subsystems with a current emphasis on protecting entertainment content from unauthorized copying. He has published over 12 papers and has been granted two patents.

Dr. Traw is a member of ACM.