Intellectual property protection systems and digital watermarking

Jack Lacy, Schuyler R. Quackenbush, Amy R. Reibman and James H. Snyder

AT&T Labs – Research, Florham Park, NJ; Red Bank, NJ

lacy@research.att.com; srq@research.att.com; amy@research.att.com; jhs@research.att.com

Abstract: Adequate protection of digital copies of multimedia content - both audio and video - is a prerequisite to the distribution of this content over networks. Until recently digital audio and video content has been protected by its size: it is difficult to distribute and store without compression. Modern compression algorithms allow substantial bitrate reduction while maintaining high-fidelity reproduction. If distribution of these algorithms is controlled, cleartext uncompressed content is still protected by its size. However, once the compression algorithms are generally available cleartext content becomes extremely vulnerable to piracy. In this paper we explore the implications of this vulnerability and discuss the use of compression and watermarking in the control of piracy.

OCIS codes: (100.2000) Digital Image Processing, (999.9999) Steganography

References and links

1. M. Bosi, K. Brandenburg, S. Quackenbush, L. Fielder, K. Akagiri, H. Fuchs, M. Dietz, J. Herre, G. Davidson, Y. Oikawa, "ISO/IEC MPEG-2 Advanced Audio Coding," presented at the 101st Convention of the Audio Engineering Society, preprint 4382, Nov. 1996.
2. D. Aucsmith, "Tamper Resistant Software," in *Proceedings of the First International Information Hiding Workshop*, LNCS 1174, Springer-Verlag, Cambridge, U.K., 317-334, May/June, (1996).
3. M. Blaze, J. Feigenbaum, J. Lacy, "Decentralized Trust Management," in *Proceedings of the 1996 IEEE Symposium on Security and Privacy*, 164-173 (1996).
4. D. Boneh, J. Shaw, "Collusion-secure Fingerprinting for Digital Data," Crypto '95, LNCS 963, Springer-Verlag, Berlin, pp. 452-465 (1995).
5. I. J. Cox and J.M.G. Linnartz, "Public Watermarks and Resistance to Tampering," Proceedings of the Fourth International Conference on Image Processing, Santa Barbara CA, October (1997).
6. F. Hartung and B. Girod, "Digital Watermarking of MPEG-2 Coded Video in the Bitstream Domain," *Proc. IEEE ICASSP*, 2621-2624, April (1997).
7. "Cryptolope Container Technology," an IBM White Paper, http://www.cryptolope.ibm.com/white.htm.
8. International Federation of Phonograph Industries, Request for Proposals.
9. *Proc. First International Information Hiding Workshop*, LNCS 1174, Springer-Verlag, Cambridge, U.K., 207-226, May/June, (1996).
10. J. Lacy, D. P. Maher, and J. H. Snyder, "Music on the Internet and the Intellectual Property Protection Problem," *Proc. International Symposium on Industrial Electronics*, Guimaraes, Portugal, July (1997).
11. J. Lacy, S. Quackenbush, A. R. Reibman, D. Shur, and J. H. Snyder, "On combining watermarking with perceptual coding," Int. Conf. Acoustics, Speech, and Sig. Proc., May (1998).
12. O. Sibert, D. Bernstein, D. Van Wie, "Securing the Content, Not the Wire, for Information Commerce," http://www.intertrust.com/architecture/stc.html.
13. J. Smith, B. Comisky, "Modulation and Information Hiding in Images," *Proc. First International Information Hiding Workshop*, LNCS 1174, Springer-Verlag, Cambridge, U.K., 207-226, May/June, (1996).
14. http://www.research.att.com/~srq/OpticsExpress

1. Introduction

Protection of digital copies of multimedia content — both audio and video — is a prerequisite to the distribution of this content over networks. Until recently digital audio and video content has been protected by its size. For example, audio on compact discs is encoded

#7085 - $15.00 US Received October 29, 1998; Revised December 04, 1998

using PCM at 1.4 megabits per second — about half a gigabyte for a 45 minute CD. Such large quantities of data are difficult to distribute and store. Modern compression algorithms provide high-fidelity reconstruction while allowing substantial size reductions. If distribution of these algorithms is controlled, cleartext, uncompressed content is still protected by its size. However, once the compression algorithms are generally available cleartext content becomes extremely vulnerable, as is evidenced by the proliferation of illegally distributed MP3 compressed music. In this paper we explore the implications of this vulnerability and how watermarking techniques can contribute to a system strategy that protects intellectual property.

2. A systemic view of IP protection

The design of secure systems should be based upon an analysis of the application risks and threats. As Figure 1 illustrates, such analysis will identify some of the risks of a particular domain. The technological net should handle many identified risks. The legal net will handle others. No matter how thorough the analysis, not all risks will be identified, and not all identified risks will be caught by the technological and legal nets. Ideally the system design includes the possibility of renewable security so that these residual risks do not undermine the foundations of the business.

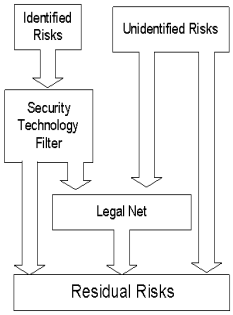

Figure 1. IP protection system

The business model for the application is one of the strongest security mechanisms. If the system is easy to use, rich in features, support and information, and reasonably priced, why should consumers go to the black market? Designing the system with this in mind will minimize the attacks from legitimate users, most of whom are willing to play by the rules. System security should not interfere with legitimate use. Finally, we want to design the system so that even if an attacker does break the system he cannot then use the same system to distribute that IP for his gain. Such a system should consist of:

1) a compression engine for managing music or video. This mechanism should discourage multiple compression/decompression cycles;
2) a mechanism for protecting the integrity of the content and for enforcing rights-to-use rules;
3) a flexible mechanism for licensing content and for granting various rights to consumers with appropriate credentials;
4) a secure client for accessing, rendering, playing or viewing content in a manner consistent with system policy and with the credentials or licenses associated with that content;

#7085 - $15.00 US Received October 29, 1998; Revised December 04, 1998

5) a mechanism for labeling the content to be distributed in a persistent manner. For example, the label might indicate ownership, name the distributor, identify the property or contain information about transactions involving the content.

Component 2 involves the use of cryptographic containers as in [10], [12] and [7]. The content is encrypted and perhaps digitally signed. The encryption keys are distributed via other channels using cryptographic protocols. A flexible licensing mechanism (Component 3), based for example upon PolicyMaker [3], manages these keys and governs their use [10]. Client security (Component 4) is what distinguishes the IP protection problem from the protected communications channel problem. That is, content must be protected in the client, not just in the channel. Protection mechanisms include tamper resistant software and hardware. These techniques are discussed in [10] and [2].

2.1 Compression

As discussed earlier, compression enables the distribution of music or video over networks. For audio, the MPEG-2 Advanced Audio Coder [1] provides CD quality reproduction for most music and most listeners at a compression ratio of 11 to 1 (128 kilobits per second). Compression may also be relevant as a protection mechanism for the following two reasons.

Attackers will always have access to decompressed output. If recompression of the decompressed content results in noticeable degradation of quality, then the 2nd generation output will be of sufficiently low quality that it is not a threat to the IP owner. Of equal importance, when cleartext content is available, nothing we do to protect compressed content matters. Should controlled degradation via compression prove possible, then a solution to this cleartext audio problem would be to compress and then decompress the music as part of the mastering process. Controlled degradation via compression is an area of current research.

Because it can easily be distributed, the compressed file is the valuable commodity. It therefore makes sense to associate labels with the compressed file in a way that is persistent in the compressed domain.

3. Digital watermarking

3.1 Overview

As stated earlier, a mechanism is needed for binding content identification to content in a persistent manner. Digital watermarking is such a mechanism. (See for example [9].) Watermarking has also been proposed as a mechanism for gating the use of content. In this case, when decisions regarding access to or use of the content are made, the mark must be retrieved in real-time and used as input in the decision-making process. No one marking algorithm is best suited for these two functions, both because of complexity issues and because different functions and different marking algorithms are resistant to different attacks. Indeed, we expect that any single album or film will be marked by a variety of different algorithms, to improve the overall resistance to attack.

System designers should think carefully before using watermarks to gate usage, since by feeding different bitstreams into the gating mechanism the attacker may be able to probe the watermark algorithm, discover mark sites and possibly generate fraudulent marks [5]. If a marking algorithm is to be used to gate usage, the algorithm should be designed in such a way that tampering with the mark should degrade the quality of the decompressed content. This suggests that the marking algorithm could beneficially be associated with the compression algorithm. We describe one such marking algorithm in section 3.4.

3.2 Desirable characteristics of watermark algorithms

The following requirements are typically expected of watermarks [8]:

#7085 - $15.00 US Received October 29, 1998; Revised December 04, 1998

1) *Imperceptibility.* A watermarked should not be distinguishable from the original signal.
2) *Information capacity*. The mark bitrate must be compatible with the rate limits imposed by the system.
3) *Robustness.* The mark must be recoverable, not only in the complete work, but also in truncated, filtered, dilated, and otherwise processed clips, in a concatenation of unrelated content, and in the presence of noise.
4) *Low complexity*. Marking schemes intended for use with real-time applications should be low complexity.
5) *Survive multiple encode-decode generations.* A watermark should survive tandem encoding-decoding.
6) *Tamper resistant or tamper evident.* It should be possible to recognize that a mark has been modified. It should not be possible to modify a mark in such a way as to create a different valid mark.
7) *Difficult to create or extract legitimate watermark without proper credentials.* In the context of the watermarking engine alone, a proper credential is knowledge of the algorithm used to insert the mark. An ideal would be a public key analogue to watermarking: hard to insert mark, easy to retrieve, hard to counterfeit.

For copyright identification every copy of the content can be marked identically, so the watermark can be inserted once prior to distribution. Ideally, detection should not require a reference because a search engine has no *a priori* way to associate the reference material with the work from which the mark is to be recovered. Not only must the watermark be short enough to be recovered in a truncated version, some means must be provided to synchronize the detection process so that the watermark can be located in the processed bitstream. Finally, any attempt to obscure the mark, including re-encoding the content, should lead to perceptible distortion.

Transaction identification requires a distinct mark for each transaction. The primary challenge of point-of-sale marking ("fingerprinting") is to move the content through the marking engine quickly. That is, the algorithm must be low complexity. One strategy is to insert the watermark in the compressed domain, in which case mark insertion should increase the data rate very little. Watermarking algorithms designed for fingerprinting must be robust to collusion attacks.

3.3 General mechanisms

Watermarks for compressed content fall into three categories: cleartext or original (PCM in the case of audio or video) marking, compressed bitstream marking which does not alter the bitstream semantics, and marking integrated with the compression algorithm in which the semantics of the bitstream are altered. We describe these below and discuss their advantages and limitations. We anticipate that in a well-designed system, each of these marking techniques will be used.

Cleartext PCM*:* We define cleartext watermarks as marks inserted in the original or during decompression into output (e.g. while writing a decompressed song to CD). Cleartext marking embeds a data stream imperceptibly in a signal. The model for many cleartext-marking algorithms is one in which a signal is injected into a noisy communication channel, where the audio/video signal is the interfering noise [13]. Because the channel is so noisy, and the mark signal must be imperceptible, the maximum bit rates that are achieved for audio are generally less than 100bps.

Cleartext marks are intended to survive in all processed generations of the work. They are therefore well suited to identification of the work. There are two major concerns with cleartext marking. Because such algorithms (usually) compute a perceptual model, they

#7085 - $15.00 US Received October 29, 1998; Revised December 04, 1998

tend to be too complex for point-of-sale applications. Second, these algorithms are susceptible to advances in the perceptual compression algorithms.

Retrieval mechanisms for cleartext watermarks fall into two classes: reference necessary and reference unnecessary. In either case the mechanism for mark recovery is generally of high complexity and is often proprietary. Further, if means for detecting these watermarks are embedded in a player, an attacker, by reverse engineering the player, may be able to identify and remove the marks. We feel that cleartext watermarks should *not* be used to gate access to content.

Bitstream Watermarking (semantic-non-altering): Bitstream marking algorithms manipulate the compressed digital bitstream without changing the semantics of the audio or video stream. Bitstream marking, being low-complexity, can be used to carry transaction information. Because the mark signal is unrelated to the media signal, the bit rate these techniques can support can be as high as the channel rate. However these marks cannot survive D/A conversion and are generally not very robust against attack; e.g. they are susceptible to collusion attacks. This type of mark can easily be extracted by clients and is thus appropriate for gating access to content; it is an example of a security measure intended primarily to "keep honest users honest".

Bitstream Marking Integrated with Compression Algorithm (semantic altering): Integrating the marking algorithm with the compression algorithm avoids an 'arms race' between marking and compression algorithms, in which improvements in hiding data imperceptibly in content are undercut by and even motivate further improvements in perceptual compression algorithms. Since the perceptual model is available from the workings of the compression algorithm, the complexity associated with marking can be minimized. Integrated marking algorithms alter the semantics of the audio or video bitstream, thereby increasing resistance to collusion attacks. An example of this approach is [6], which however does not use perceptual techniques. We now present another example.

3.4 Integrating the watermarking algorithm with compression

We have developed a first generation system that combines bitstream and integrated watermarking. It can be configured to support the three marking functions mentioned above. It does not include but is compatible with use of a front-end cleartext-marking algorithm as well. We assume that the cleartext original is not available except possibly to auditors seeking to recover the watermark. In particular, the cleartext original is not available to attackers. The decompressed and marked content will generally be available to everyone.

Our method relies on the fact that quantization, which takes place in the encoder, is a lossy process. By combining mark insertion with quantization we ensure that the attacker cannot modify the mark without introducing perceptible artifacts. The fact that marking data is present is indicated by characteristics of the bitstream data. Our marking technique involves the perceptual modeling, rate control, quantization, and noiseless coding blocks of a generic perceptual coder. The algorithm can be used for either audio or video. We concentrate on audio here. Results for video can be found in [11].

In MPEG AAC spectral lines are grouped into 49 "scale factor" bands (SFB), each band containing between 4 and 32 lines. Associated with each band is a single scale factor, which sets the quantizer step-size, and a single Huffman table (AAC employs 11 non-trivial Huffman tables). The coefficient for each spectral line is represented by an integer (i.e. quantized) value.

Let $A = \{f_i, H_i, \{q_{ij}\}\}$ be the set of triples of scale factors f_i, Huffman tables H_i, and quantized coefficients $\{q_{ij}\}$. We assume that we have selected some set of scale factor bands into which mark data will be inserted. The marking set will generally be dynamic. Let M be the set of indices associated with the set of SFB chosen for marking.

#7085 - $15.00 US Received October 29, 1998; Revised December 04, 1998

DOCKET ALARM

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS
Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS
Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS
Sync your system to PACER to automate legal marketing.