The "ticket" concept for copy control based on embedded signalling[1]

J.P.M.G. Linnartz
Philips Research, WY8, Holstlaan 4
5656 AA Eindhoven, The Netherlands
linnartz@natlab.research.philips.com

***Abstract* - This application-oriented paper discusses the use of electronic watermarks (also called embedded signaling) for copy control. Playback Control and Copy-Once are described. The 'ticket concept' is presented to provide these functionalities. Although the ticket shows similarities with a digital signature, there are essential differences. For instance, the ticket allows typical modifications of the content, which are common practice in transmission, storage and presentation of video. The concept is part of a proposal currently under investigation for standardization of DVD / CPTWG copy control. This paper also compares the ticket concept with other solutions, such as embedding a secondary mark at the recorder and using a signature scheme**

Key words: Protection of IPR, watermarking, embedded signaling, copy control, copy once, playback control, Copy Generation Management System (CGMS), Digital Versatile Disc (DVD).

1. Introduction

While digital multimedia technology opens many opportunities for new applications and services, content owners are afraid to loose revenues as copies of digital content can be generated rapidly, perfectly, at large scale and without limitations to the number of generations of copies. As copy management involves demanding and conflicting requirements, the issue has come on the "critical time path" of the market introduction of several digital products, including the Digital Video Broadcasting (DVB), Digital Versatile Disc (DVD), the IEEE 1394 digital interface and improved digital audio carriers (e.g. Super-Audio CD). The image quality of digital video disks provides a significant improvement over the quality of existing home video equipment, such as VHS recorders. For content providers, there is a greater risk of illegal copying using perfect digital reproduction. The problem of protection against illegal copying has been recognised decades ago, but adequate solutions have not yet been found. The first aim of copy protection schemes is to prevent illegal copies from being made. Failing that, the aim is to reduce the value of illegal copies, either by reducing their quality (hopefully to the point of being unwatchable) or by restricting their use. Copy protection critically determines the viability of many business concepts in the Information Society, and it is receiving increasing pressure to find better solutions. Tools for copy protection in the digital world are sought in two directions: cryptography and embedded signalling.

[1] J.P.M.G. Linnartz, "The ticket concept for copy control based on embedded signalling", ESORICS '98, 5th. European Symposium on research in Computer Security, Louvain-La-Neuve, September 1998, Lecture Notes in Computer Science, 1485, Springer, pp. 257-274.

The old cryptographic paradigm in which Alice communicates with Bob and is wiretapped by Eve, does not hold for copy protection. Rather Alice wants to sell information to an unreliable Bob, but meanwhile she wants to restrict the use of that content. At the same time, Bob must be able to use and copy any similar works of art that he created himself, without any restrictions. It can easily be understood that encryption, as for instance applied to DVD disc sectors, only addresses part of the issue of illegal copying. It avoids that the user has direct access to valuable, highly compressed digital content. Instead, the user (or more precisely his electronic device) must make use of patented decryption algorithms. Conformance to copy protection measures is enforced in licensing contracts. Often such contracts state that a playback device may only offer the content in analogue form to the end user.

Watermarking is a another technique, not only useful for active protection of intellectual property against illegal copying but also for new multimedia trading mechanisms. Yet there are several technical issues to be resolved. In the process of determining a watermarking standard for enhanced DVD copy protection, the "copy-once" and "conditional playback" became important requirements. This has lead to the development of methods to signal and dynamically modify the copy state of watermarked content, which we cover in this paper.

The outline of this paper is as follows: Section 2 discusses the DVD embedded signalling standardisation requirements. Section 3 addresses play control and copy generation management in a broader context. Section 4 covers the proposed 'ticket' solution, and applies it to play control and generation control, in Section 4.1 and 4.2, respectively. Two other solutions have been proposed for the DVD copy-once requirement, namely the secondary watermark, and the digital signature, which are discussed in Section 5 and 6, respectively. Section 7 covers some generic attacks and potential weaknesses of these systems and Section 8 compares the solutions. Section 9 concludes this paper.

2. DVD Copy Control[2]

DVD video material is encrypted on disc and the decryption keys are stored in a manner such that an ordinary copier who does not have access to the internal hardware of the disc drive can neither read or write these. Moreover, the decryption, MPEG decoding and D/A conversion are conducted in a more or less tamper resistant environment. This is easier to achieve in a consumer electronic device that is closed black box, but more difficult for personal computers.

Hence an attacker has difficulty in getting the digital plaintext. Moreover, if he has access to the ciphertext only, he cannot make an encrypted copy that also has the appropriate keys. A digital copy of the encrypted content will not play unless the keys are also copied. To ensure exportability, the key length has been restricted, which exposes the system to cryptanalytic attacks. An important aspect of this approach is that its places a group of 'compliant devices' in the market, which internally adhere to copyright rules and externally communicate over protected links (e.g. using the IEEE 1394 interface). Lacking globally uniform copyright laws, licensing of cryptographic technology is essential to enforce the compliance to copy

[2] This section heavily relies on an earlier paper by Cox and Linnartz [1]

control rules. Hence, we observe that cryptography is used more as a tool to bind manufactures to copyright rules than as a copy protection mechanism by itself.

Another possible weakness is the analogue signal. Content must eventually be converted into an analog form to present it to a human viewer. Neither cryptography nor licensing contracts protect these analog signals. The in-the-clear video signal is available at a variety of interfaces, including the NTSC or the RGB signal output. To prevent analog copying, DVD players are equipped with an analog protection system (APS). This is a proprietary technology developed by Macrovision that modifies the generated NTSC signal such that most VHS video recorders cannot record a high quality copy despite the fact that the same signal does not affect the TV display. Unfortunately, this system does not protect RGB signals, which are common to PC's and in the Europe SCART connectors, from analog recording and is therefore easily circumventable. Thus, copyrighted video material will find its way into the analog domain.

Considering the level protection to digital content, the most likely method by which a causal copier attempts to make a digital copy is through the digitization of an analog signal. Neither encryption nor the APS signaling prevent playback or recording of this illegal copy, unless A/D video grabbers are equipped with APS detectors, to voluntarily disallow digitization of APS formatted video signals. To provide enhanced protection, a watermark is inserted into the copy-restricted video sequence. It is intended to prevent illegal copying by telling a compliant device not to copy it. Hence, the watermark should survive MPEG-2 compression and digital-to-analog-to-digital conversions, i.e. if the video fidelity remains high, then the watermark should remain detectable. It can also reduce the value of illegal copies by preventing them from being *played* on compliant devices. This means that consumers will have a choice between a) compliant devices, which can play legal, commercially released discs that were encrypted, but cannot play pirated discs, and b) non-compliant devices, which can play pirated material, but cannot play encrypted discs.

According to requirements formulated by the DVD Copy Protection Technical Working Group, all possible (video)-content should fall into one of four categories:

- **Free Copy** (no copy restrictions whatsoever, e.g. home productions)
- **Copy Never** (no copies allowed; e.g. for DVD titles or for films rented from a video store). Despite the legislation in many countries that a customer is allowed to create backups for personal use, the content industry currently intends to predominantly publish content as "never copy".
- **One Copy Allowed** (one *generation* of copies may be made). Users will be able to make a digital copy, but the system should prevent copies of this copy (or subsequent generations of copies) being made. Applications are primarily for Time Shift Viewing, that is, recording a feature film to be watched at a later time. There is no limit to the number of first generation copies that can be made from the original content. However, this original typically is only available once on the digital output of a Pay TV Set Top Box. Making a backup of purchased content is another legal application, at least in those countries where a Copy Never status is unlawful. According to several court cases, the user implicitly buys the right to copy if royalty fees are being levied on blank recording media.
- **Copy No More** the copy-state of a recording after being copied a first generation.

Copyright data describing the restrictions on that video's usage should not only signalling of the copy state, but also it should trigger the APS system. Instruction bits for usage of the analog protection system (APS) are

00 Don't use APS
01 Use type 1 APS
10 Use type 2 APS
11 Use type 3 APS

In the past the Copy Generation Management System (CGMS) was introduced to provide copy-once functionality. In order to implement the copy once functionality of CGMS, it will probably be necessary to have one or more additional bits in the watermark that can easily be changed by consumer recorders. Ignoring certain extensions, the CGMS bits are

00 Video may be copied freely
01 not used
10 Video may be copied once
11 Video may never be copied

A well known weakness of CGMS is that hackers can easily flip a copy state bit to grant themselves the possibility to copy. Even worse, cheap black-box devices can be sold on the market that set the copy state of any content to "free copyable", though these have recently been outlawed in the U.S.

Embedded signalling should strengthen APS trigger data and generation management. The Copy Protection Technical Working Group (CPTWG) initiated the Data Hiding SubGroup (DHSG), which released a call for proposals in May 1997 [3]. The requirements placed on watermarking algorithms for the above application differ from those for other applications that are currently in the market, such as identification of ownership. The application of watermarking for copy protection requires a low bit-rate and allows the use of many frames for watermark detection. However, since watermark detectors must be built into millions of low-cost, consumer devices, and since these detectors must work at video rates, there is a very strong requirement that the detector be extremely simple and cheap. Furthermore, since the DVD standard employs MPEG coding, the watermarking method must work well with MPEG. These last two requirements are challenging design specifications.

The requirements for the playback control and copy generation system are:

- Detectable in the baseband and/or compressed video
- Detector should be very inexpensive both in terms of gate count (hardware) or MIPS (software). Typically, the detector should be implementable in "a few thousand gates", although common belief is that "a few tens of thousands of gates" are needed to satisfy the basic requirements.
- Extremely low false positive rate. Consumer equipment may not fail to work because of an erroneously trigger watermark detector.
- Detection in 10 seconds or faster.
- No visible artifacts, i.e. very high image fidelity
- Tamper resistant, i.e. it should not be easily circumvented or removed.
- The watermark should survive color representation conversion from YUV to RGB
- Low data rate
- Support of generation control

As some of the above requirements appeared a challenge to designers, other requirements were relaxed, in particular those addressing robustness to image transformations. Nonetheless, the watermark should also survive:

- Compression
- Decompression
- Digital-to-Analog
- Analog-to-Digital
- Standards conversion, e.g. analog video recorder (VHS)

Several solutions have been submitted in response to the DVD Call for Proposals [3], including systems designed by Hitachi, HP, IBM, Macrovision/Digimarc, NEC, SONY, Philips and Pioneer. While the original CFP only recognised the embedding of a secondary watermark as the technical solution for the generation control, other solutions surfaced. Ignoring some details that address the coexistence of encryption and watermarking, we distinguish three categories of approaches:

1. Embedding a secondary watermark when the copy is made. The NEC proposal follows this approach.
2. The video that may be copied contains a permission ticket, which is stripped by the recorder. The Philips proposal follows this approach.
3. The video that may be copied contains a signature-type of permission "token", which is stripped by the recorder. The signature solution resembles the ticket concept in that both add a signal to the content that can be removed easily. The IBM proposal follows this approach.

In the following sections, we present the design rationale behind the ticket concept and compare it with the other solutions.

3 Record and Play Control

Formally speaking, this copy protection approach relies on a form of *public-key* watermarking, i.e., user devices must be able to read the watermark, but this should not reveal how a watermark can be erased. Such schemes have not yet been found [1, 2] and most currently proposed systems do formally not satisfy this requirement. Current thoughts (as for instance expressed in DVD / CPTWG / DHSG standardisation) are that the watermark detector is embedded as a tamper resistant element of the electronic chip which controls the record and playback engine. In the remainder of this article, we assume that an appropriate watermarking scheme exists such that consumer devices can verify the watermark, but not erase the watermark from the content. Watermarks can only be embedded by agencies that have access to the embedding (algorithm and its) keys. However, Linnartz et al. have shown that such solution potentially is vulnerable to a sensitivity attack [6], in which the tamper-proof detector box is used as an oracle that reveals up to one bit of information about the watermark secret per experiment.

The most basic and most common approach is record control. The recording device detects the presence of a watermark and inhibits copying this content. Record control prevents a casual consumer from copying copy-protected, i.e., watermarked material onto a recording device.

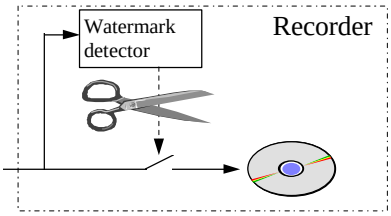

Figure 1: *Fundamental weakness of Record Control.*

A pirate who is interested in illegal copying may not only attempt to tamper with the watermarked image, but can also attempt to circumvent the copy control mechanism while leaving the watermarked content unchanged. The most trivial attack is to tamper with the output of the watermark detector and modify it in such a way that the copy control mechanism always sees a "no watermark" situation, even if a watermark is present in the content (Figure 1). Since hackers and pirates more easily can modify (their own!) recorders but not their customers' players, playback control is a mechanism that detects watermarks during the playback of discs. The resulting tape or disc can be recognized as an illegal copy.

An essential strengthening of the system is to prohibit *playback* of content if it appears to be an illegal copy. In its simplest from, watermark content is played back only if it comes from original media (e.g. stamped 'silver' discs), but playback is rejected if the watermarked content is played from recordable media (e.g. 'golden' discs). The player recognises the copy state of the content, e.g. by detecting the watermark and comparing this with a physical mark on the disc. Only if the carrier's physical properties correctly match with the watermark, the device is authorised to play.

If the above hacker would illegally copy discs by modifying a recorder, the watermark will remain in the content and the playback drive of his customer detects the fraud. For the pirate it is usually economically and technically not feasible to modify the installed base of players of all his potential customers. In play control, the medium must contain a *physical mark* which at least distinguishes between ROM and RAM discs, but a stronger protection that also eliminates counterfeiting of ROM discs is to be preferred. According to as recent market investigation was economically more attractive to publish on stamped (silver) CD ROM instead of (golden) CD recordable discs, if the batch size is above some 200 discs. Price erosion in CD-R has recently shifted this turnover point, but prices of disc are likely drop as well. Hence, piracy of CD or DVD-ROM discs is not adequately countered if the watermark detector in consumer electronics players only distinguishes between RAM and ROM DVD discs, without checking the relation between the origin of the ROM medium and its content in a more sophisticated manner.

Small-scale pirates typically do not run their own ROM disc-pressing equipment, but they have access to commercial pressing facilities. In particular, a protection mechanism that requires a physical modification to the disc press machine effectively hinders many small-scale pirates who do not own the press plant themselves. Enhanced protection can be provided against an attacker who only has access to consumer or unmodified professional equipment (but not to dedicated reverse engineering tools) by making the bit contents of the physical mark inaccessible. Moreover, knowledge of the bit contents should not allow the attacker to press copies of a discs on equipment operated by an commercial company with unmodified pressing machines.

4 Ticket concept in play control

In abstract terms, the ticket method (which is the prime focus of this paper) addresses a method to associate a message (which we shall call the ticket) to a watermarked piece of content (say, an image or motion picture), such that

- the recipient of the image can detect with high reliability whether or not such an associated ticket file has been issued by the copyright owner
- if the recipient has access to a ticket that is presented as being associated with an image, this recipient can verify its integrity and authenticity. That is, he can verify with high reliability whether this ticket is the correct one, that truly belong to this particular image and has been issued by the content owner, and
- the above two properties still hold after typical signal processing operations (e.g. MPEG compression) have been performed on the content.

The third requirement is at odds with the property of 'hard' *digital signatures* (see Section 6).

To provide generation management, we modify the associated data in a computationally irreversible manner. Our ticket concept can be viewed as a cryptographically secured Copy Generation Management System (CGMS). (See Section 4.2). Moreover, as explained in the next section, it can also strengthen the play control protection.

4.1 Ticket concept in play control

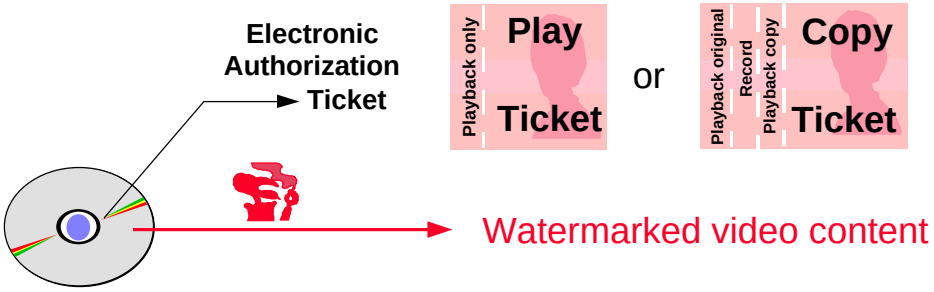

Figure 2: Basic elements of play control: if a player detects watermarks, it checks the presence of an appropriate authorisation mark.

In its basic implementation, play control allows playback of watermarked content from pressed (silver) discs, but not from recordable or rewritable (silver or green) discs. As we argued before a stronger, say cryptographic mechanism is needed to relate the content to the physical carrier.

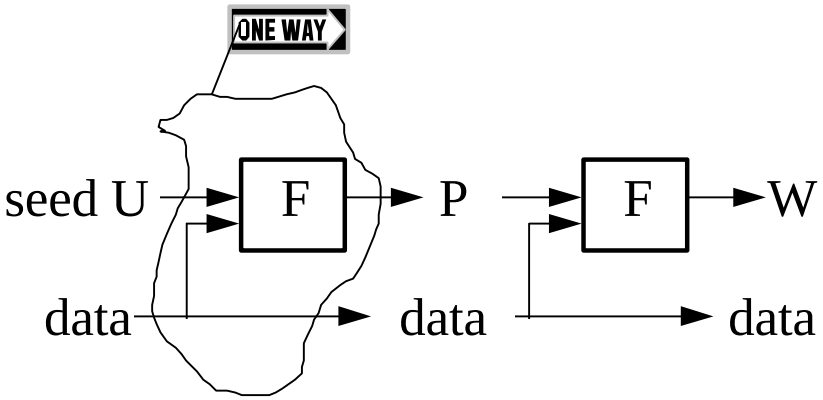

Figure 3*: Relation between copyright data, random seed U, Physical mark P and watermark W. F is a cryptographic one-way (hash) function.*

In the ticket concept, the physical mark carries bit string *P* which is related to the watermark *W* according to

$$W = F(P),$$

where $y = F(x)$ denotes a cryptographic one-way function. The bit content of the physical mark is embedded on the disc but cannot be read or recovered externally of the drive. During mastering, the physical mark and watermark are created from a seed *U*, according to

$$P = F(U) \text{ and } W = F(F(U)),$$

The predominant requirement of the one-way function $F()$ is that it is computationally unfeasible to compute the inverse. Computing an inverse means finding which particular x_0 leads to a given y_0 with $y_0 = F(x_0)$. With unfeasible we mean that the most efficient method to find such x_0 is to exhaustively search all possible bit combinations of x_0 and to compute and verify $F(x_0)$ for each attempt. The one-way function is calculated within the drive and recorder, so simple hardware solutions are preferred. From a security point of view, the one-way function itself may not need to be kept secret. No other system secret is present in the consumer products, except the secret behind the watermark detection process.

One suitable realisation of a physical mark is the *wobble* groove in optical discs. The basic principle of this concept is covered in the Orange Book standard for CD-R [4]. It is a superimposed small, periodic variation on the normally linear variation of the radius of the spiral with data on the disc. The wobble bit content cannot be recovered from the output of the disc drive, but it can be detected from the control circuit which stabilises the optical pick-up above the track. The wobble is too fast to be tracked by the servo motor, or its control current, but the wobble signal is present in the feedback loop that positions the pick-up. Wobble marks can be inserted on pressed discs, but custom writing these on a recordable disc with normal writers is physically impossible. The other requirements are satisfied by performing a cryptographic one-way function in the control hardware.

4.2. Ticket concept in generation control

To allow one-copy from disc, a special form of playback occurs. The drive must pass a copy ticket to its output. This ticket must allow a recorder to copy the content and a next player to play the copied content. After these transitions, no further copies must be possible. Requirements are

- It should be difficult to retrieve the one-copy-allowed value of that mark from copy-no-more content.
- The consumer equipment should not carry a secret which reveals how tickets can be generated for existing watermarked content. Only the content owner can generate a ticket. Users only can 'clip' the ticket.
- Content may undergo transformations, during which the copy state must be preserved.

As illustrated in Figure 3, the ticket changes state during every passage through 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. We implement this state changing process in such a way that it becomes (computationally)

impossible for attackers to `increment' the counter again. We may do this using the cryptographic one-way function from the previous section: initially *T* is some multi-bit number, and during every passage, we apply *F(.)* to this number and call the result the new ticket.

The ticket is a volatile piece of data which can be stored and transferred either in *embedded* or *associated* form. This distinction is particularly relevant for internal processing of the signal on platforms such as the PC, where only embedded signals will be retained and associated data can easily be lost.
In storage, the ticket will typically be *associated* as a physical marker of the storage medium. That is, the ticket is stored at locations that are inaccessible by normal hardware products, thus separated from the content. The wobble is a typical example, suitable for professionally released content on stamped discs. For recordable media, potential methods to carry the ticket are by introducing intentional bit errors in a predetermined manner, or to modulate EFMP modulation codes which determine the relation between user bits and pit and lands on the disc [5].
In transit, the ticket is *embedded* into the signal. Examples are MPEG `user_data` bits, or data in the blanking intervals of the PAL and NTSC television standard.

For a never-copy signal, *T* specifies that only playback is allowed, i.e., the ticket reduces to the physical mark as discussed in Section 4.1. A "one-copy" signal on a professionally-produced disc carries a ticket for 3 passages: playback, recording, followed by another playback. When in transit from a player or a Set Top Box to a recorder, such the video signal carries a ticket for 2 passages.

Ticket *T* in the stream is replaced by *T' = F(T)* during each recording or playback operation, whereby *F(.)* is a publicly known cryptographic one-way function. Neither the player nor the recorder pass *T* transparently. The system does not rely on a global secret.

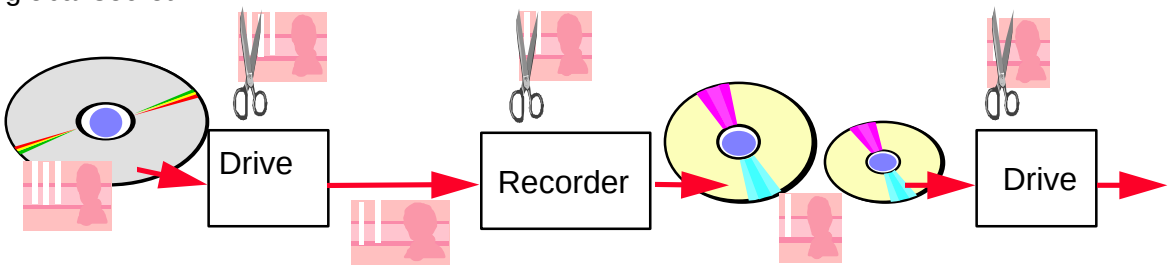

Figure 4: *The ticket is clipped (cryptographically modified) during each playback or recorder passage.*

Playback

Playback will only take place if one of the following conditions is met:

- The disc contains a ticket in the form of a physical mark *P* reserved for stamped media. The content on the disc contains a watermark *W*. The player further checks the validation ticket. One of the following conditions must be satisfied:
  "Never-copy": The relation *W = F(P)* is satisfied.
  "One-copy": *T* is present and *W = F(F(F(T)))* is satisfied. In this case, *F(T)* is made available at the output of the drive.

- The disc contains a physical mark *P* reserved for recordable media. The content contains a valid *W* watermark used for professional recording. The validated one-copy *T* is present, and *W = F(T).*
- The disc contains a physical mark P reserved for recordable media. The content is identified as a home recording of a user's personal creation by checking the absence of a watermark.

Recording
Recording of copyrighted content is allowed only if the watermark in the stream matches *W = F(F(T)).* The recorder always passes the copy control ticket through the one-way function before transferring it to disc. If an attacker manages to modify his recorder and record video even if the appropriate *T* is not present, a normal player will reject to playback the disc.

Professional Publishing
A professional title is produced by initially generating a seed *U*. From this seed, the following variables are computed: *P = F(U),* and *T = F(F(U))).* The one-way function *F* and variable *P* is specified such that *P* also contains an identifier for the publisher or a serial number of the mastering machine. If a pirate publisher attempts to write a particular *P*, in order to make a bit-exact copy of a copyright disc, that pirate must know *U* (which he cannot learn from the store bought product) or tamper with his DVD press.

Summary of Authorisation
The control of the ticket and physical mark is summarised in Figure 5. The content starts with generating a seed. From this seed physical marks, tickets and watermarks are obtained.

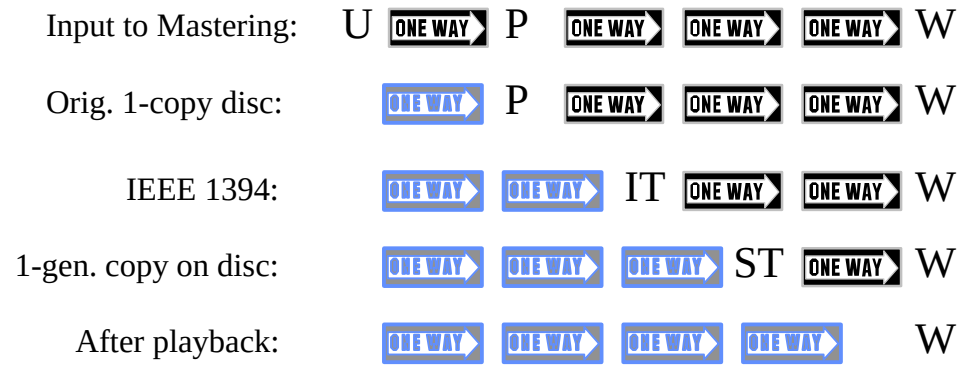

Figure 5: *Relation between Physical Mark, Watermark and corresponding authorisation Ticket on a medium indicated in the first column. Black one-way functions are to be verified. Grey one-way functions have been applied in the past, and cannot be undone.*

5. The Secondary Watermark Solution

In the concept formulated in the Call for Proposals of the DVD CPTWG [3], copy generation control was suggested to be performed by embedding a *secondary watermark* in the recorder. A recorder would accept to record content with a primary

watermark only, but would not accept content that has both a primary and a secondary watermark.

This approach has the disadvantage that consumer recorder must be able to embed a watermark. This implies that content must made accessible in a form that allows embedding. Typically at least partial MPEG decoding would be required, even if MPEG decoding would neither be required for watermark detection or for the recording function. Reliable embedding, thus with sufficiently strong watermark energy presumably must the adapted to image properties so it requires evaluation of the video using a perceptual model. Maes [8] argued that fixed depth watermarks are also very sensitive to "Twin Peak Histogram Attacks". Section 7 will address some other security aspects associated with this scenario.

6. The signature Solution

In another context (e.g. [7]) it has been proposed to make it mandatory that any digital content, watermarked or not, should be accompanied by a digital signature from an authorized agent. A problem with such a copy control concept is the potential leak that occurs because private users must be able to create, store and copy their own personal works of art. Such works must then also be signed. Hence, any user can sign content and can attempt to sign and create illegal copy permissions for copyrighted content. In this paper we address scenarios in which signatures are required only if a watermark is present.

In this solution, content is watermarked if copy restrictions apply. Moreover, a hash function is performed over the digital representation of copy-once content and this hash is signed, i.e., encrypted[3] by the content owner. Consumer devices are allowed to copy the watermarked content only if this signature is present. When such devices make a copy, the signature is removed. Two enhancements of this proposal exist

1. Soft signatures:

Since signal processing ruins the validity of the signature, the recipient would not be able to verify the integrity of copy state of processed video. In practice, such modifications often occur, for instance during conversion from one format to another (e.g. digital to analogue, U.S. NTSC into European PAL television standard), transmission using lossy compression such as JPEG or MPEG. It would be inappropriate if the user looses permission to copy due to legitimate processing. Hence, the integrity of the ticket must be verified, even if the image undergoes permissible modifications. Therefor, the hash function is not performed over the digital MPEG representation directly, but over a set of extracted features of the images. These features are chosen such that these are unlikely to change in typical signal processing operations.

2. Countermeasure to replay of the ticket by intentional modification of the content: An attacker can grab the ticket and store it separately and use this to make a second generation copy. To invalidate the ticket and to avoid that the ticket can be misused to restore the one-copy-allowed state, the video can be modified intentionally by a recorder such that the feature extraction mentioned in the previous paragraph yields a different result.

[3] To satisfy complexity requirements, a symmetric encryption algorithm was proposed to the CPTWG.

7. Weaknesses and potential attacks

Watermarking weaknesses

All solutions addressed here have one common vulnerability, namely that the copy protection is lost if the attacker can tamper with the primary watermark. The attacker can attempt to transform (scale, tilt) the image such that the watermark detector is not triggered. Moreover, he can try to find the watermark secret and erase the watermark pattern. The presence of a watermark detector in every consumer device can be exploited by the attacker to estimate the watermark pattern using the sensitivity attack. See [1, 2] for an overview of attacks and vulnerabilities.

In the case of the secondary marking method, an additional weakness occurs because the consumer has access to a watermark embedder, though possibly only in tamperproof encapsulation. Thus, attackers can experiment by the marking random inputs. Next, we will argue that the embedder must have properties that dictate substantial linearity. Thus there is a conflict between robustness and security requirements, which may weaken the security of the secondary marking method.

Mathematically, given an image *I* and a watermark *W*, the watermarked image *I'* is formed by $I' = I + f(I,W)$ such that the perceptual differences between *I* and *I'* are less than "just noticeable", i.e. the watermarked image is constrained to be visually identical (or very similar) to the original unwatermarked image. In theory, the function *f* may be arbitrary, but in practice robustness requirements pose constraints on how *f* can be chosen. One requirement is that watermarking has to be robust to random noise addition. Therefore many watermark designers opt for a scheme in which image *I* will result in approximately the same watermark as a slightly altered image $I + \varepsilon$. In such cases $f(I,W) \approx f(I + e,W)$. If the recorder changes the watermark status from "one-copy allowed" to "no more copies allowed" by embedding a watermark, .the attacker has access to the content before and after this marking. That is, he can create a difference image $f(I,W)$, by subtracting the unmarked original from the marked content. An obvious attack is to pre-distort the original to undo the mark addition in the embedder. That is, the attacker computes $I - f(I,W)$ and hopes that after embedding of the watermark, the recorder stores

$$I - f(I,W) + f(I-f(I,W),W)$$

which is likely to approximate $I - f(I,W) + f(I + e,,W) \approx I$ because watermarks are small modifications themselves, thus $f(I,W) \gg f(I-f(I,W),W)$.
Although some countermeasure exist, many schemes are vulnerable to this attack or sophistications of it. Detection of the watermark *W* is often achieved by correlating the watermarked image with a locally stored reference copy of the watermark. Such correlator is vulnerable to the pre-subtraction attack discussed here. Note that the attacker can not only create a copy that plays on his own or his customers recorders, he can also sell generic circumvention devices.

Scrambled recording attack

When evaluating the security of the various solutions, it is relevant to consider attacks that appear so fundamental that these are unlikely to be solved by any system. Most importantly it appears unlikely that protective measures can be found to avoid that a hacker can build a storage device that stores 'random bits'. For

instance, a hacker can take his copyrighted video sequence and create a copy for his personal use by weakly encrypting all bits. The resulting sequence of 'random bits' can be stored on his device. 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 simple 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 and manufacturers might argue that the devices serve a legitimate purpose in protecting a user's privacy. Similarly, digital MPEG can easily be converted into a file of seemingly random bits. 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.

Steganographic recording

Moreover, it would not make sense to outlaws format non-cognizant bit-copiers because a more subtle circumvention of the copy control mechanism can be used. This attack exploits the technique of steganography (or data hiding) to bypass the watermark detector in the recorder. The method of attack is similar to the argument that laws against the private use of cryptographic encryption can be evaded by steganography. The copyrighted work is hidden in an innocent-looking file of a known and recognized format. For instance the digital MPEG video representation allows the user to embed additional user_data without any significant limitation. Stuff bits may be misused by a pirate to embed a complete stolen MPEG video film as user_data of another video sequence. During playback, the playback platform, e.g. the PC must perform a few additional functions, but this does not need to cause significant performance problems.

Thus, any play-control system can be circumvented by a pirate who can insert functionality (e.g. de-scrambling) between the drive and the MPEG decoder, or between de MPEG decoder and the display. Hence, in the evaluation of the various solutions for generation control, we must conclude that systems perform equally in this regard. This makes replay attacks of the ticker or signature less interesting, as the circumvention devices is more sophisticated than a generic black-box, because it needs to store data for every piece of content that is copied.

8. Comparison

Cryptographically, signatures (or "tokens") and tickets behave essentially different their vulnerability to collisions. Collisions can not be exploited to attack the ticket. The philosophy of the ticket purely rests in the difficulty to modify the watermark The attacker has to find a ticket value that matches with this watermark. The one-way function is designed such that finding any input signal for a *given* output (i.e., an indelible watermark) is not feasible with reasonable amount of hardware and within limited time. This effort would allow him to copy one title.

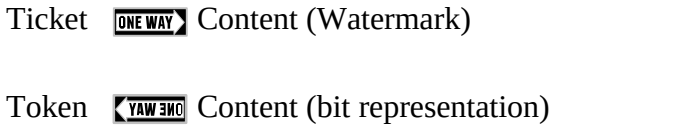

Figure 5: The *conceptual difference between signature and ticket solution rests partly in direction of one-way operation, and in whether the bit representation or the watermark is considered.*

It is critically important to observe that we use the ticket as *input* to our one-way or hash operation, and that the watermark challenges the *output* of the one-way function (Figure 5). Through this design the ticket is invulnerable to potential collisions of the one-way function. This concept is essentially different from a proposal based on signatures, in which the content is used as input and the authorisation token is a result of the output of a cryptographic function. In such scenarios, an attacker can attempt to modify the content slightly, in the hope that he finds it matching with a particular token.

We summarize relevant differences between the approaches in Table 1.

Table 1: comparison of various solutions for generation control

	Secondary mark	Hard Signature	Soft Signature	Ticket
Reliability: permission to copy remains after processing the video	**Yes**	**No**	**Partially, highly depends on type of transformation and design of feature extraction**	**Yes**
Vulnerable to replay attacks of associated tag	**N/A**	**Yes**	**Yes, but recorder can distort image a little to render tag invalid**	**Yes**
Vulnerable to comparison of input and output of recorder	**Yes**	**No**	**No, provided that the soft signature is cryptographically strong**	**No**
Embedded secrets in recorder, besides watermark detection	**secondary watermark detector, watermark embedder**	**signature verification keys**	**signature verification keys, feature extractor**	**none**
System secrets distributed among content providers, besides primary watermark		**signature generation method**	**signature generation method**	**each providers keeps his own secret seed. No global secrets**
Protection against piracy	**no**	**no**	**no**	**Yes, pirate must modify mastering**

	Secondary mark	Hard Signature	Soft Signature	Ticket
via ROM media				**machines**
Vulnerable to scrambled recording attack	**yes**	**yes**	**yes**	**yes**
Vulnerable to black-box that modifies state back to copy once	**Box must strip secondary mark, box designer must know secondary watermark secret**	**Counterfeit ticket can be generated by any malicious content provider who has a licence to release content**	**Counterfeit ticket can be generated by any malicious content provider who has a licence to release content**	**Counterfeit ticketing requires inversion of one-way function**
Perceptual artefacts	**Secondary embedder must work with limited complexity, artefacts are more difficult to avoid than with primary mark**	**none**	**Might occur if image is modified to avoid replay of signature**	**none**
Complexity in CE product, in addition to primary watermark detector	**secondary watermark detector and embedder**	**hash function, encryption function**	**Image feature extractor, hash function, encryption function**	**one-way function**
Required watermark payload	**2 bits**	**1 bit**	**1 bit**	**40-64 bits, but may be signalled at low rate**
Ability to support both record and play control	**record control only**	**record control only**	**record control only**	**Both**

9. Concluding Remarks

The introduction of the Digital Versatile Disc (DVD) has initiated a substantial effort in enhanced copy control mechanisms. Presumably it will be the first mass-market use of embedded signaling and electronic watermarking. This paper has covered some system concept aspects that occur if watermarks are used for copy control.

The ticket concept for record, playback and generation control has been presented. The basic assumption is that the watermark remains fixed throughout the entire chain for transferring content. In each step an authorisation ticket (or physical mark) has to be present. A cryptographic counter value is decremented every time the ticket is clipped.

An implementational difference between the ticket and signatures is the 'direction' of the cryptographic one-way function. In the ticket concept, random seed data is used as input, the content properties (i.e., the watermark values) are output. In concepts

based on digital signatures (as proposed by others), the content is used as input, while the permission item (i.e., the signature) is an output of the one-way function.

References

1. Cox and J.P.M.G. Linnartz, “Public Watermarks and resistance to tampering”, IEEE Journ. of Sel. Areas in Comm., 1998
2. I.J. Cox and J.P.M.G. Linnartz, “Public Watermarks and resistance to tampering”, ICIP 97, Santa Barbara, CA, October 1997
3. DVD Copy Protection Technical Working Group (CPTWG), DataHiding SubGroup (DHSG), Call for Proposals for Embedded Signaling, Burbank, CA, May 1997.
4. Orange Book, CD R, CD-RW and CD-MO standard, Coordination Office Optical & Magnetic Media, Philips Consumer Electronics, SWA-1, Eindhoven.
5. J. Hogan, "Exploiting Modulation Code Redundancy", invited paper in Proc. Optical Data Storage, Tucson, AR, April 1997, pp. 88-94
6. J.P.M.G. Linnartz and M. van Dijk, "Analysis of the sensitivity attack against electronic watermarks in images", Workshop on Information Hiding, Portland, OR, 15-17 April, 1998.
7. G. Itkis, "Copyright Protection, Analysis and Authorization Chains", Response to DAVIC CFP 7, London, May. 1997.
8. Maes, "Twin Peaks: The histogram attack on fixed depth image watermarks", Preliminary Proceeding of the 2nd International Information Hiding Workshop, Portland OR, April 1998