STEPTOE & JOHNSON LLP
ATTORNEYS AT LAW

1330 Connecticut Avenue, NW
Washington, DC 20036-1795

Telephone 202.429.3000
Facsimile 202.429.3902
www.steptoe.com

Rhonda M. Bolton
202.429.6495
rbolton@steptoe.com

ORIGINAL

EX PARTE OR LATE FILED

RECEIVED
NOV 14 2001
FEDERAL COMMUNICATIONS COMMISSION
OFFICE OF THE SECRETARY

November 14, 2001

Via HAND DELIVERY
Ms. Magalie Roman Salas
Secretary
Federal Communications Commission
445 Twelfth Street, SW – Room TW-A325
Washington, D.C. 20554

Re: ***Ex Parte* Presentation** -- *In the Matter of Review of the Commission's Rules and Policies Affecting the Conversion to Digital Television*, MM Docket No. 00-39; *In the Matter of Commercial Availability of Navigation Devices*, CS Docket No. 97-80; *In the Matter of Compatibility Between Cable Systems and Consumer Electronics Equipment*, PP Docket No. 00-67

Dear Ms. Salas:

In accordance with Section 1.1206 of the Commission's Rules, 47 C.F.R. § 1.1206, Veridian Corporation ("Veridian"), through its undersigned counsel, writes to report two *ex parte* presentations to the Commission in the above-referenced dockets. On October 31, 2001, representatives of Veridian met separately with: (1) Commissioner Kathleen Q. Abernathy and Stacy Robinson, Mass Media Legal Advisor to Commissioner Abernathy; and (2) Susanna Zwerling, Media and Consumer Protection Legal Advisor to Commissioner Michael J. Copps.

The purpose of the meetings was to inform the Commission of digital copy protection technology developed by Veridian and the role that such technology may play in aiding a swift transition to digital television. Also discussed was the need for the Commission to ensure the compatibility of navigation devices both with multichannel distribution systems using Veridian's technology and with systems using the methodology set forth in the current Interim Specification proposed by Cable Labs' OpenCable project.

The matters addressed in the meetings are summarized in detail the attached letter, which was provided during the meeting to the participants.

An original and two copies of this *Ex Parte* Notice for each docket number are provided to your office as required by Section 1.1206 of the Commission's Rules.

Respectfully submitted,

Rhonda M. Bolton

Rhonda M. Bolton
Counsel for Veridian Corporation

Enclosure

cc: Commissioner Kathleen Q. Abernathy
Stacy Robinson
Susanna Zwerling

VERIDIAN
Corporate Headquarters

1200 South Hayes Street
Suite 1100
Arlington, Virginia 22202
703.575.3100
www.veridian.com

October 25, 2001

The Honorable Michael K. Powell
Chairman
Federal Communications Commission
445 Twelfth Street, SW
Washington, D.C. 20554

Re: *In the Matter of Review of the Commission's Rules and Policies Affecting the Conversion to Digital Television*, MM Docket No. 00-39; *In the Matter of Commercial Availability of Navigation Devices*, CS Docket No. 97-80; *In the Matter of Compatibility Between Cable Systems and Consumer Electronics Equipment*, PP Docket No. 00-67

Dear Chairman Powell:

I write on behalf of Veridian Corporation ("Veridian") to inform you of a digital copy protection technology developed by Veridian that has the ability to protect content broadcast over-the-air as well as delivered by cable or satellite. The technology can aid significantly in the nation's transition to over-the-air digital television ("DTV"). To fulfill that potential, it is essential that the Commission remain vigilant to ensure the compatibility of navigation devices both with multichannel distribution systems using Veridian's VeriFIDES™ method technology and with systems using the method set forth in the current Interim Specification proposed by Cable Labs.

The CableLabs interim specification does not cover digital over-the-air reception; rather, it depends upon the availability of a back channel between the distributor and the reception equipment and supports only 'conditional access' (whether or not content may be viewed) and elemental 'copy control' (allowing no copies, a single-generation copy, or unlimited copying). By contrast, in Veridian's technology the content can be denied or can be obtained only with full copy protection (controlling the use of any copies that may be made). While the CableLabs proposal is a step in the right direction, it is incomplete and might effectively be picking a winner: the inclusion in the specification only of the "5C" protocol means as a practical matter that equipment manufacturers would be disinclined to also incorporate an alternative digital copy protection technology. A consumer with a 5C-only navigation device would then not be capable of switching it from a cable system using the 5C protocol to a system using the VeriFIDES™ technology, in violation of the statutory requirement of interoperability. The proposal could also be picking a winner that does not offer the digital protection needed to stimulate the availability of content for DTV. The Commission should ensure that the specification is open to both technologies. This will help achieve both Congress's DTV

transition objectives and the Commission's statutory mandate under the navigation devices provision – a mandate that the Commission has set out to implement in the above-captioned proceedings.

Congress and the Commission have chosen to prescribe a "rapid"[1] and rigorous timetable for the transition to DTV in light of its well-recognized benefits to the public.[2] At the same time, Congress has recognized the importance of removing the barriers that "could slow the transition to digital television and undermine localism and the future of free, over-the-air television." Letter from United States Representative Billy Tauzin, et al. to Michael Powell, Chairman, FCC (dated Mar. 2, 2001) at 1.

One such barrier has been the absence of a digital copy protection technology that works both for cable and satellite systems utilizing conditional access and for the reception of digital over-the-air signals.[3] For example, the Interim Specification developed by CableLabs for the OpenCable POD Copy Protection System is intended to protect content on the interface between the Point of Deployment ("POD") module and the host navigation device, and does not protect "free access off-air broadcast content." *See* Cable Television Laboratories, Inc., OpenCable POD Copy Protection System Interim Specification (dated May 15, 2001), at 1. While over-the-air digital broadcasts are

---

[1] *See* 47 U.S.C. § 309(j)(14) (mandating recovery of analog spectrum by 2006); *see also In the Matter of Review of the Commission's Rules and Policies Affecting the Conversion to Digital Television*, Report and Order and Further Notice of Proposed Rulemaking, MM Docket No. 00-39, 16 FCC Rcd. 5946 (2001) ("*DTV Conversion Review*") (stressing the Commission's desire for "rapid" conversion to DTV); *In the Matter of Advanced Television Systems and their Impact Upon the Existing Television Broadcast Service*, Fifth Report and Order, 12 FCC Rcd. 12806 (1997), ¶ 76 (establishing a May 2002 deadline for all commercial television broadcast stations to complete construction of their digital broadcast facilities).

[2] *See In the Matter of Commercial Availability of Navigation Devices*, Further Notice of Proposed Rulemaking and Declaratory Ruling, CS Docket No. 97-80, 15 FCC Rcd. 18199, 18204 (2000) ("*Navigation Devices*" proceeding) (as the transition continues, "consumers will benefit from the myriad advantages offered by digital content").

[3] The Commission has noted the broadcasters' "belief in the critical importance of adequate copy protection for broadcast television," for, among other reasons, "the rapid implementation of DTV." *In the Matter of Compatibility Between Cable Systems and Consumer Electronics Equipment*, Report and Order, PP Docket No. 00-67, 15 FCC Rcd. 17568, 17579 (2000) (quoting comments of the four largest commercial broadcasting networks) (internal quotation marks omitted).

delivered to the consuming public without need for a subscription, the need for digital copy protection remains important for such broadcasts as well – it is needed to spur content providers to allow the free digital broadcast of content to consumers without the fear of commercial duplication.[4] Moreover, with a "one-stop" digital protection technology, the consumer electronics community can develop reception equipment and navigation devices that protect digital content no matter how transmitted, whether over the air, by cable or satellite, over the telephone line, via DVD or by Internet (which itself may be delivered by cable, satellite, or telephone). In my view, the protection of broadcast content and the deployment of electronics equipment equipped to incorporate that protection may help jump-start the now vexingly slow pace of the DTV roll-out.

Veridian's VeriFIDES™ technology accomplishes these objectives. VeriFIDES™ can augment systems using conditional access, or can be used in the absence of conditional access, while the method currently endorsed by the CableLabs specification does not protect content either broadcast over the air or sent to the PC. VeriFIDES™ technology does not require a "conversation" and/or a "handshake" between connected pieces of equipment. Rather, a "ticket" (an electronic license) is used to control what and how content may be copied and accessed.

Equally important to content owners, VeriFIDES™ allows ongoing control even after (encrypted) content is passed across to personal computers or the Internet. VeriFIDES™ technology allows encrypted programs to be digitally recorded and played again, *under the control of a ticket*, with full content-owner or distributor authorized access control. The CableLabs interim specification precludes such time-shifting, which has become an expected consumer 'right'. While allowing recording and *authorized* playback, VeriFIDES™ capability would prevent the unauthorized access to copies of streamed video content as has occurred with audio content. I believe that these characteristics are required in the quest for "a commonly accepted copy protection system [that] will be a major factor in promoting the rollout of digital television via cable and ensuring that content providers make high value content available to the cable industry." Letter from Gary Shapiro, President and CEO, Consumer Electronics Association, to Honorable Michael K. Powell, Chairman, FCC (dated May 10, 2001). It is important to note that, unlike other systems, the protection of content using VeriFIDES™ would not require that

---

[4] The Commission acknowledged the importance of digital copy protection to DTV transaction in the *Navigation Devices* proceeding: "Commenters note that the copy protection and security provisions at issue are critical to ensuring that content providers will supply high quality programming necessary to further the digital transition." 15 FCC Rcd. at 18206 (citing comments of broadcast and cable interests). And the Commission has explicitly confirmed the importance of increased availability of digital programming to a rapid transition, stating "We agree that the wide availability of digital programming . . . will help speed the transition to DTV." *DTV Conversion Review*, 16 FCC Rcd. at 5950.

all consumer equipment incorporate the technology – only those consumers wishing to view protected content would need to acquire such equipment.
While the specification is a good step in the right direction, it is incomplete in not expressly incorporating a technology such as VeriFIDES™ that covers digital over the air broadcasts and also allows time-shifting. The consequence of this may be the employment of a system that is not only incomplete in terms of providing protection (as it offers only conditional access and elemental copy control), but also effectively precludes adoption of a more complete protection system that enables time shifting and continuing access control on any copies. The entrenchment of an incomplete system of protection in this manner will only exacerbate content owners' reticence to increase the flow of high quality digital programming recognized as necessary to further the digital transition.

The inclusion of VeriFIDES™ will therefore help significantly with DTV transition both because it covers digital over-the air transmissions and because the specification would not preclude the choice of the best digital copy protection system by the market.

In addition, an open specification that incorporates VeriFIDES™ is also essential to ensuring navigation device compatibility, *see* 47 U.S.C § 549.

Incorporating only the back channel/conditional access technology would mean that devices employing only this technology would then be incompatible with systems using a technology such as VeriFIDES™. If, for example, a cable system adopts only the back channel/conditional access technology, which then must be contained in the cable boxes operating within its system, a consumer who owns such a box cannot use the box if he or she decides to switch to another competing programming distributor that utilizes VeriFIDES™. The consumer must instead purchase or lease yet another set-top box. This lack of device compatibility would be inconsistent with the goals of both Congress and the Commission. As discussed above, these goals would best be promoted by allowing the market to choose between available systems for access and copy protection, and giving the market an opportunity to sort out which technology, or combination of technologies, consumers will find to be most convenient and desirable in navigation devices.

In sum, the Commission should ensure that the final released specification will extend to all modes of digital content transmission and will also incorporate the VeriFIDES™ technology.

Veridian is a privately held $700 million corporation, providing information technology solutions and engineering services to the intelligence community, Defense Department, civilian agencies, and private sector. The World Airline Entertainment Association has adopted a specification for content protection of digital in-flight movies based on VeriFIDES™ technology. We would be pleased to provide additional information to you, your staff, or to others. You can reach me at 703-575-3135. For your reference I have

enclosed a reprint of "Persistent Access Control to Prevent Piracy of Digital Information", Proceedings of the IEEE, July 1999.

Yours truly,

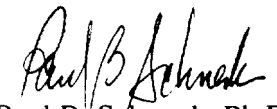

Paul B. Schneck, Ph.D
Chief Technology Officer

cc: Commissioner Kathleen Q. Abernathy
Commissioner Michael J. Copps
Commissioner Kevin J. Martin
Bruce Franca, Office of Engineering and Technology, FCC
Rick Chessen, FCC DTV Task Force
Barbara Esbin, FCC DTV Task Force
Tom Horan, FCC DTV Task Force
Keith Larson, FCC DTV Task Force
Jonathan Levy, FCC DTV Task Force
Mary Beth Murphy, FCC DTV Task Force
Alan Stillwell, FCC DTV Task Force
Gary Shapiro, Consumer Electronics Association

Persistent Access Control to Prevent Piracy of Digital Information

PAUL B. SCHNECK, FELLOW, IEEE

Marshall McLuhan's statement that "the medium is the message" captured the reality of communications until the digital age. Communicating digitally involves the transfer of information as streams of bits. No longer does matter need to be moved. In the case of data representing continuous waveforms, the communication is complete when the recipient possesses a digital representation of the original signal with a given fidelity criterion. This allows one to embed watermarks to deliver hidden information. In the case of discrete data files, the communication is complete when the recipient possesses the identical stream of bits delivered by the source: at that moment there are two, identical, streams of bits. It is impossible to characterize one as the "original" and the other as the "copy." Consequently, one loses control of digital information after transmitting it to another. The recipient is free to use their "original" as they see fit. The ubiquitous PC and Internet provide the means easily and inexpensively to copy and distribute data files to anyone wishing to do so. This paper introduces a framework for "persistent access control" over digital data. That is, data remain under the control of their originator, even after another has received them. This requires the use of multiple levels of encryption, transmission of ciphering keys through licenses associated with the data, and of trusted hardware and trusted software (access control mechanism) to ensure that the originator's restrictions are not circumvented. We show how a PC or consumer electronics appliance can accommodate the required system—while remaining compatible with previously produced data. This material is the subject of a pending U.S. Patent Application (allowed) and foreign applications.

Keywords— *Access control, copyright control, cryptography, license management, secure processor, trusted hardware.*

I. INTRODUCTION

From the stone-chiseled art of the Paleolithic age to the incunabula of the fifteenth century, the message was closely bound to the medium. Copies would easily be recognized as such. Copying required as much effort as making the original. Widespread distribution of printed works had not yet become a reality. In the case of poetry, which could be separated from the medium that carried it, the early Hebrew authors devised an effective manner of ensuring acknowledgment of their authorship: they wove their names into their poems. Removing the author's identity would damage the poem.

As the new technology of printing allowed wide distribution of material, particularly of "dangerous ideas," royalty felt the need to authorize and control printing. The Statute of Anne provided printers with a monopoly on production in exchange for the Crown's control over what might or might not be printed.

Advances in the technology of printing, and later advances in the technology of copying, have made it possible to separate the message from the medium and to make high-quality copies at relatively low cost. Fortunately, technology also offers tools uniquely to identify individual printed copies by marking techniques, those techniques being incorporated in some current commercial products.[1] Copyright law has also been developed to preserve the property interest of originators in their ideas and creations.[2]

The use of digital media, where ideas are expressed as streams of ones and zeros, raises a new set of issues. In the digital world the concept of "copy" is irrelevant. When two files consist of identical streams of bits, how can one be called an "original" and the other a "copy"? Such a distinction might attach to media, but it makes no sense to label their contents in this way.

In the digital world, we have two kinds of sources to be envisaged.

- Waveforms (audio, image, or video signals) for which the transmission criterion is characterized by a fidelity criterion, which means that there exist several digital representations of a waveform signal for a given fidelity criterion. That variability is used by many watermarking techniques.
- Discrete sources, e.g., software or data files, for which an exact replication of the source at the destination is demanded.

Previous attempts at protection have focused on ensuring that only authorized users have access to particular files. However, even if one could limit access of files to only those users with authorization, a problem arises as a corollary of the fact that all bit streams are effectively "original." Thus, anyone possessing a bit stream can manipulate it in

Manuscript received March 3, 1998; revised April 8, 1999.
The author is with MRJ Technology Solutions, Fairfax, VA 22030 USA.
Publisher Item Identifier S 0018-9219(99)04969-5.

[1] For more information see http://www.xerox.com/xsis/countdet.htm.

[2] For more information see http://www.house.gov/judiciary/4.htm.

0018–9219/99$10.00 © 1999 IEEE

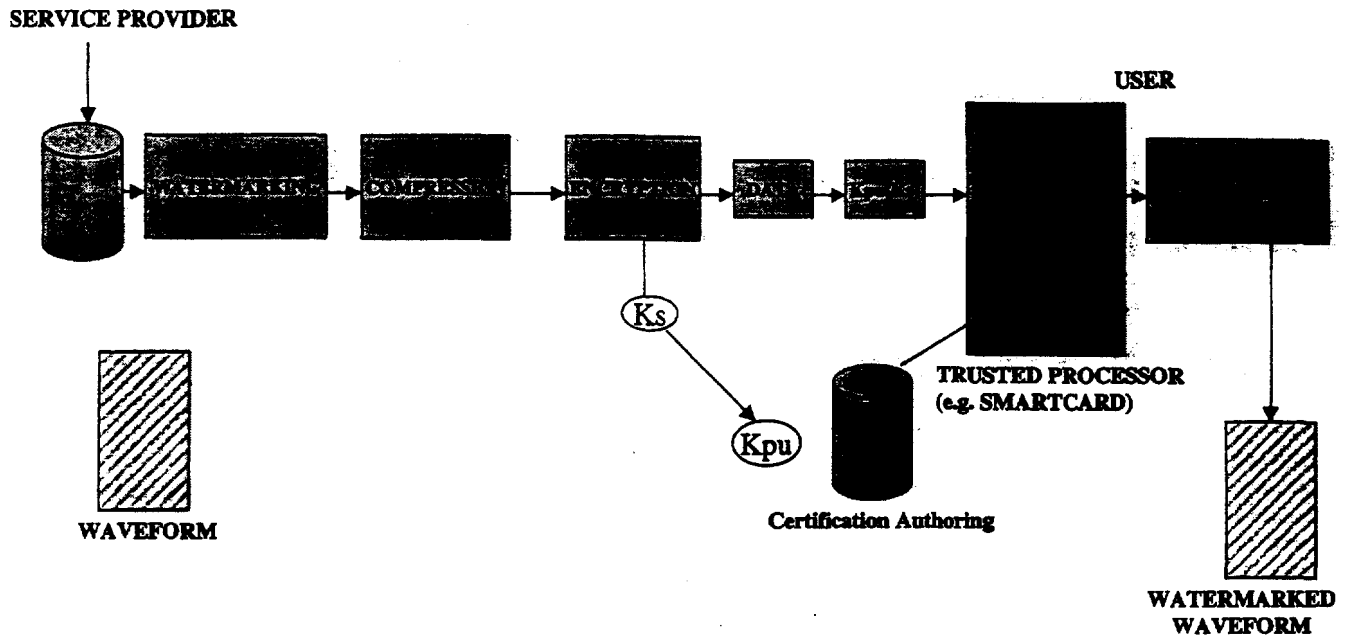

Fig. 1. Analog waveform data delivery.

precisely the same ways as could its creator. This means that any user of a file would be free to redistribute the file, print its images, etc. We introduce in this paper a framework based on the management of licenses associated with data files, which, when used with encryption and trusted platforms (hardware and software), ensures persistent protection of the files. In Section II we review some of the techniques used to limit access of files. Section III describes our proposed framework. Section IV gives some application examples, and Section V concludes the paper.

II. Access Control for Digital Information

Access control [12] is generally based on ciphering (encryption) of data by a secret-key algorithm (typically DES or IDEA) with a service, or session, key. This service key is sent to the recipient through a secure channel, often a public-key infrastructure (PKI), managed by a certification authority (see [1] and [2]). In the case of delivery of waveforms (see Fig. 1) persistence of control on the original signal is achieved through the introduction of watermarks, remaining after deciphering. In the case of delivery of data files (see Fig. 2), persistence of access control is achieved through limiting deciphering to (very) restricted conditions (trusted hardware and software). We begin with an overview of classical approaches to persistent access control, namely watermarking for waveforms and persistent cryptographic coding for data files.

A. Watermark-Based Persistent Control

Watermarks, familiar as the faint images on bond paper, are subtle patterns, barely perceptible (if perceptible at all) to the user. They may contain information that identifies the content owner and the purchaser or user of protected material. They are inserted in material in such a way that they do not interfere with the intended use of the material. That is, they do not materially degrade the aural or visual qualities of a sound track or video. Depending on the specific algorithm employed to place a watermark, users may not be able to find the watermark or, even if they locate the watermark(s), should not be able to remove them. The content owner, knowing the algorithm and parameters used to place the watermarks, is able easily to locate them and check that the material has been legitimately acquired. Content owners frequently use this technique to deter unauthorized placement of material on websites (see [16]).

Perhaps the earliest example of a watermark-like technique occurs in the Hebrew poems of the Jewish High Holidays liturgy. A poet would weave the spelling of his name into the structure of a poem. Although the poem might be copied, there would always be an acknowledgment of the author. Any attempt to remove or alter the author's name would require the substitution of words that match in rhythm, meaning, part of speech, as well as style—a most difficult proposition.

The very nature of watermarks as subtle patterns is the source of a significant limitation in their use. Compression techniques, employed to reduce the space and transmission time of files representing sounds or images, often fail to preserve the low-order (least significant) bits. These are precisely the bits on which the watermark is carried. Use of other bits would render the image (too) degraded. The design of imperceptible watermarking techniques strongly resistant against several compressions and decompressions leads to interesting results [4].

For audio signals, there is deeper understanding of how to generate information that is not perceptible to listeners [5]. Identification of the purchaser can be encoded on these areas. To further make the marking robust despite compression or other transformations that do not preserve low-order bits, the watermark information is typically encoded many

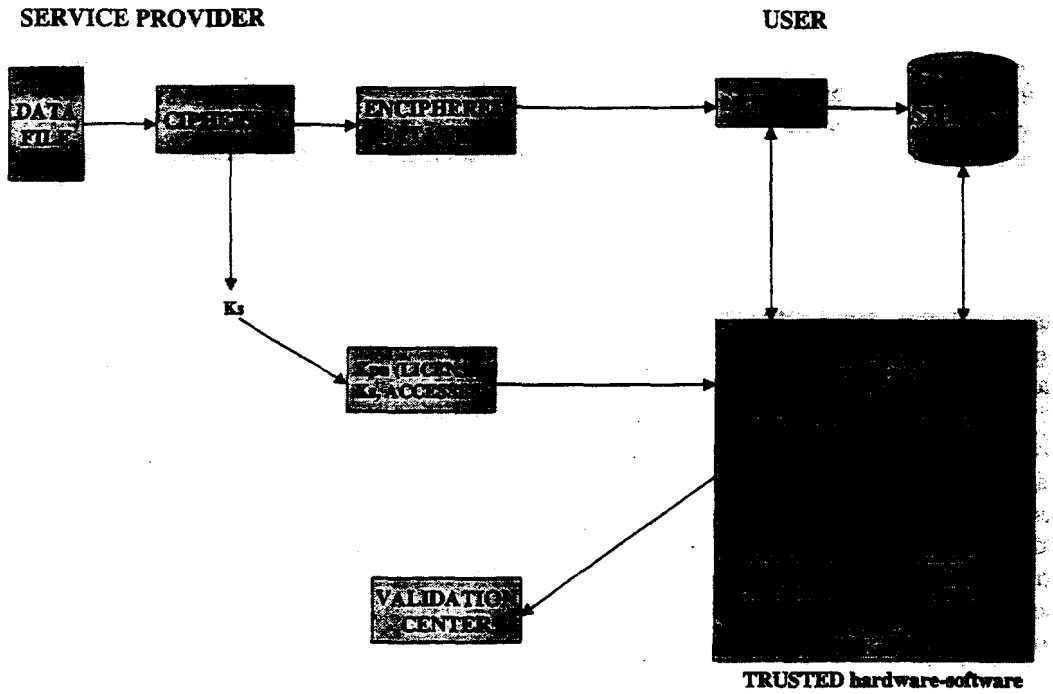

Fig. 2. System-level data flow.

times in small sections. Without knowledge of the encoding process and the parameters used for a specific recording (comparable to encryption algorithm and session key) it is difficult to locate (and remove or alter) the watermark. The intellectual property owner, knowing the encoding algorithm and parameters, can easily scan copies of material of suspected provenance to determine their authenticity.

When dealing with text, there are no low-order bits on which to impress a watermark as each byte is the coding of a particular letter. There are three ways in which information can be hidden in a text document: by changing intercharacter spacing; by changing interline spacing; and by changing features of the character font. Such changes are generally not noticeable to humans. Changes in spacing are generally impervious to copying, even including changes in size during copying. Changes produced by adding noise to character bitmaps is not robust under copying but is harder to detect. Scanning text and performing optical character recognition will remove information hidden by these three techniques. Consequently, these techniques are generally applicable only to identifying printed textual information or photocopies.

It is difficult to add a watermark to an executable program as there are few, if any, low-order (i.e., insignificant) bits that may be modified without impairing the function of the program. Consequently, the use of watermarks is generally restricted to "content," and we will not follow this path in this paper.

B. Persistent Cryptographic Coding

Encryption by efficient ciphering algorithms (such as DES or IDEA see [1] and [2]) is becoming widely used to aid in the prevention of piracy. The owner encrypts information and transmits it to the user. The decryption key (presumably much smaller in size and therefore more easily protected than the information to be transmitted) is provided to the user through a secure channel (e.g., opened through a PKI). However, by itself, encryption only provides a mechanism for ensuring that while information is in transit from owner to user a third party can not steal it. The user, having obtained the key, is able to copy or use the information as if he or she were the owner. Consequently, additional mechanisms have been explored to prevent copying the deciphered data, such as securing disks and I/O devices.

1) Dongles: The etymology of "dongle" is unclear, however the word may be a corruption of "dangle," as the device is a small unit that typically plugs into the printer port of a PC. Software to be protected checks for the presence of the dongle. If the dongle is not found, the software (or some of its features) will not operate. "Passive" dongles are relatively easy to defeat by patching the software to take the path corresponding to having ensured the existence of the dongle. To make hacking more difficult, some dongles are "active" and, in response to data, provide a cryptographic transformation that is used to unlock or decrypt portions of the software. In this case, the would be hacker must obtain the software and dongle and traverse each segment of code that requests a key, recording the key and patching the software. The software may not be usable until all execution paths have been traversed and captured. The level of effort and required expertise to accomplish this is usually sufficient to discourage small-scale theft.

2) "Uncopyable" Disks: Uncopyable disks represent a low-cost, convenient implementation of a dongle. The protected software (or some of its features) will not operate

without the disk. In addition to suffering from the same vulnerabilities as the dongle, hackers are usually able to copy these supposedly uncopyable disks.

The disks were uncopyable by the standard driver routines of the operating system. The software manufacturer typically provided special driver routines capable of reading the disks. For example, a specific error condition was forced and the special driver routine ignored that normally fatal error. Alternatively, a nonstandard format is used, such as recording data at a higher density than is possible with standard driver routines. No matter what the specific approach, the hacker can reverse engineer the driver routines and copy the "uncopyable" disk.

3) Cryptographic Containers: Cryptographic containers represent an attempt to extend the influence of the owner beyond furnishing the decryption key to the user. So-called "trusted software" is used to access the information sent in the container. The trusted software is supplied by the information owner and requires that a key or other identifier be obtained from the owner prior to decrypting the information in the container. In addition, the trusted software may implement some restrictions imposed on the use of the decrypted information. Two significant weaknesses are immediately evident. First, the key or password supplied by the owner may be passed along to another user. As a disincentive to passing the key on to another, the user's identity may be encoded within the key. Second, the trusted software may be corrupted so that it provides the key and/or decrypted information to the user or so that it no longer restricts use of the decrypted information (e.g., it allows copying or printing). In some implementations, the information, once obtained, is available for reuse. This business model is followed when the information provider charges for each initial access to a datum because:

- the long-term value of the information is low;
- the cost to duplicate and store information is nearly the same or even greater than its purchase cost;
- the user is deemed trustworthy;
- the economic risk due to piracy is deemed sufficiently low.

Improvements to cryptographic containers by the distribution of a part of the access keys through a clearance center has been proposed in papers [3] and by some products, such as IBM's "cryptolope".[3]

4) Consumer Electronics Platforms: In consumer electronics appliances, proprietary solutions offer an easier way of ensuring persistent cryptographic coding than is the case for the ubiquitous PC. As those devices are generally used for reproducing signals that originated as analog waveforms (sound, image video), watermarking is an applicable protection technology.

a) Digital video disks (DVD's): DIVX recently announced a proprietary system that provides copy protection for DVD's. In the DIVX approach, the initial fee paid by the consumer is a license fee entitling the purchaser to a fixed period (48 h) of viewing. That viewing period may begin at any time. Once begun, the consumer may view the movie as many times as desired within the fixed period. If the movie is viewed after the conclusion of the first fixed period, another fixed period begins, for which the consumer will be billed an additional fee.

Copy protection is provided by virtue of proprietary data formatting and a secret encryption scheme protected by tamperproof circuit packaging. The circuit module is designed so that it will not be usable, nor will it reveal useful information, if it is taken apart. Output from the system is available only in analog form such as NTSC video produced by an analog protection system such as Macrovision.[4] In addition to these technical measures, patents cover the technology necessary to read the DVD. This provides a mechanism for dealing with unlicensed manufacturers of players.

In order to track usage of DVD's, the player must transmit the player owner's identity and billing information to a processing center. This information must be made available at the time of purchase, when a user account is established. After each transaction the processing center provides an (encrypted) acknowledgment to the DIVX player. The player will cease functioning if, after a reasonable period of time, say a month, it has not received an acknowledgment from the processing center (e.g., because it was prevented from contacting the center). The integration of a part of those access control messages through watermark signals is proposed in [7].

b) Set-top boxes for cable TV: The threat model for cable distribution of analog television signals has been focused on unauthorized access to the cable signal and not on piracy and copying [8]. The set-top box contains the unique customer address that is used to send programming instructions and keying information. A modest level of tamper-resistant or tamper-detecting packaging is frequently used to safeguard the address information. If tampering is detected the device is rendered inoperable. When the customer requests a repair of the "faulty" set-top box the technician can determine that tampering took place. Under this threat model, the loss of revenue is limited to those individuals that have defeated the tamper-resistant packaging.

A paying subscriber can easily record the signal as it leaves the set-top box and enters the television receiver. As described above, analog protection schemes are used to render the signal "uncopyable" by most VCR's. This approach is generally successful against the individual subscriber that is the focus of the threat model.

The situation will change dramatically when digital program material is broadcast. Today's approach ensures that a user must pay to receive access to a particular program. When digital program material is received, a viewer can copy and redistribute the digital stream. The threat model

[3] For more information see http://www.software.ibm.com/security/cryptolope.

[4] For more information see http://www.macrovision.com/dev.html.

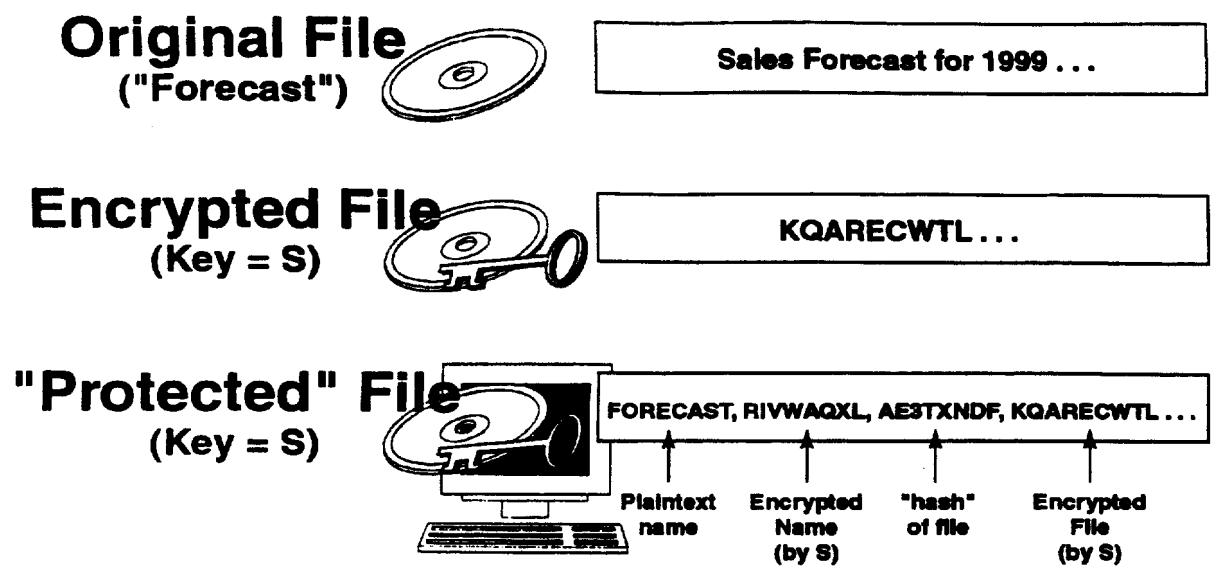

Fig. 3. Structure of a protected file.

must now also include unauthorized copying and redistribution. An adequate solution has not been available. Consequently, many studios have stated that they do not intend to release program material in digital form until they can be assured of protection from copying and piracy.

III. A Persistent Access Control Framework

We develop a specific model which is compatible with previous hardware and software and achieves copy-protection security through the mandatory use of trusted I/O mechanisms.

A. The Concept

Our framework is based on a combination of trusted hardware and a low-overhead license management mechanism that ensures persistent control of digital data. The system allows the open distribution of data ("content") in a protected (encrypted) form. Users may be sent data by a distributor (the "push" model) or may retrieve them from a server (the "pull" model). The transmission medium is irrelevant. In order to access the data, an electronic license must be present. The license contains the information necessary to determine whether or not an access is permitted and the information needed to decrypt the data. The "access control mechanism" in the user's PC or consumer-electronics device controls all I/O operations and, using the information in the license, allows or disallows specific accesses.

A license may have been sent with the data or the user may need to obtain a separate license. In a consumer scenario, a user might purchase a CD and a license for access at the same time. The price of the license would be a function of the types and duration of accesses and uses requested. In another scenario, a user might obtain data from an Internet server, along with limited free access. To obtain further use of the material the user would purchase a license using the Internet, by phone, by mail, or in person.

Transactions within an enterprise or between enterprises might focus on access control. In these cases the content owner would likely distribute data and licenses to specific addressees. If the license permits, derivative products could be redistributed to others. These secondary recipients would each need to obtain a license—either from the originator or from the creator of the derivative product, depending on the license arrangements made between the originator and the redistributor.

Executable software that is aware of the access control mechanism can interrogate the mechanism to determine whether or not a feature is allowed by the license controlling the program. This enables distribution of multiple levels of software conveniently within a single package.

B. Implementation Overview

The objective of the system is to retain control of data for the data owner while allowing users freedom to utilize data to which they are granted access. The basic concept is to secure data by encryption whenever they are external to the trusted hardware system. When data are within the system they may be decrypted and made available for processing or reviewing by the user (as the owner authorizes). The system intercedes to ensure that no data may be released (output) unless the user is authorized to do so. The mechanism for capturing the data owner's limits and delivering them to the user's system is the electronic license. Without a license a user cannot access the data.

1) Encryption: Protected data are encrypted whenever they are outside the trusted hardware system. The data owner selects a secret key and an encryption algorithm. The secret key will not be exposed to the user. It will be used to decrypt the data within the trusted hardware. Because protected data are encrypted they may be openly transmitted and distributed.

The data owner encrypts the data elements that are to be protected (see Fig. 3). A header record is prefixed to the

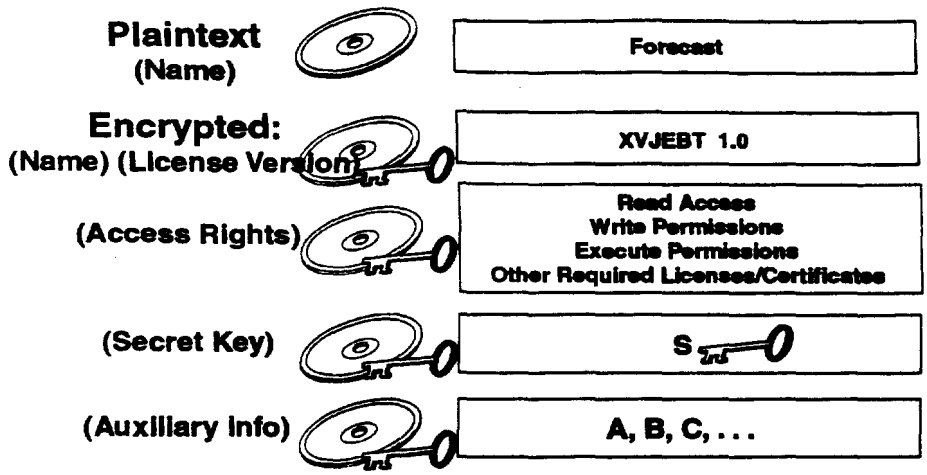

Fig. 4. The structure of the electronic license.

data. The header record contains an identifier[5] denoting the file's protected status, the unencrypted name of the file, the structure of the file—which parts are plaintext and which parts are protected (encrypted), and an encrypted hash value (of the file) that is used to bind the header to the encrypted file. This binding ensures that the only the valid header can be associated with the encrypted file.

A file is recognized as containing protected data when a header record precedes it. Removal of the header record will prevent the system from recognizing a file as protected. Although this may appear to be a vulnerability, it is not. Without a header the access mechanism will not recognize the data as protected and will not decrypt the file. Only the data owner can digitally sign and bind a header to a protected file that will match a license created by the data owner and containing the secret key to decrypt the protected file.

2) Electronic License: The electronic license contains the name of the file, a list of access permissions, and the secret key needed to decipher the encrypted file (see Fig. 4). One or more of the access permissions may require that the license be signed by the data owner or by other parties. The license may also contain or require other licenses or digital certificates.

In order to allow rapid processing, access permissions may be specified in three different formats.

- A bit-vector provides rapid access to the most frequently used access criteria (e.g., print permission).
- A parameter list provides the values used in determining whether or not access is to be granted (e.g., a date window).
- Executable code allows arbitrary computations to determine whether or not access is to be granted. Note that the data owner is responsible for providing the code and ensuring its accuracy, reliability, and trustworthiness.
- The license is encrypted and then prefixed with the plaintext file name. Doing so allows the system rapidly to locate a license. The plaintext name is used only to improve performance. The name within the encrypted license is used to verify the identity of the license.

The key used by the data owner to encrypt the electronic license is described later.

3) Trusted Hardware System: Secure hardware includes devices with embedded secret keys allowing a service provider to require a recipient to authenticate himself before receiving data. That kind of hardware is largely described in the literature [9], [10]. A low-cost solution can be achieved by the use of, e.g., smart cards (see [15]) in which the secret key may be stored, providing mobility to users.

The security of the system depends on the hardware to retain protected data within the system unless the data owner (as expressed in one of the permissions granted in the online license) permits release. Each PC will contain a unique private key and an associated public key. The public key is used to encrypt electronic licenses intended for the PC. The trusted hardware system also protects the cryptographic variables—they are never made available to the user. To simplify this exposition, the ubiquitous personal computer is used as the frame of reference. A subset of this approach would be used in consumer electronic devices.

C. Access Control Mechanism

The firmware (in the case of the use of a BIOS) or the drivers and the I/O elements of the OS that provides basic services to application programs are replaced by new modules including an access control mechanism that mediates all requests for input or output (see Fig. 5). Any request for an input or output operation from an application is handled by the new I/O routines including the access control mechanism. Therefore, any attempt to execute directly input or output operations results in the execution of the access control mechanism.

Implementing this would depend on establishing a secure environment (e.g., memory protection, privileged instruction limitations, etc.) and would be based on a trusted operating system (at least with respect to I/O routines).

[5] The identifier can be made large enough (i.e., bit length) so that the probability of an unintentional match is insignificant.

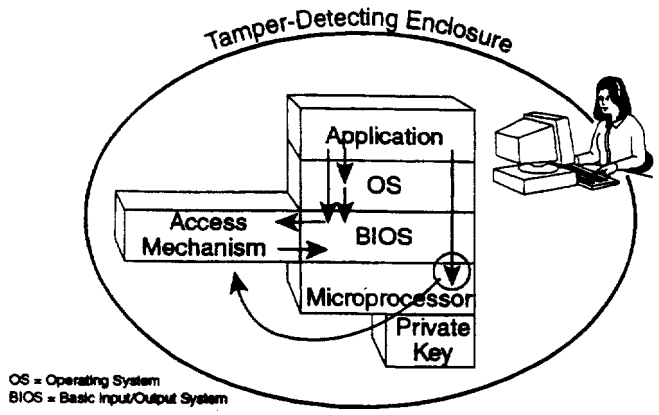

Fig. 5. BIOS-based access control mechanism.

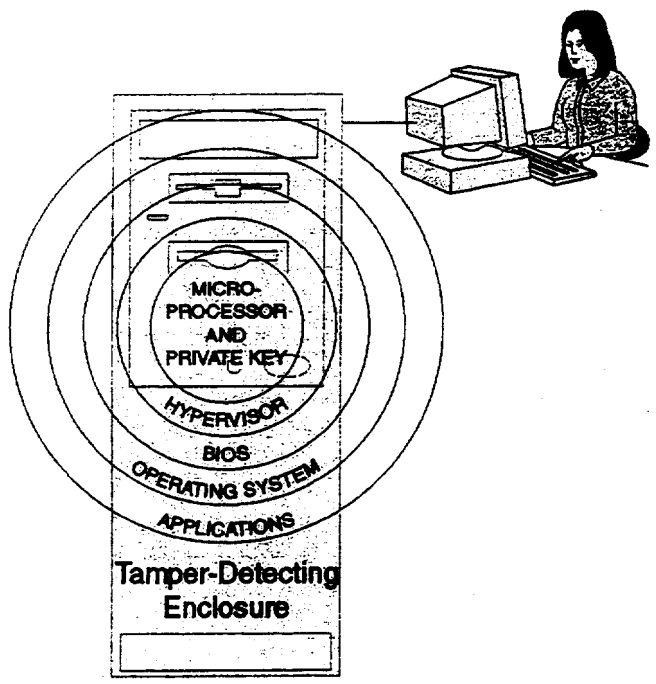

Fig. 6. Hypervisor.

Additionally, each new operating system release (and each different operating system) would require coordination and verification of the access control software.

An alternative approach, that is independent of the operating system, is for the access mechanism to be part of a "hypervisor," an operating system-like program that creates a virtual system image (in the same fashion as the operating system creates a virtual memory image for application programs). The hypervisor is interposed between the operating system and the hardware and is unseen by the operating system (see Fig. 6). The resulting virtual system appears to function as the actual system except that all input or output operations and all privileged operations are trapped and analyzed by the hypervisor. Any attempt to execute directly input or output operations results in control being passed to the access control mechanism where the restrictions specified in the electronic license are enforced. The specific actions taken by the access control mechanism are described below. Because the access mechanism has no awareness of application tasks, it must treat the entire computer as if only one task were running. This means that the access restrictions imposed on any active task are applied to each of the tasks in the system.

In either approach, the access mechanism, although implemented in software, is protected by hardware and is part of the trusted hardware system. The access control mechanism ensures that protected data are accessed only as provided in the electronic license. The access control mechanism will be small and can be made secure.

1) Input: The first time an input file is accessed it must be opened. The access control mechanism reads the first record of the file and determines whether or not it is a header.

If the first record of the file is not a header, the file is not a protected file and the access control mechanism returns to the original kernel code. Processing continues as in an ordinary operating system (OS). This retains compatibility with all unprotected files. Programs and files may continue to be used without modification.

If the header of a protected file has been removed in an attempt to circumvent the protection system, the file will not be recognized as protected. This causes no loss of protection as the system will not seek a license and will not decrypt the data read from the file.

If the first record of the file is a header, the file is a protected file. The access control mechanism first verifies that the header and file are correctly linked and have not been altered. If verification fails, an error action follows. If verification succeeds, the access mechanism searches for a matching electronic license. There are three cases.

- If a matching electronic license cannot be located, the access control mechanism displays a dialogue box stating: "Provide the electronic license path, abort, deny access, or cancel." The user can provide the path to the electronic license, at which time the access control mechanism again searches for a matching electronic license. The "abort" option is self explanatory. The "deny access" option allows the program to continue running, but without access to the denied file (e.g., the system returns the status "file not found" to the program).
- When more than one matching electronic license is located, the access control mechanism displays a dialogue box listing the licenses and stating: "Identify the electronic license to be used, abort, or cancel."
- When exactly one matching electronic license is located (or identified), the access control mechanism proceeds to interpret the license.

The first step in interpreting the license is decryption. This is accomplished with the PC's unique private key, available only to the access control mechanism. After the license is decrypted, it is checked for integrity (e.g., to ensure that it matches the external plaintext name used for rapid matching). If the license cannot be decrypted (generally because it was intended for a different PC) or fails

the integrity check, the user is asked to provide a different license. If the license passes these tests, the access permissions are examined to determine whether or not the file may be opened. If access to the file is denied, the system will either abort the program or return "file not found" as above.

When access is granted, a system-wide flag is set, indicating that protected files are in use. This flag is used to ensure that ensuing output requests are subject to the owners' access controls.

On input operations, the file owner's secret key is obtained from the associated license and used to decrypt data as they enter the machine. Commodity chips are available to perform encryption and decryption on the fly (e.g., as they are read from disk) without performance penalty (see [11] as an example). From this point on, processing proceeds as if the data had never been protected. Further intervention occurs only when there is a request to output data.

2) Output: When a write request is processed, the access control mechanism first checks the system-wide flag to determine whether or not protected data are involved. If not, special processing is not necessary and control is returned to the ordinary output routines. If the system-wide flag is set, all open licenses are checked to determine whether or not the write operation may proceed and what, if any, restrictions might be imposed. The conditions of each license are intersected, yielding the most restrictive access. If output is denied, the program may be terminated or the access control mechanism may return a "successful completion" status without performing the write operation.

If output is permitted, a header is appended when a protected output file is opened or when the file's status changes to protected. The header indicates the structure of a protected file (as described above).

The secret key used to encrypt the protected portions of a protected output file can be a key from one of the input files or can be a new secret key generated by the access mechanism. In the former case, the user of the output file would obtain a license from the owner of the input file. In the latter case, the user/creator issues licenses for this "derivative" file.

3) Memory Protection: The access mechanism ensures that each program operates in data-fetch-protected and in data-store-protected memory segments. Doing so ensures that users cannot write programs to "see" or copy executable code. Nor can users modify executable code. Any attempt to do so would result in a memory-protection interrupt that would turn control over to the access control mechanism.

Instruction–fetch–protection is also utilized. Protected programs will only run when invoked by an authorized user. This ensures against running code as an unauthorized subroutine of a coresident program.

4) Software Features: A new feature is the ability to request the status of software features. The access mechanism responds by checking the electronic license(s) and returns a value indicating whether or not a feature is available.

For example, a program may inquire as to the current availability (there may be a time- or situation-dependent response) of a feature. A zero response indicates that the feature is not currently available to the program. A nonzero response indicates that the feature is available. The program can use the value of the response for further information.

Software cannot be "hacked" or modified because it is distributed in encrypted form and is not visible to users. Further, even if users knew the binary structure of the software they would be unable to modify the software if the license required that it be executed at the highest level (i.e., not as a subroutine) or in a separate, protected area of memory. The system realizes the vision of "trusted software" that has until now been only an abstraction.

D. Trusted Hardware

1) Tamper-Detecting Enclosure: The system as described so far is vulnerable to physical intrusion. A determined thief can obtain the plaintext version of a file by capturing the bit stream as it is transferred to memory. Alternately, the thief could reset the system and then start a program that would capture the contents of RAM. Protection from this class of attack requires barring entry to the system.

A tamper-detecting enclosure provides a signal shortly before an attacker can gain physical access to the system. When tampering is detected, the private key, secret keys, and plaintext data are destroyed. An example of how this is done is given below.

The smaller the period of advance notice that is required, the higher the degree of protection that can be provided. The requirement for only a few microseconds of advance notice ensures that the tamper-detection signal will result in effectively blocking access. The tamper-detection signal triggers a stand-alone circuit to write a series of random values over the private key and over any decrypted secret keys. This sequence takes less than a microsecond for up to several secret keys. Power is then removed from these circuits (possibly "crowbarred," to ensure immediate loss of stored data). A long-life battery or other energy source provides the energy for erasing the keys. If the system is operating, power is removed from the main RAM storage, protecting any decrypted plaintext data.

In the event of tampering, the private key will be destroyed and the system will no longer be able to access any protected data. In order to restore the system to full service, it must be taken to an authorized repair facility and a new private/public key-pair created.

The tamper-detecting case must prevent against a straightforward opening of the case, against cutting or drilling, and against thermal changes designed to render it unable to detect tampering. The design must also provide radio frequency protection so that a receiver cannot copy data flowing within the system.

2) Nonvolatile Storage: To protect against a "sledge hammer attack" where the PC might be destroyed in order to gain unfettered access to the hard disk, all protected data (as indicated by the system-wide flag) are encrypted before writing and decrypted upon reading. Writing or reading internal storage is treated as a special case of an output or

input operation, allowing data accesses to continue to be controlled by any license restrictions.

3) Peripheral Equipment: Whenever data leave the system (i.e., exit the tamper-detecting case) they must be protected. In some cases an electronic license may allow plaintext data to be output, however we will describe the situation where protection is required. Data are encrypted prior to being sent to an output device.

If the device is an intermediary (e.g., storage device or transmission device) then data must be encrypted in the same format as an input file. Then, whenever the data are read, an electronic license will be required. The encryption key used may be that of one of the data providers or may be a new key created by the system that wrote the data. The user of these data would acquire a license from one of the original data providers or from the value-added providers responsible for the final product. The decision governing which key is used is a function of the licenses governing all input data. Typically, this would be negotiated by the parties and implemented within the electronic licenses.

If the device is a transducer producing a physical product, such as a printer, display, or audio system, the data must be protected up to the time they are converted into the end product. This requires encryption for transmission as well as protection within the peripheral.

Two schemes come immediately to mind. In the first method, modeled after the basic concept described above, the data and license are packaged together. The peripheral would implement a subset of the access mechanism sufficient for carrying out its functions. In the second method, the peripheral and PC would establish a secure link (based upon a certificate within the peripheral) and a session-encrypted output stream would be sent to the peripheral. This is the type of approach that has recently been proposed to the Copy Protection Technical Working Group for allowing transfer of DVD movies among authorized consumer electronics devices.

4) Performance: Protection is not free. Some overhead is associated with the operations necessary to maintain a secure environment. One of the hallmarks of this approach is that the overhead is quite low.

When the system is operating with unprotected data there is no perceivable overhead. A few instructions check for the presence of a header when a file is opened, or check the status of the system-wide flag. If protection is not in use, special processing is not required.

When a header indicates that a file contains protected data, the associated electronic license must be accessed. This entails an additional disk access and a public-key decryption. Nevertheless, as this occurs only once for each file, this is not a source of significant overhead.

Reading or writing protected data requires encryption and decryption, respectively. If implemented in software, this would add overhead. Implementers will likely use a commodity chip to perform encryption and decryption on the fly, as data are being input or output, resulting in little overhead.

E. PKI

The need to encrypt each license with the public key of the system that will be reading the data results in an associated requirement to verify the authenticity of the public key. An authentic public key is one that is associated with the private key kept securely within the system. If the private key corresponding to a public key is known, that would enable the owner of the public key to read the license outside the PC, to obtain the secret key, and to decrypt the protected file. To prevent this from occurring, the system generates public/private key-pairs within the PC and registers the public key with a validation authority. No identification of the user or of the PC is required, although manufacturers may choose to associate a machine with the public key.

The public/private keys may be generated by the manufacturer and registered prior to a machine's sale, or they may be generated after the sale when the customer first uses the machine (e.g., based on a stream of erratically timed keystrokes provided by the customer) and transmitted to a validation authority (encrypted and signed by the PC).

Each time a license is requested, the data owner must validate the public key. This step is similar to obtaining a credit card authorization, although simpler. The data owner merely needs to supply the public key (or its fingerprint, or a pseudonym) and the validation authority responds by indicating the validity of the public key (a 1-bit response) or sends the data owner the public key to use.

F. Required Characteristics of the Global System

A system that provides persistent access control must ensure both that data may be delivered securely to the intended user and that, even after delivery, the data remain secure. The system described here meets each of these requirements by maintaining data in encrypted form except when in use.

In order to ensure a long period of protection, a system must minimize the amount of information that is kept secret. System security depends only on protecting the private key stored in each user computer. It is likely that the private key, along with any active (i.e., decrypted) secret keys will be stored in special hardware. Detection of tampering (i.e., violating the outer case) or an attempt to gain access to the special hardware will result in erasure of the keys. Licensed material held in RAM is also protected. Detection of tampering will result in removal of power and loss of the contents.

Data owners are responsible for maintaining the security of their secret keys. Although a conventional personal computer could be used to package data for secure distribution, it is likely that data owners will require additional measures in order to safeguard their secret keys. The hardware used to safeguard the system's private key can also be used for holding a data owner's key. This would make it impossible to obtain the secret key and use it to release the protected data. Instead, there would be a much smaller risk: unauthorized access to the system containing the secret

key would allow creation of licenses—a point-vulnerability rather than a system wide compromise.

Finally, with tens of millions of personal computers and tens of billions of dollars of software investment (programs and data), it is essential that a system to protect data be compatible with current files, current equipment, and current software. This system provides both forward and backward compatibility.

The presence of a header at the front of a file indicates that the file is protected. Programs using only unprotected files experience no change in processing. Even programs that read protected files are not aware of any change in processing other than the time required to obtain the license and decrypt or encrypt input and output files. Programs that operate on today's PC's will continue to operate on machines that implement the protection scheme described here.

Protected files may be processed on systems that do not implement the access control system. Although such files cannot be decrypted, they can be copied, stored, and transmitted. During a transition period, as the system is being put into place, intermixed machines can communicate with one another although protected data would be available only on machines with the persistent access control capability.

IV. Example Applications

A. Software Distribution/Feature Control/Versions

A software vendor could provide a single version of a popular software program, which implemented multiple features and levels of capability. A user would purchase only the levels that they desired. If the user wished to upgrade to a higher level of capability, they could do so immediately by obtaining an upgraded license. There would be no need to download a new version of the program. As the license size is on the order of a thousand bytes, this process would be convenient and very rapid.

Vendors would not have to develop and maintain separate versions of software. A single release would contain all features, controlled by the licenses issued to users. Not only would development and logistics be simplified, but also there would be no issues of different behavior or bugs in different versions, and the load on the vendors distribution facility would be significantly lower.

Software upgrades could be accommodated within the license structure so that a user would automatically be eligible for upgraded software (for some period of time or number of upgrades). The vendor need only place the upgrade on a server and announce its availability.

B. Key Escrow—A Special Case of Feature Control

Some encryption products are restricted from export out of the United States unless their use is restricted to those times when the encryption key has been deposited with an approved escrow authority. The access control mechanism can enforce a requirement that a certificate (from an approved authority, attesting to the escrow of the encryption key) be present and can ensure that the associated key is used for encryption. Not only would such a program be encrypted when distributed, but it would also be protected against modification when running. Without such protection, such programs can be reverse-engineered and modified to defeat any protections provided by their developers.

C. Digital Entertainment Products

The paucity of digital video entertainment products currently available that provide digital output is understandable in light of Hollywood's fears that digital content released today, when adequate protection is not available, would soon be posted on the Internet, as is being done with MP3 files of digital sound recordings (notwithstanding that doing so would be in violation of copyright law). Some proprietary systems have been announced that begin to address this situation. However, they do so either by restricting the data to analog form when delivered to the user (thus sacrificing some of the advantages of digital source materials) or by restricting the ability of a PC to have general access to data. The system described here allows the full interconnection of PC's and entertainment appliances. It would open up new markets, for example, by enabling the purchase of an electronic license that allows the printing of single frame posters from movies. Making the power of the PC available for entertainment products while protecting the interests of the product owners will lead to many new opportunities.

The system described above allow the distribution of digital entertainment products on, e.g., DVD's or web servers. The consumer would purchase a license providing the access rights desired. This could be accomplished at a retail store (the store would prepare the license or would obtain a license directly from the content provider) at the same time as the DVD is picked up. Alternatively, the consumer might obtain the product from a web server and purchase the license on the Internet or by telephone or e-mail. The license could be delivered through a browser (FTP), by e-mail, or by parcel delivery.

In the case of a rental, no return would be necessary—a significant advantage to the consumer. A secondary advantage (both to the consumer and to the content provider) is the ready availability of the product for use at a later date without requiring a trip to the rental store. The user would merely purchase another license. Even simpler, vendors might offer prepaid licenses that could be debited as various programs were viewed.

D. Program of the Week

In a manner similar to the operation of current book clubs, a future audio/video club could send subscribers the "program of the week." Transmission would be done at night, when telephone line or cable usage is low and bandwidth is inexpensive. The club would allow a free sample to be heard/viewed by the subscriber. If the subscriber wished to access the entire program, he or she would then purchase a license or use a prepaid debit license as above.

This approach enjoys the cost savings associated with avoiding the delivery of a physical product—no production costs, no delivery costs, and no return costs. Further, subscribers might save the program for a future date (as the cost would only be storage media) or might pass the program along to a friend or colleague. The customer becomes part of the distribution mechanism.

E. V Chip

The functions required of the V Chip can easily be made a part of electronic licenses that govern access to digital television or to digital movies.

F. Enterprise Access Control

An enterprise can use this system to control access to information within its domain or among its customers and partners. Doing so would take advantage of the services provided by the access control mechanism and would reduce the complexity of many current systems. A common interface would be used throughout the enterprise, replacing the *ad hoc* arrangements in use today.

Electronic licenses could be distributed either by name or by role. In either case the recipient would receive only those accesses granted by the data owner. The data owner would control who and when a user could access data as well as the user's ability to copy, to edit, or to redistribute the data.

G. Private Networks

Using software similar to that employed in current e-mail systems, it would be straightforward to implement a virtual private network (over the Internet) within which each message would be encrypted and accessible only by authorized recipients.

For especially sensitive information or applications, a private network based on routers and gateways using persistent access control technology would allow both message content as well as addresses and headers to be encrypted. This would make it virtually impossible within the network to eavesdrop on or to alter a specific message or messages from a specific individual.

H. Encyclopedias, Collections, and Databases

Access to information might be marketed on a per-use basis rather than as an entire collection. This might result in a low selling cost for an encyclopedia, coupled with a cost for each access. Market elasticity in this situation would result in an increase in sales of the encyclopedia and a subsequent increase in use. Many families that would not purchase a full price encyclopedia, either because it is too expensive or because they feel they would be only infrequent users, would purchase a lower cost item. Consumers would view the marginal cost of each access as within reach. Volume discounts would provide incentives for those who are heavy users of the information.

A publisher might provide a "library rate" or an "educational rate" for works that would be used in specific settings. Indeed, market forces would be expected to exert pressure on publishers to provide such rates in lieu of the no-longer-applicable "fair use" rights pertaining to copyright protection.

I. Copyright Clearance

Publishers of journals and other reference materials would provide prepaid licenses that would be debited on a per-copy basis. This would enable corporate and other institutional users to obtain electronic copies conveniently, while ensuring collection of revenue associated with each copy.

J. Derivative Products

Allowing an end user to create a derivative product, the sales of which generate additional revenue for the original content owner, will create new markets and opportunities. The movie-poster shop (see above) is just one such example.

V. Conclusion

We have presented a global approach to protecting digital intellectual property that allows the user access to material only in accord with a license obtained from the property owner. The material can be distributed openly in protected form but can only be viewed or used within a system that processes the owner's license restrictions and protects the data. The implementation of the system can be effected within a PC or consumer electronics appliance. The resulting device, in addition to its new capabilities, remains compatible with existing data and retains complete functionality. One may compromise between the security level and the demanded modifications (secure hardware, secure OS). That kind of system will encourage the growth of electronic commerce solutions for the sale and delivery of licenses. Additionally, the ability to disseminate content and yet provide the owner with continuing control and revenue opportunities will enable many new markets and applications.

Acknowledgment

The author gratefully acknowledges the careful review and improvements to the paper made by the Guest Editor of this Special Issue and his referees.

References

[1] B. Schneier, *Applied Cryptography*, 2nd ed. New York: Wiley, 1996.
[2] D. Stinson, *Cryptography Theory and Practice*. Boca Raton, FL: CRC Press, 1995.
[3] T. J. Alexandre, "A Java-based platform for intellectual property protection on the World Wide Web," in *Proc. 7th World Wide Web, (Special Issue of Computer Networks and ISDN Systems)*, vol. 30, no. 1-7, paper SP1.
[4] *Signal Processing (Special Issue on Copyright Protection and Access Control for Multimedia Services)*, vol. 66, no. 3, 1998.
[5] L. Boney, A. H. Tewfik, and K. H. Hamdy, "Digital watermarks for audio signals," in *Proc. EUSIPCO 1996*, Trieste, Italy, Sept. 1996.

[6] S. Low, N. Maxemchuk, J. Brassil, and L. O'Gorman, "Document marking and identification using both line and word shifting," in *Proc. Infocom'95*, Boston, MA, Apr. 1995.
[7] J. P. Linnartz, "The ticket concept for copy control based on embedded signalling," in *Proc. ESORICS'98 (Lecture Notes in Computer Sciences)*, vol. 1485. Berlin, Germany: Springer-Verlag, 1998, pp. 257–274.
[8] B. Macq and J. J. Quisquater, "Cryptology for digital TV broadcasting," *Proc. IEEE*, vol. 83, pp. 944–957, June 1995.
[9] P. Devanbu and S. Stubblebine, "Research directions for automated software verification using trusted hardware," in *Proc. IEEE Int. Conf. Automated Software Engineering*, 1997, pp. 274–279.
[10] U. Linqvist, T. Olovsson, and E. Jonsson, "An analysis of a secure system based on trusted components," in *Proc. 11th Annu. Conf. Computer Assurance, Systems Integrity, Software Safety and Process Security, COMPASS'96*, pp. 213–223.
[11] B. Yee and J.-D. Tygar, "Secure coprocessor in electronic commerce applications," in *Proc. 1st USENIX Workshop Electronic Commerce*, New York, 1995.
[12] R. S. Sandhu and P. Samarati, "Access control: Principles and practice," *IEEE Commun. Mag.*, pp. 40–48, Sept. 1994.
[13] P. A. Loscocco, S. D. Slalley, P. A. Muckelbauer, R. C. Taylor, S. J. Turner, and J. F. Farrell, "The inevitability of failure: The flawed assumptions of security in modern computing environments," in *Proc. 21st National Information Systems Security Conf.*, Oct. 1998.
[14] C. P. Pfleeger, *Security in Computing*, 2nd ed. Englewood Cliffs, NJ: Prentice-Hall, 1997.
[15] N. Rankl and W. Effing, *Smart Card Handbook*. New York: Wiley, 1997.
[16] F. C. Mintzer, L. E. Boyle, A. N. Cazes, B. S. Christian, S. C. Cox, F. P. Giordano, H. M. Gladney, J. C. Lee, M. L. Kelmanson, A. C. Lirani, K. A. Magerlein, A. M. B. Pavani, and F. Schiattarella, "Toward on-line worldwide access to Vatican Library materials," *IBM J. Res. Develop.*, vol. 40, no. 2, pp. 139–162, Mar. 1996.

Paul B. Schneck (Fellow, IEEE) received the B.S. and M.S. degrees from Columbia University, New York, and the Ph.D. degree from the Courant Institute of Mathematical Sciences, New York University, New York.

He has been the Assistant Director for Computer and Information Sciences at Goddard Space Flight Center, Head of the Information Sciences Division at the Office of Naval Research, and a Fellow and Director of Information Systems at Mitretek Systems (and its parent, The MITRE Corp.), where he was responsible for the Technology Research Program and the Corporate Intellectual Property Group. From 1985 to 1993, he was founding Director of the Supercomputing Research Center, Institute for Defense Analyses. He currently is Vice President, Information Analysis Solutions, MRJ Technology Solutions, Inc., Fairfax, VA. He has taught at Columbia University, the University of Maryland, and Johns Hopkins University. He is author of the book *Supercomputer Architecture* (Norwell, MA: Kluwer, 1987) as well as several book chapters and numerous technical papers on hardware, software, algorithms, and numerical analysis for high-performance computing. He has an allowed patent pending for persistent access control.

Dr. Schneck is a Fellow of the Association for Computing Machinery, Inc. and a member of the British Computer Society. He was Chairman of the IEEE Technical Committee on Computer Architecture, and he is a member of the Editorial Board of this PROCEEDINGS.