Digital rights management for digital cinema

Marcus PEINADO, Fabien A. P. PETITCOLAS, Darko KIROVSKI

Abstract—**There is a wide consensus among the feature film production studios that the Internet era brings a new paradigm for film distribution to cinemas worldwide. The benefits of digital cinema to both producers and cinemas are numerous: significantly lower distribution and maintenance costs, immediate access to film libraries, higher presentation quality and strong potential for developing new business models. Despite these advantages, the studios are still reluctant to jump into the digital age. An important concern regarding digital and conventional cinema is the danger of widespread piracy. Piracy already costs Hollywood an estimated two billion dollars annually and digital cinema without proper copyright enforcement could increase this number. In this paper, we present a copyright management system that aims at providing the set of necessary security tools: standard cryptographic primitives and copyright protection mechanisms that enable a reliable and secure feature film delivery system.**

I. INTRODUCTION

Despite the adoption of many innovative techniques to improve the process of film production, such as the use of state-of-the art computer graphics technology to prepare special effects, the distribution of films to cinemas has hardly changed over the past century. Films are still sent to duplication houses and then delivered to cinemas through distribution chains. This has several inconveniences for all parties, from film producers down to the cinema-going public.

Once the final version of a film is ready, it is transferred into a final master print that is duplicated for distribution and projection. Copies are sent via special courier service to cinemas. Such a process is very expensive and release-prints usually cost around $ 2,000 to which an insurance premium is added – leading to a total distribution cost of up to several million dollars [5]. Defining the number of prints to produce is a very difficult issue as one never knows how popular a film will be. Another important problem with anolog films is the quality: in today's cinema, copies have a very good quality. However, the medium deteriorates fairly quickly and has to be replaced to maintain a good show quality. Typical prints suffer degeneration through repeated use, colour drift, cracks in audio etc. This means that the cinema-going public who missed the first show in a cinema will never be able to enjoy the same quality and indeed may never see the film the way the director saw it (of course, lack of standardisation across cinemas worldwide is another factor).

Digital cinema does not need prints and, thus, avoids much of the distribution cost and reduces significantly the risk of making 'too many' or 'too few' copies. However, the cost of upgrading systems may hinder the development of digital cinema and will have to be addressed somehow. Another important advantage of digital cinema is improved image quality. All the above mentioned problems vanish: what the director sees is what the cinema-going public will see and this will never change over time. Also, digital cinema gives much more flexibility to managers of cinemas for scheduling films as they can allocate films and screens on a per show basis. Finally, content protection measures, should help reduce piracy. Several other advantages of digital cinema are detailed by Morley [34].

The feature film production studios are aware of the inevitable change of distribution technology. For example, the Movie Producers Association of America has already created the working group Digital Cinema DC 28 within the Society of Motion Picture and Television Engineers to establish a standard for digital cinema [45]. The challenges involved in creating the future of digital cinema include: reliable and fast content distribution from data centres to cinemas, development of projectors capable of displaying high-fidelity digital imagery and audio and development of security mechanisms that would prevent an explosion of piracy and various forms of fraud that could appear in this new setting. In this paper we will address the security aspects of digital cinema and propose a distribution system that represents a combination of existing digital rights management technology together with fingerprinting techniques.

Digital rights management (D.R.M.) is a set of technologies to enable access controlled data distribution. In most instances, D.R.M. is a system that encrypts digital media content and limits access to only those people who have acquired a proper license to play the content. For digital cinema, D.R.M. technology is the core system that will allow the owners to distribute their films in a controlled way. The owner specifies, in which ways and under which conditions each cinematic asset may be accessed (digital rights, licensing) and the D.R.M. system will try to ensure that each asset can only be accessed as specified by the owner (enforcement).

The *content* or *asset* we consider in this paper is very high value entertainment content of a cinematic title, including video, audio but also text and metadata; this also includes older valuable releases that may be stored in a digital library. From the data management perspective, a typical two hour 35 mm feature film scanned at a standard high-quality resolution of 1920 by 1080 pixels and 24 frames per second (used for high-end H.D.T.V. as well) would, in its uncompressed form, require more than one terabyte of disk space. Assuming a compression rate of 30:1 using MPEG, a typical feature film would incur a transfer of up to several tens of gigabytes of data from the film library. These numbers are likely to increase in the nearest future as no-loss 35 mm scans for editing require at least a 4000 by 4000 pixel resolution (examples of such scanners include Kodak Cineon and Quantel Domino) and cinemas such as IMAX already display content at up to 75 frames per second. Although today these numbers require impressive computing and networking systems, many companies such as Sony, Qualcomm and Microsoft are actively developing their digital cinema projection and distribution technologies [39], [46].

II. Attack model and environment

This section aims to define the goals of the anti-piracy system described in this paper. At an abstract level, D.R.M. or similar anti-piracy technology is quite ubiquitous. For example, a given authentication protocol can be used independently of the type of content which is being distributed (e.g., video, audio, books, etc.) and of the participants in the system (e.g., studio to cinema or Internet retailer to consumer). However the ultimate success of any given content protecting system, its economic feasibility and the appropriateness of certain anti-piracy measures (e.g., fingerprinting) depend strongly on the environment, in which the system has to operate. In general, there will be substantial piracy if the value of the good (to the adversary) exceeds the cost of piracy (including legal threats). Critical environmental factors include the value curve of the content (for the owner and for the adversary), the number of participants in the system, the availability of redistribution channels for copyrighted goods, the difficulty of identifying adversaries, the cost of breaking the content protection measures, the cost of recovering from a compromise and legal penalties for piracy. Technology can influence only some of these factors.

Different kinds of content protection systems have been deployed over the years with varying degrees of success [16] (satellite TV subscription [28], [51], copy protection for video games [19], content distribution over the Internet [37], various content protection measures in consumer electronics devices and digital home networks [2], [3], [11], [48]). However, the operating environments of these systems differ substantially from that of digital cinema.

A first important difference is the value curve of the content. The initial value of each asset is extremely high (up to hundreds of millions of dollars for newly released films) and declines very rapidly (millions of dollars per day). Most exhibition revenues are made during the first couple of weeks after release. Subsequent exhibitions are expected to reach much smaller audiences. For example, the feature film ‘Titanic’ by James Cameron has grossed up $ 600M in the United States only, with $ 400M in box office revenue within only two months of its release [47]. The rapidly declining value curve limits the time span during which protection by the digital cinema system is critical. For example, after a film has been released through other channels (e.g., digital broadcast or D.V.D.), these more weakly protected channels will remove the piracy pressure from digital cinema.

A second environmental parameter, which differentiates digital cinema from the well-studied environments mentioned above, is the relatively small and constrained set of participants (several hundred thousand projectors worldwide versus tens or, possibly, hundreds of millions of satellite T.V. receivers). The projectors contain expensive optical equipment and moderately complex anti-piracy components would not impact the total cost noticeably.

A. System and attack model

The overall system can be modelled as a directed acyclic graph (DAG). The vertices correspond to devices that handle (e.g., store, transform, display) the content. The source nodes of the DAG correspond to authoring facilities that inject content in to the system. The sink nodes are projectors in cinemas. Intermediate nodes have mostly distribution and management functions. The edges of the DAG correspond to communication channels between these devices. We assume the existence of low-bandwidth point-to-point (unicast) channels between nodes (for cryptographic authentication and key exchange) as well as a high-bandwidth multicast channel (for content distribution). The size of the multicast is a system parameter, which encompasses unicast as well as broadcast.

The goal of an attack is to extract content from the system. Without protection measures, any point of the system is subject to attack. That is, content can be extracted at any vertex or at any edge. Typically, edges (communication channels) are protected by means of cryptography (e.g., SSL). Protection measures for vertices (devices) can be categorized as attempts to build trusted computing environments, i.e. environments, which try to assure the integrity and confidentiality of their computation. In this context, our attack model makes the following assumptions:

- Strong cryptography: It is not feasible to circumvent the relevant cryptographic functions (e.g., encryption, signatures) when set up properly.
- Limited strength of trusted environments: Any implementation of a trusted computing environment t can be subverted at a finite cost C_t.

In this model, the devices are the main points of attack. In concrete terms, possible attacks on the system take the following forms:

- The attacker extracts the film (plaintext) from a legitimate projector. For this purpose, the attacker has to overcome the tamper-resistant hardware protecting the projector.

- The attacker extracts cryptographic keys stored in the projector (see below). This allows an arbitrary device to impersonate the compromised projector.

The periphery of the content protection system is another area of attack. These attacks exist irrespective of the content protection system and must be averted by other means:

- The content is attacked before it is injected into the content protection system (e.g., social engineering in the production studio).
- The content is attacked after it leaves the content protection system (e.g., camcorders in the cinema). This attack is difficult to execute without the collusion of cinema personnel. Fingerprinting techniques can help to identify cinemas, where this type of attack occurs frequently.

B. *Content protection objectives*

The general goal is to ensure secure distribution of the content and enforce conditional access to it. In particular, the system should prevent adversaries from obtaining free versions of the original master copy or copies received by cinemas. At the same time, one must recognise that it is not possible to enforce perfect protection. Any given projector (and the films it can access) can be compromised at a fixed cost. Furthermore, the attack will remain undetected, at least, until the compromised content is re-released. We assume that in the case of a commercially relevant re-release of copyrighted films we can detect that a break has occurred.

More precisely, we have the following objectives

- The content protection system should not be subject to global breaks. While our attack model assumes that any given device t can be broken at a cost C_t, the effort in breaking n>1 devices should be about $C_t n$.
- The cost C_t of subverting a given device *t* should be sufficiently high to make attacks of this type infrequent.

A quantitative study of the minimally required value of C_t is beyond the scope of this paper. However, the following measures can be taken to increase C_t.

1. Raise the cost of the initial attack by means of tamper-resistant hardware. As stated above, the high cost of the optical equipment as well as the constrained set of participants makes it possible to deploy more sophisticated security hardware than would be feasible in a retail environment. This might involve mechanical barriers around the projectors.
2. Make adversaries identifiable. Given a compromised copy of a film, it should be possible to identify the compromised projector from which it was extracted. Our system implements this by means of robust fingerprinting. A complementary approach lies in the use of tamper-evident hardware in conjunction with an audit procedure.
3. Enable cheap and easy renewal of the system. After a compromise has been detected, the system must prevent the compromised projectors from receiving new content. More generally, even in the absence of a compromise, the security components of the projectors should be renewed (changed) periodically, in order to present a moving target to potential attackers.

As stated above, the piracy rate depends furthermore on a number of non-technical parameters, such as the legal environment and, most importantly, the policy of the content owner for making films accessible to different cinema operators. Indiscriminate distribution of films will inevitably lead to more frequent compromises than a highly selective policy.

III. System description

This section describes the proposed system. The system consists of a set of *secure repositories* or *nodes*, which implement D.R.M. functionality and which are operated by different participants (studios, distributors, cinemas). D.R.M. enabled projectors are one type of nodes. Section A defines the nodes and describes their critical properties and the functionalities they have to implement. We pay special attention to the nodes inside cinemas in Sect. B. Section C describes how the nodes of the different participants can interact under different configurations to implement sophisticated distribution chains. Finally, we specify the protocols by which nodes communicate in Section IV.

A. *Node capabilities*

A node is a secure repository for protected content. A node can be given the capability to access (e.g., display) given pieces of content. The node will only access the content in accordance with a description of access rights, which originates from the content owner. We call the combination of the cryptographic keys, which allow content access (e.g., decryption) and the description of access rights a *license*. A *D.R.M. system* is given by a collection of nodes and their interactions, which allow content to move between different nodes.

This section will describe the generic capabilities, which are required for D.R.M. enabled nodes. In addition, we will focus particularly on the nodes in cinemas, including the interactions between a central server and individual projectors and speakers.

In general, participating nodes have to implement the following capabilities, in order to meet the functionality and anti-piracy goals stated so far:

- Authentication;

- Rights management (licensing);
- Content encryption and decryption;
- Fingerprinting.
- Provisions for renewability
- Provisions for tamper-resistance

1) Authentication

Depending on its place in a distribution chain, a node can act as a sender or receiver of content. When acting as a sender, the node must ensure that the receiving node, to which it is granting content access, is an authorised (legitimate) node, which will enforce the access rights. Conversely, when acting as a receiver, an authorised node must be able to prove to the sender that it is indeed authorised. This requires authentication capabilities in the nodes. We base authentication on public-key cryptography. Each node is required to store (and hide) a private key, to have an associated public-key certificate and to implement the basic public-key operations (encryption, decryption, signing and verification). Given these primitives standard cryptographic authentication protocols can be used (see [33] for an overview). Section IV will provide more details on the authentication protocol in the proposed system.

2) Rights management and licensing

One of the main purposes of D.R.M. is to allow asset owners to specify how their assets may be accessed after they have been electronically distributed. This requires the definition of a formal language (the digital rights language), in which owners can express these access rules and conditions (license). Before granting access to any given asset, a D.R.M. node will verify if the access is permitted by the license.

A typical license specifies *access rights* or actions, which may be performed with the asset (e.g., ‘play’ or ‘transfer to another node’). Each action is typically accompanied by a set of *conditions*, which restrict the action (e.g., time specifications, counted operations, payment). Enforcement of some of these restrictions may require the availability of secure counters or a secure clock. In addition, the license may specify that certain actions (e.g., fingerprint insertion) have to be performed on the content. Other policy elements of the license may include the following:

- Period of validity of the license;
- Limitation of the show to particular days and hours (e.g., midnight shows);
- Enforcement of audit logs, that is automatically recording of the title, time, duration, etc. of each show for each projector;
- Enforcement that the movie is shown in its entirety. The cinema cannot skip certain scenes and sections, such as cast and acknowledgments;
- Use of specific types of projectors and rooms (e.g., guaranteeing quality, size of display, audience capacity)
- Requirements for certain projector features (e.g., audit, fingerprinting).

Furthermore, a license typically specifies a set of principals or entities to which the license is tied. Typically, a principle will specify a node or a class of nodes.

A license should be protected by cryptographic means, in order to ensure its integrity as it is sent over unprotected channels from the owner across intermediate nodes to the projector. Standard digital signatures are sufficient for this purpose. The content decryption keys can be tied to the associated license rules by the same mechanism. The XrML rights language [7] has all the properties described in this section.

3) Content encryption and decryption

As stated above, nodes are secure repositories, which can access assets (i.e., films). In the untrusted space between nodes, the assets are protected by means of encrypting them with keys, which are only accessible by the intended destination node. Thus, nodes must be able to decrypt and encrypt assets. Given the ability of nodes to hide private keys (cf. Section *1)*), this is a matter of implementing standard cryptographic ciphers (cf. [33] for an overview). Performance requirements and the format, in which the film is stored, may impose additional constraints on the cipher.

4) Fingerprinting

As described in Section II, we have to account for the possibility that individual nodes may be compromised. Fingerprinting of assets is our main tool for the identification of compromised nodes. Thus, at least some nodes have to be able to insert fingerprints into assets. Section V will give a detailed description of fingerprinting in the proposed system.

5) Renewability and key management

The purpose of renewability is to allow the overall system to recover from different types of attacks, which are part of our attack model. The key management strategy should be such that possible breaks are confined to individual nodes and recovery is fast and cheap. Given an arbitrary set of compromised nodes, it must be possible to prevent them from receiving new content without affecting the remaining nodes. This mandates that each node should have *unique* cryptographic keys. Thus, if a single node is compromised (e.g., by a malicious operator), its key can be revoked without affecting any other node. Revocation is implemented during authentication, where the validity test of the receiver's public key certificate includes a test whether this public key is contained in a revocation list of compromised keys. The revocation information should be cryptographically tied to the content.

The existing watermarking and fingerprinting algorithms are typically considered to be far less robust than cryptographic algorithms. Thus, the fingerprinting component of the nodes should be easily field-upgradeable. For example, signed watermarking code could be distributed with each asset.

6) Platform security and tamper resistance

A node will only be able to participate in the system if its public key is certified by a trusted authority. The trusted authority will only certify a node, if its initial hardware-software configuration meets certain requirements. These requirements should include that the nodes are *closed* platforms, where the term closed means that – without significant hardware tampering – it should not be possible to install arbitrary (untrusted) software or hardware on the node. Thus, in the absence of significant hardware manipulations, an authorised node is initially trusted (precondition for authorisation) and will remain trusted (closed platform). Note that closed platforms do not preclude field upgrades. For example, nodes may allow installation of software, but only if the software is certified (e.g., digitally signed) by a trusted authority.

An attack on a node will involve attempts to discover the node's private key (such that the node can be impersonated by an arbitrary untrusted device) or attempts to gain access to content, which is not permitted by a license (e.g., extracting decrypted video). Based on the requirements of initial integrity of the node and of being a closed platform, these attacks will involve hardware tampering. Defence against attacks of this type will involve tamper-resistant or tamper-evident hardware, in combination with an audit procedure. Given the very high price of digital projectors, adding an extra tamper resistant processing unit to store the projector key, do the decryption and enforce the license given will have a minor impact on the final cost of the projector. Details lie beyond the scope of this paper.

B. Cinema setting

In the proposed system, each cinema has a central server and a number of projectors. The Cinema Server is the central switch for the cinema from which a single operator can start shows, select screens on which films are displayed, etc. Reception of the encrypted content by the cinema is done on this server. It stores the encrypted content together with the activation data that is the content encryption key, itself encrypted under the projector public key and the D.R.M. rules.

Encrypted content and D.R.M. rules are uploaded onto the projector before the show. The projector verifies that the license is genuine, checks that the request of the Cinema Server is compatible with the license, decrypts the content encryption key using its private key and starts decrypting and displaying the content. There should be an option for the projector to inform the server whether it will be able to perform such a task at a given time in the future, such that the cinema operator can ensure that shows will run smoothly. Optionally, verification of the integrity of the content can be done using 'linked' hash values of each block of the content [18]. This works both if the content is 'streamed' or available in full.

C. Distribution chain

From the point of view of most asset owners, digital cinema opens a new distribution channel. Adoption of digital cinema will be facilitated if this new distribution channel can mimic the properties of existing distribution channels. Typically, these distribution channels transfer physical copies of the film. A typical chain could have studios, several tiers of distributors and cinemas as participants. The film moves from the producer or studio through the links in the chain to the cinema. Usage data and, possibly, payment moves in the opposite direction. The system architecture should be able to model real world distribution chains. Indeed, its flexibility has the potential of going far beyond the conventional physical distribution chains.

It is possible for D.R.M. to model the entire distribution chain of Figure 2a). The participants in the chain have the functionality of nodes (access to content and decryption keys, rights interpretation, generation of a modified license for the next participant in the chain). The flow of payment information in the opposite direction can be handled within the D.R.M. system or outside of it (separate commerce functionality). Figure 2b) shows a distribution chain, in which a central *clearing house* issues licenses and potentially collects and redistributes usage data or payments.

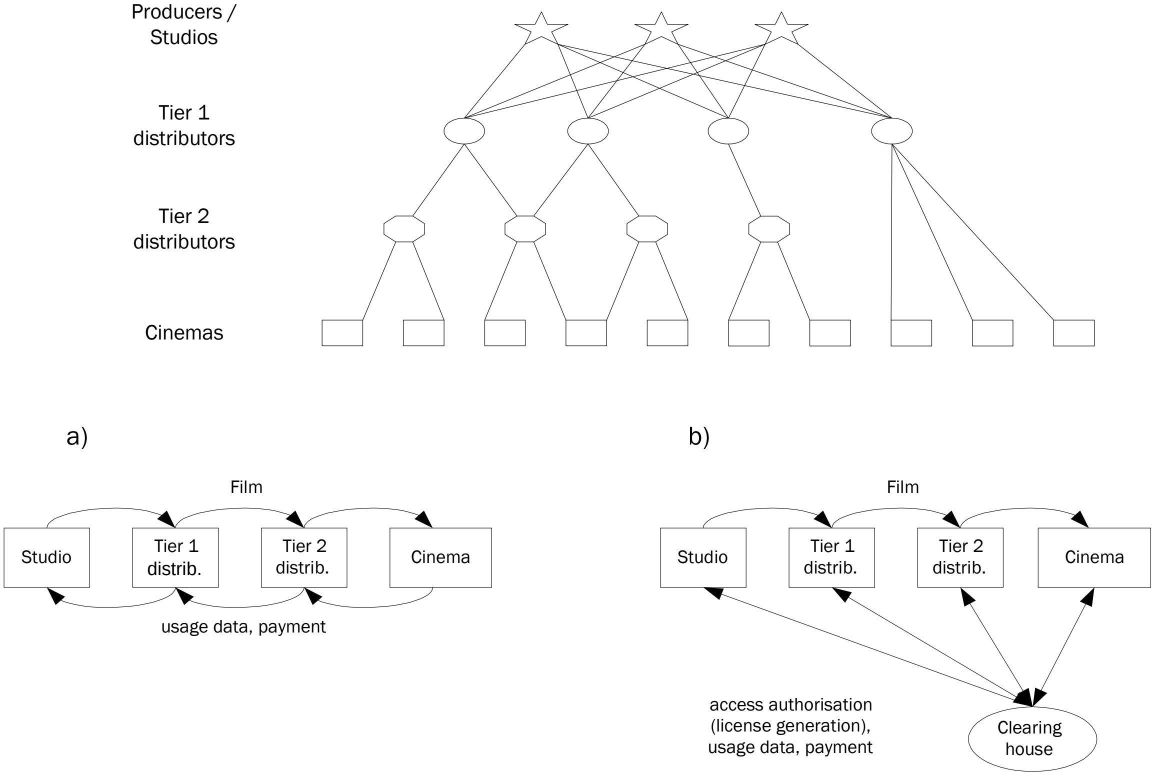

Figure 1 – Two possible film distribution chains.

It should be noted that, under D.R.M., the distribution chain is more than a sequence of independent links. The license rules do not only allow the source to specify how the content may be accessed. In addition, each node can specify within which limits the nodes downstream in the chain can grant rights to their respective downstream nodes. For example, a studio may make a film available to a distributor. The associated license specifies not only how many copies of the film each distributor is allowed to make available to cinemas and which price the distributor has to pay to the studio, but also which rights the distributor can issue to each cinema server. The license can also specify the terms of the license that can be issued to the cinema. For example, it can specify how many projectors in any given cinema can simultaneously play the film. The XrML rights language [7] is an example of language which allows the specification of such *derivative rights*.

In summary, the following functionalities are required to support the distribution chains described above:

- The rights language must be able to express derivative rights, and nodes in the distribution chain must be able to interpret them.
- Clearinghouses or other nodes (depending on the configuration) must be able to generate new licenses based on the derivative rights from a source license and additional specification from the node operator.

It must be possible to track the path of an asset or its associated license in a tamper proof way. The implementation of the system suggested in the next session does not depend on the distribution chain chosen.

IV. Concrete system description

The previous sections have outlined a generic digital cinema system. The goal of this section is to define a concrete system by specifying parameters, algorithms and protocols, which have been left open so far.

Nodes are built largely on standard PC hardware. However, the security functionality of each node is encapsulated in a central security component, which is implemented by a secure coprocessor (S.C.P.) [53], [40]. The S.C.P. implements the following functions

- The standard set of public key operations (encryption, decryption, signing, signature verification, key generation) for the R.S.A. algorithm on 2048 bit keys. Each S.C.P. s contains a unique decryption key pair (P_s, P_s^{-1}) and a unique signature key pair (Q_s, Q_s^{-1}). The public keys P_s and Q_s are certified by the manufacturers of the S.C.P. The standard options regarding public key infrastructure and certification authorities exist.
- A hash function (SHA-1 [14]) and a symmetric cipher (A.E.S. [15]).

- An engine for the evaluation and manipulation of XrML [7] documents (licenses).
- A signal processing module for the insertion of fingerprints.
- Control logic, which invokes the functions mentioned so far, as described in this paper.

The security component implementing these functions should be protected by hardware against tampering. In particular, the S.C.P. must not be openly programmable in a way, in which such programming could subvert the S.C.P. A detailed description of mechanisms for hardware tamper resistance is beyond the scope of this paper. Various accounts can be found in the literature [27], [40].

The overall system requires a ubiquitous low-bandwidth two-way communication channel for authentication and key transport and a high-bandwidth channel for the transmission of the movies themselves. The latter can, in principle, be restricted in different ways (e.g., one-way, high latency, etc.). We rely on the Internet as the low-bandwidth communication channel between geographically distributed nodes (i.e. studios, distributors, cinemas). This requires each cinema to be able to access the Internet (at least through a low-speed telephone connection). There are no special requirements on the high-bandwidth channel. Any channel, which provides the required bandwidth, can be used by the system (e.g., satellite, high-speed Internet, dedicated networks, or physical distribution). For the communication between the projectors and the central server within each cinema, our system assumes the existence of a high-bandwidth intranet in the cinemas.

We propose a layout of the distribution chain as described in Figure 2a). Content flows from a studio through one or more distributors and a central cinema server to a projector without the intervention of a clearing house. The exact protocols and information flow are described below. It should be noted, though, that this decision constitutes policy of the system operator. Technically, the nodes and protocols could easily be adapted to work with a clearing house.

A. *Content preparation and distribution*

As stated in Section 3.1, one of our critical goals in increasing the overall robustness of the system is to make adversaries identifiable. For this purpose, our system supports individualisation of distributed content by means of fingerprinting. In addition, different copies of a given asset may be encrypted with different keys. The system permits this individualization to occur at the granularity of each individual copy of an asset. At the same time, the system permits the operator of each server to choose the granularity of individualisation and thus the trade-off between robustness (uniqueness) and cost.

Broadcasting the same encrypted content to all projectors (so no fingerprints are used) is the cheapest solution but it does not allow the content owner to trace infringers. For this reason we prefer unicast distribution or, at least, multicast: content is encrypted and fingerprinted by the content owner on a per-projector basis to a group of *n* projectors depending on the level of traceability required by the distributor. With this choice of granularity, the distributor can use D.V.D., tapes or satellite broadcast as the content distribution channel, when network bandwidth to the cinema is not sufficient. Later on projectors add their own fingerprint during the projection of the film.

Moreover, the fact that the set of identical copies is limited and distributed over geographically different regions helps to narrow down the search significantly when tracking is done. In particular, it is possible to uniquely identify an infringing cinema after it redistributes a small number of movies. More precisely, let *k* be the number of cinemas or projectors. In order to send identical copies of a given movie to approximately *n* recipients, the distributor produces $\lceil k/n \rceil$ differently fingerprinted and encrypted versions of any given movie. A simple randomized model is to choose a uniformly distributed random copy from the set of $\lceil k/n \rceil$ versions for each recipient. This model is equivalent to the birthday attack model, whose analysis carries over. For example, the expected number of pairs of recipients, who are sent identical copies of *m* different movies, is

$$\binom{k}{2}\left(\frac{n}{k}\right)^m$$

With parameter choices of $k = 100{,}000$, $n = 10{,}000$ and $m = 10$, this number is approximately ½ and 10 different versions of each asset have to be produced. More sophisticated models have been studied in the context of traitor tracing [12],[13] and achieve certain degrees of collusion resistance (i.e. the ability to uniquely identify infringing recipients, even if groups of adversaries combine their copies).

Clearly encryption keys and fingerprinting patterns should not be shared across different cinematic titles, so that breaks are limited to one title only.

B. *Protocols*

Beyond the transmission of encrypted content, the communication between the different participants – especially between the distributors and the cinema servers – must achieve the following: (a) authentication, (b) key transport (of content decryption keys), (c) transmission of content access rules and (d) transmission of usage information.

Our authentication protocol is the Needham-Schroeder public-key protocol [33],[36]. This protocol provides mutual entity authentication and key transport [33]. In order to facilitate the detection of compromised keys, which are being shared by several adversaries to request decryption keys from distributors, we augment the basic authentication protocol as follows. For each sender-receiver relationship, the sender and the receiver maintain a monotonic counter, which is increased at the beginning of each sender-receiver transaction. The sender transmits its value to the receiver, who will service the transaction only if the received value is the expected value. This method will detect compromised keys, which are being shared without coordination.

Figure 2 displays the overall protocol where the sender is a distributor D and the receiver a cinema server S and a projector P. The sender D generates a license for the receiver S as follows. The content decryption key K_C is encrypted by the sender D with the public key of the receiver S: $X_{CS} = \mathrm{E}(K_C, P_S)$. X_{CS} is cryptographically tied to the content access rules R_C by means of two mechanisms. Firstly, the sender signs the concatenation of X_{CS} and R_C using his signing key Q_D^{-1}: $S_{CDS} = \mathrm{S}(R_C|X_{CS}, Q_D^{-1})$. Secondly, he applies a message authentication code (MAC) with key K_C to the concatenation of X_C and R_C: $M_{CS} = \mathrm{M}(R_C|X_{CS}, K_C)$. The former mechanism ensures that only authorised servers can issue licenses while the latter mechanism ensures that the issuer of the license had knowledge of K_C and makes the binding of X_C and R_C robust even in the presence of compromised signing keys. That is, in order to forge a license, an adversary must have access to the content decryption key K_C and to a compromised signing key. We call the structure consisting of X_C, R_C, the signature S_{CDS} and the MAC M_{CS} a license L_C.

Upon receiving a license, the S.C.P. of S decrypts the content key $\mathrm{K}_\mathrm{D} = \mathrm{E}(X_{CS},P_S^{-1})$ and verifies the MAC (Is $M_{CS} = \mathrm{M}(R_C|X_{CS}, K_C)$?) and the signature (Is $S_{CDS} = \mathrm{S}(R_C|X_{CS}, Q_D)$?). In general, the latter test requires S to establish the validity of D's public key Q_D. Typically, S will trace a certificate chain to a root key that is known to the S.C.P. of S and consult a revocation list. The certificate chain and the revocation list must be protected by means of digital signatures and contain a freshness indicator (e.g., expiration data, version number). The same protocol is used to transfer K_D from the cinema server S to a projector P and to transfer usage data from the cinemas back to the distributors and the studios.

Figure 2 - Distribution protocol for digital content.

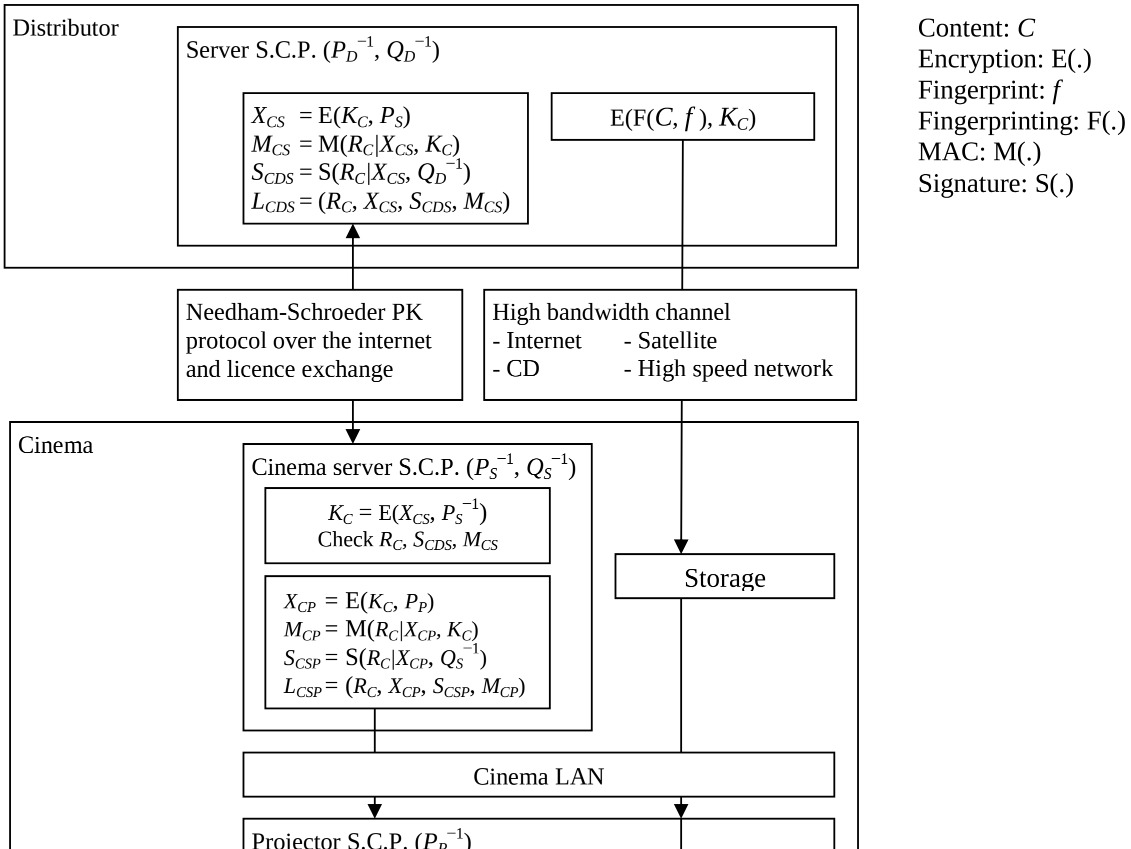
Distributor
Server S.C.P. (P_D^{-1}, Q_D^{-1})
$X_{CS} = \mathrm{E}(K_C, P_S)$
$M_{CS} = \mathrm{M}(R_C|X_{CS}, K_C)$
$S_{CDS} = \mathrm{S}(R_C|X_{CS}, Q_D^{-1})$
$L_{CDS} = (R_C, X_{CS}, S_{CDS}, M_{CS})$
$\mathrm{E}(\mathrm{F}(C, f), K_C)$
Needham-Schroeder PK protocol over the internet and licence exchange
High bandwidth channel
- Internet
- Satellite
- CD
- High speed network
Cinema
Cinema server S.C.P. (P_S^{-1}, Q_S^{-1})
$K_C = \mathrm{E}(X_{CS}, P_S^{-1})$
Check R_C, S_{CDS}, M_{CS}
$X_{CP} = \mathrm{E}(K_C, P_P)$
$M_{CP} = \mathrm{M}(R_C|X_{CP}, K_C)$
$S_{CSP} = \mathrm{S}(R_C|X_{CP}, Q_S^{-1})$
$L_{CSP} = (R_C, X_{CP}, S_{CSP}, M_{CP})$
Storage
Cinema LAN
Projector S.C.P. (P_P^{-1})
Content: C
Encryption: E(.)
Fingerprint: f
Fingerprinting: F(.)
MAC: M(.)
Signature: S(.)

V. FINGERPRINTING DIGITAL AUDIO/VIDEO CONTENT

As we have mentioned above, we cannot solely rely on traditional data protection techniques such as encryption or scrambling, because multimedia will eventually be played in an unscrambled or decrypted format. Therefore, in all scenarios it is possible to record the decrypted content; in the worst case by recording the analogue output of the playback device. An approach that can survive such re-recording attacks is insertion of watermarks/fingerprints in the content itself [44].

Fingerprints are used both by the content owner and by the projectors to enforce content copyright by enabling the copyright owner to trace back the source of a piracy act [23]. In a typical content protection scenario, all users are given different copies of the content, where each copy contains a fingerprint – a user-specific watermark. If an unauthorised client redistributes the fingerprinted content, its uniqueness is used to trace back the malicious client. Fingerprints, just as watermarks, must be:

- **reliable:** the probability of falsely accusing a user of piracy should be as small as possible (at least 10^{-12}) while preserving a solid likelihood of detecting malevolent users even after strong malicious attacks on the fingerprinted content (at least 10^{-3}).
- **robust** to common editing (e.g., compression, format conversion, filtering) and various malicious attacks (e.g., de-synchronisation, crop-and-paste [38]); at the moment most existing algorithms will not survive distortions introduced by the projection onto screen followed by re-recording using a law quality camcorder but research is being carried to invert these distortions [54].
- **easy-to-detect:** as oppose to watermark detection which is done at the client before playing the content, fingerprints are detected at the server after the piracy has been committed, thus they do not have to be detected in real-time; in addition, the fingerprint detector is allowed to compare the 'attacked' fingerprinted content to the original;
- **imperceptible** to the target audience and any analytical tools; it is important to stress that full imperceptibility against statistical tools may be difficult to achieve especially in cases when the probability density function of the original signal is well known [41].

Since it has been demonstrated that a clique of clients can be effective in removing the secret marks by colluding their copies [10], it is important that fingerprint encoding enables as good as possible:

- **collusion resistance:** defined as the number of copies that can be colluded with an arbitrary best-known collusion algorithm to result in a new copy that still reveals at least one of the colluders,
- **traceability:** the resulting copy of the content should identify at least one of the colluders with non-negligible probability regardless of any malicious attacks superimposed to the collusion process (e.g., de-synchronisation, jamming); and
- **frameproof:** no coalition of users should be able to create a copy that frames an innocent user.

An important asymptotical upper bound on collusion resistance of fingerprinted material has been established by Ergun et al. [10]:

$$K \square O\left(\sqrt{\frac{n}{\ln(n)}}\right)$$

where K is the collusion resistance and n corresponds to object length. Obviously, this upper bound puts a strong limit on the efficacy of any fingerprinting mechanism. In the remainder of this section, we briefly review watermarking techniques as core technologies for the content marking layer and fingerprint encoding methods as techniques for maximising collusion resistance.

A. Watermarking

Watermarking schemes rely on the imperfections of the human perception system (H.P.S.). Numerous secret hiding techniques explore the fact that the H.P.S. (for both auditory and visual) is insensitive to small amplitude changes, either in the time [1] or frequency [9], [43], [44] domains, as well as insertion of low-amplitude time-domain echoes [20]. Information modulation is usually carried out using: spread-spectrum (S.S.) or quantisation index modulation (Q.I.M.). Advantages of S.S. and Q.I.M. watermarking include: (*i*) testing for watermarks does not require the original and (*ii*) it is difficult to extract the embedded information using optimal statistical analysis under certain conditions [41]. In addition, S.S. watermark detection is exceptionally resilient to attacks that can be modelled as time- and frequency-axis scaling with fluctuations [24], [25] and additive or multiplicative noise.

Disadvantages include: (*i*) the watermarked signal and the watermark have to be perfectly synchronised at watermark detection and (*ii*) to achieve a sufficiently small error probability, signal length may need to be quite large, increasing detection complexity and delay. By far the most significant deficiency of both schemes is that they are not BORE-resistant (BORE – break once run everywhere), i.e. by breaking a single player (debugging, reverse engineering, or the sensitivity attack [30]), one can extract the secret information (the key used to generate the SS sequence or the hidden quantisers in Q.I.M.) and recreate the original (S.S.) or create a new copy that will induce the Q.I.M. detector to treat the content as unmarked. To address this problem, recently several asymmetric watermarking schemes have been developed with little success [17], [22]. Fortunately, the BORE-resistance issue does not play a significant role when watermarks are used as fingerprints, as the detection process is fully performed at the server side.

B. *Fingerprint Encoding*

Fingerprint encoding for maximised collusion resistance is a notoriously hard problem. Fingerprint systems are commonly based on a *marking assumption*: existence of a robust watermark technology that disables the adversary to remove any bit of the fingerprint while preserving the perceptual characteristics of the original recording. The difficulty behind providing high collusion resistance stems from the fact that a coalition of users can compare their copies and at each position of the content (e.g., time domain sample) where at least one copy differs from the others, the coalition can conclude the value of the embedded mark or even the value of the original non-marked sample. With the extracted information, a sufficiently large coalition can remove the fingerprint and/or frame another innocent user.

Boneh and Shaw have introduced an encoding mechanism that achieves given collusion resistance c with object length l proportional to:

$$l \in O\left(c^4 \log\left(N/\varepsilon\right)\log\left(1/\varepsilon\right)\right)$$

where ε is the probability of fingerprint misdetection and *N* is the total number of users [4]. In other words, such codes result in asymptotical collusion resistance proportional to the fourth root in object size: $c \in O\left(\sqrt[4]{l}\right)$.

Several other fingerprint encoding techniques have been developed which improve the basic Boneh Shaw mechanism by superimposing it to spread-spectrum [52] and with different attack modelling strategies [21].

VI. Evaluation

We briefly evaluate the system we have proposed. A number of design choices were made, and it is natural to review possible alternatives. The choices that could lead to alternative system designs can be categorized as

- choices regarding the network topology (What types of nodes are required in the system. What special capabilities must each type of node have?)
- choices regarding algorithms and protocols.

Network Topology: We have chosen to include a central server for each cinema into the design to allow central control over all projectors. Alternatively, the projectors could be operated autonomously, eliminating the need for a central server and a local area network in the cinema. Our design has the benefit of offering a central control point to the cinema operator, allowing him to coordinate the shows in the different show rooms more easily. The central server design also allows the projectors to be simplified. Certain functionality, such as storage, only has to be implemented in the central server and does not have to be replicated in every projector. On the other hand, the projectors implement most of the security related functionality described in Sect. III. This makes it possible to rely on the same mechanisms that are used throughout the system to protect the movies up to the optical equipment.

As stated in section III, the network topology outside the cinemas (distribution chain) is mostly a matter of system configuration, rather than system design. Our architecture can support a wide range of system configurations beyond the examples in Sect. III. For instance some cinemas have already been equipped with digital projectors and receive the films on hard-drives drives send to them without further protection than the physical protection in place for distributing films on celluloid. The system described in this paper could fully accommodate with this and would require the film to be encrypted and fingerprinted.

Regarding algorithms and protocols, section III defines the functionality the system relies on without specifying particular algorithms. Design decisions at that level include the use of public key cryptography and of fingerprinting. Public key cryptography in connection with an appropriate public key infrastructure facilitates key management. The alternative – symmetric cryptography – typically requires a central trusted server (e.g., Kerberos). We prefer a system design that does not have this infrastructure requirement. Section IV instantiates the requirements of Sect. III by naming several concrete algorithms and protocols. We note that these are not the only choices. Many good alternatives exist.

VII. Conclusions

In this paper we have described a possible distribution technique for digital cinematic titles. We have detailed the basic components of the system including the general distribution protocol as well as the use of fingerprinting to trace infringements. Our aim was to provide an overview of the general principles of content protection for digital cinema, rather than giving a detailed description of any particular design. We have assumed a fairly powerful and high-level model for secure hardware. While we have argued that the model is realistic in the context of digital cinema, we have omitted a detailed description of how it could be implemented. Such a description as well as a more detailed description of the certification procedure and more exact parameters of the operating environment will be the focus of another analysis.

VIII. References

[1] P. Bassia and I. Pitas, 'Robust audio watermarking in the time domain,' Proc. EUSIPCO 98, vol. 1, pp. 25–28, Rodos, Greece, Sept. 1998.

[2] A. Bell, 'The dynamic digital disk,' IEEE Spectrum, vol. 36, no. 10, pp. 28–35, October 1999.

[3] J. A. Bloom, I. J. Cox, T Kalker, J.-P. M. G. Linnartz, M. L. Miller, C. B. S. Traw, 'Copy protection for DVD video,' Proceedings of the I.E.E.E., vol. 87, no. 7, pp. 1267–1276, July 1999.

[4] D. Boneh and J. Shaw, 'Collusion-Secure Fingerprinting for Digital Data,' In Advances in Cryptology – CRYPTO 95, pp. 452–465, 1995.

[5] J. L. Butt, Jr and C. D. Howe, 'eCinema – Projecting the future of movies', The Forrester Brief, 19 July 1999.

[6] Call for Proposals, Phase I, http://www.sdmi.org, 1999.

[7] ContentGuard, 'XrML: Extensible Rights Markup Language,' http://www.xrml.org

[8] B. Chen and G. W. Wornell, 'Digital watermarking and Information embedding using dither modulation,' Proc. IEEE Workshop on Multimedia Signal Processing, Redondo Beach, CA, pp. 273–278, Dec. 1998.

[9] I. J. Cox, J. Kilian, T. Leighton and T. Shamoon, 'A secure, robust watermark for multimedia,' in R. Anderson, ed. *Information Hiding*, proceedings of the 1st international workshop, Lecture Notes in Computer Science, pp.185–206, Cambridge, England, 1996.

[10] F. Ergun, J. Kilian and R. Kumar, 'A note on the limits of collusion-resistant watermarks,' Eurocrypt, pp. 140–149, 1999.

[11] A. M. Eskicioglu and E. J. Delp, 'Overview of multimedia content protection in consumer electronics devices,' IEEE Signal Processing: Image Communication, vol. 16, no. 5, pp. 681–699, April 2001.

[12] A. Fiat, 'Tracing Traitors,' In Advances in Cryptology – CRYPTO'94, Lecture Notes in Computer Science, Springer-Verlag, pp. 257—270, 1994.

[13] A. Fiat and T. Tassa, 'Dynamic Traitor Tracing,' In Advances in Cryptology – CRYPTO'99, Lecture Notes in Computer Science, Springer-Verlag, pp. 354—371, 1999.

[14] FIPS 180-1, 'Secure hash standard,' Federal Information Processing Standards Publication 180-1, U.S. Department of Commerce/N.I.S.T., National Technical Information Service, Springfield, Virginia, 17 April 1995.

[15] FIPS 197, 'Advanced Encryption Standard (A.E.S.),' Federal Information Processing Standards Publication 197, U.S. Department of Commerce/N.I.S.T., National Technical Information Service, 26 November 2001.

[16] B. Fox, 'The pirate's tale,' New Scientist, December 1999. http://www.newscientist.com/ns/19991218/thepirates.html

[17] T. Furon and F. P. Duhamel, 'Robustness of an Asymmetric Watermarking Method,' Proc. IEEE International Conference on Image Processing, Vancouver, Canada, vol. III, pp. 21–24, 2000.

[18] R. Gennaro and P. Rohatgi, 'How to sign digital stream,' in B. S. Kaliski Jr, ed., *Advances in Cryptology – Crypto'97*, proceedings of the 17th annual international cryptology conference, pp.180–197, vol. 1294 of Lecture notes in computer science,' 17–21 August 1997..

[19] D. Grover, 'The protection of computer software: its technology and applications,' Cambridge University Press, Cambridge, England, 1992.

[20] D. Gruhl, A. Lu and W. Bender, 'Echo hiding,' Information Hiding, Springer Lecture Notes in Computer Science, vol. 1174, pp. 295–315, 1996.

[21] H.-J. Guth and B. Pfitzmann, 'Error- and collusion-secure fingerprinting for digital data,' in A. Pfitzmann, ed., *Information Hiding*, proceedings of the 3rd international workshop, vol. 1768 of Lecture Notes in Computer Science, pp. 134–145, Dresden, Germany, September/October 1999.

[22] F. Hartung and B. Girod, 'Fast public-key watermarking of compressed video,' Proc. of the IEEE Int. Conf. on Image Processing, Santa Barbara, CA, October 1997.

[23] S. Katzenbeisser and F. A. P. Petitcolas, Eds. Information Hiding Techniques for Steganography and Digital Watermarking. Boston, MA: Artech House, 2000.

[24] D. Kirovski and H. Malvar, 'Robust Spread-Spectrum Audio Watermarking,' Proc. of the IEEE Int. Conf. on Acoustics, Speech and Signal Processing, May 2001.

[25] D. Kirovski and H. Malvar, 'Robust Covert Communication over a Public Audio Channel Using Spread Spectrum,' in I. Moskowitz, ed., *Information Hiding*, proceedings of the 4th international workshop, vol. 2137 of Lecture Notes in Computer Science, pp. 354–368, Pittsburgh, PA, U.S.A., April 2002.

[26] P. Kocher, 'Television and Tamper Resistance,' Cryptographic Research, 2000.

[27] P. Kocher, J. Jaffe and B. Jun. Differential Power Analysis. In Advances in Cryptology – CRYPTO'99, Lecture Notes in Computer Science, Springer-Verlag, 1999.

[28] M. G. Kuhn, 'Attacks on Pay-TV Access Control Systems,' Security Seminar, Computer Laboratory, Cambridge 9 December 1997. http://www.cl.cam.ac.uk/~mgk25/vc-slides.pdf

[29] J. Lacy, J. Snyder, D. Maher, 'Music on the Internet and the intellectual property protection problem,' Proc. ISIE, July 1997.

[30] J. P. Linnartz and M. van Dijk, 'Analysis of the sensitivity attack against electronic watermarks in images,' Proc. of The Information Hiding Workshop, Portland, Oregon, April 1998.

[31] P. D. Lubell, 'A coming attraction: D-cinema.' IEEE Spectrum, pp. 72–78, March 2000.

[32] F.J. MacWilliams and N.J.A. Sloane, 'The Theory of Error-Correcting Codes,' North-Holland, Amsterdam, 1977.

[33] A. Menezes, P. van Oorshot and S. Vanstone, 'Handbook of applied cryptography,' CRC Press, 1997.

[34] S. A. Morley, 'Making digital cinema actually happen – What it takes and who's going to do it', presented at SMPTE 140th technical conference, Pasadena, California, October 1998.

[35] P. Moulin and J. A. O'Sullivan, 'Information-theoretic analysis of information hiding,' submitted to IEEE Transactions on Information Theory.

[36] R. Needham and M. Schroeder, 'Using Encryption for Authentication in Large Networks of Computers,' Communications of the ACM, vol. 21, pp. 993—999, 1978.

[37] M. Peinado, 'Digital Rights Management in a Multimedia Environment,' SMPTE Journal, 2002.

[38] F. A. P. Petitcolas, R. J. Anderson and M. G. Kuhn, 'Attacks on copyright marking systems,' Information Hiding Workshop, Portland, Oregon, USA, 1998.

[39] SightSound Technologies and Microsoft Debut Digital Cinema Trial Using Windows Media, 19 March 2001. http://www.microsoft.com/presspass/press/2001/Mar01/03-19SightsoundPR.asp

[40] S. Smith and S. Weingart. Building a high-performance, programmable secure coprocessor. Computer Networks. 31(8):811—860, April 1999.

[41] J. K. Su and B. Girod, 'Power-spectrum condition for energy-efficient watermarking,' Proc. IEEE Int. Conf. Image Processing, October 1999.

[42] M. Sudan, 'Decoding of Reed Solomon Codes Beyond The Error-Correction Bound,' Journal of Complexity, vol.13, (no.1), pp. 180-193, March 1997.

[43] M.D. Swanson, B. Zhu, A.H. Tewfik and L. Boney, 'Robust audio watermarking using perceptual masking,' Signal Processing, vol.66, pp. 337–355, 1998.

[44] W. Szepanski, 'A signal theoretic method for creating forgery-proof documents for automatic verification,' Proc. Carnahan Conf. on Crime Countermeasures, Lexington, KY, pp. 101–109, May 1979.

[45] The Motion Picture Association of America. http://www.mpaa.org/dcinema.

[46] The Qualcomm Digital Cinema. http://www.qualcomm.com/digitalcinema

[47] The Internet Movie Database. http://www.imdb.com

[48] C. B. S. Traw, 'Protecting digital content within home,' IEEE Computer, pp. 42–47, October 2001.

[49] L. Vaitzblit, 'A high-resolution video server for cinema of the future,' IEEE Multimedia, pp. 65–69, Fall 1995.

[50] H. L. Van Trees, Detection, Estimation and Modulation Theory, Part I. New York: Wiley, 1968.

[51] A. Wool, 'Key management for encrypted broadcast,' ACM Transactions on Information and System Security, vol. 3, no. 2, May 2000, pp. 107–134.

[52] Y. Yacobi, 'Passive Fingerprinting,' DIMACS Workshop on Management of Digital Intellectual Property, 17–18 April 2000, DIMACS Center, Rutgers University, Piscataway, NJ, U.S.A.

[53] B. Yee. Using Secure Coprocessors. Ph.D. Thesis, Carnegie Mellon University, 1994.

[54] D. Delannay et al., 'Integrated fingerprinting in secure digital cinema projection,' in A. G. Tescher Ed. *Applications of Digital Image Processing XXIV*, col. 4472, SPIE, San Diego, California, U.S.A. , 31 July–3 August 2001.