

© 2004 IEEE. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or lists, or to reuse any copyrighted component of this work in other works must be obtained from the IEEE.

Watermarking to Track Motion Picture Theft

Jeffrey A Bloom and Christos Polyzois

Sarnoff Corporation
Princeton, New Jersey, USA

Abstract - The music industry suspects that unauthorized P2P trading of music files has had a negative impact on revenue. The motion picture industry fears the same thing will happen to movies. However, the piracy problems in these two domains are different. In this paper we take an interesting look at movie piracy as we compare it to music piracy in the two areas of the piracy source and the potential impact on revenue. We do not directly address issues of politics, business, ethics, sociology, or copyright law, but identify relatively uncontroversial areas where technology, specifically digital watermarking, can make a significant contribution.

I. INTRODUCTION

There are important differences in piracy of music and movies that stem from fundamental differences in the way in which these products are brought to market. Music typically has one release date, the date on which the recording is made available for purchase. Release for radio broadcast typically coincides with, or slightly precedes, the release date and, around the same time, the performer may begin a tour of live performances to promote the product. In contrast, motion pictures have many release dates or windows. A key date is the date of theatrical release and the piracy problem can be considered as two: pre-release piracy and post-release piracy. Section II below discusses the sources of piracy of both music and movies drawing a distinction between piracy due to consumers and piracy due to *insiders*. These two sources of piracy also have different potential impact on the revenues of the content owner and this is discussed in Section III. The conclusion drawn from these two sections is that an important segment of motion picture piracy is due to insiders prior to theatrical release. The remainder of this paper discusses the use of digital watermarking to address this problem and, more specifically, the use of low frequency watermarking.

II. SOURCES OF PIRACY

The rise of the peer-to-peer network has led to file sharing, and one of the most commonly shared file types is digital music. The music files that are commonly traded fall into two general categories: bootlegs and rips. Bootlegs are copies of recordings made at live performances. When these recordings are unauthorized (the typical case), such trading constitutes a form of piracy. Rips are music files copied off of a store bought recording such as a compact disc.

The music industry suspects that unauthorized trading of music files on P2P networks has had a negative impact on revenue. The motion picture industry fears that the same thing will happen to movies as compression technologies improve and bandwidth to the home increases.

The sources of motion picture content traded on P2P networks are different from the sources of traded music content. Different sources imply differences in quality, potential impact on revenues, legal and business implications, and technologies that can be adopted to address the issues.

A. Sources of Pirated Music

Once music files appear on a P2P network, they are distributed by the members of the network. However, the initial piracy, the unauthorized bootleg or rip, is performed by a legitimate customer. Concertgoers who have purchased tickets surreptitiously record a concert with a hidden microphone and then make copies of that recording available for trade. Other customers purchase a CD, copy music files from the disc, and make these files available for trade. The music industry has, to date, targeted the P2P traders in their attempts to suppress the unauthorized distribution of music files.

The quality of a bootleg can vary greatly from one recording to the next. However, the quality of a rip is essentially identical to that of the CD source. Ripped songs are often compressed prior to redistribution, however current compression technologies such as MPEG Layer 3 (mp3), provide very good quality.

B. Sources of Pirated Movies

There are many sources of pirated motion pictures. One distinction that can be made is between piracy of movies prior to the theatrical release date and post-release piracy. Pre-release sources of piracy include production and post-production, pre-release distribution, and pre-release screenings. After theatrical release, motion pictures can be pirated from the theater, during controlled *small-screen* releases, from DVD or video releases, Internet distribution, or from broadcast television.

During production, the film passes through many hands, both internal to the editing room and outsourced to third parties, before the final release version is ready. This process provides many opportunities for piracy. When such a pre-release copy is pirated, the culprit is a trusted insider: an employee of the studio or one of its post houses. The quality can be very good. However, the content is often an unfinished product.

Once the final release version is complete, the movie may be distributed on VHS or DVD to a select (but large) group of individuals including members of the Academy of Motion Picture Arts and Sciences and critics as well as marketing and publicity professionals. These copies are often referred to as

screeners. Any of these trusted parties could be tempted to make their copy available for wider, unauthorized distribution, again prior to theatrical release. The quality of screeners can be as good as the DVD source. However, screener copies typically have a text overlay identifying the movie as a pre-release copy that is not intended for distribution.

Still prior to theatrical release, the film may be shown in private screenings to critics, sponsors, and VIPs. There are two potential sources of piracy during this time. Members of the audience, who have been invited as the guests of the studio, may capture the movie on camcorder and make the bootleg available for unauthorized distribution. Such copies are typically poor quality with poor sound. However, the theater operator, again working as an agent of the studio, can surreptitiously copy the movie at very high quality. This can be done simply, by showing the movie after hours with an empty house and a camcorder setup on a tripod in the back of the theater. The sound can be taken directly from the house sound system. Or, an even higher quality copy can be made directly from the film through a digital scanning process known as telecine.

Once the film is released in the theaters, the potential for camcorder capture and theater operator piracy increases drastically. This time the camcorder pirates are paying customers and this situation is analogous to the bootleg piracy of music. The camcorder copies are typically poor quality. The theater operators, who can be considered semi-trusted industry insiders, can get high quality camcorder captures or very high quality telecine copies.

Shortly after theatrical release, the film is released in a number of controlled *small-screen* settings. These include airline release (IFE), pay-per view (PPV), and video on demand (VOD). The IFE copies are entrusted to airline employees who have the opportunity to make unauthorized duplicates. PPV and VOD can be copied by consumers. The quality of these copies is DVD or VHS quality, however, the IFE copies typically have a visible text overlay identifying the airline to which the copy was distributed.

Once the movie is released on VHS and DVD, it becomes vulnerable to ripping. Analogous to the ripping of songs from a music CD, a legitimate customer could copy the movie file from a purchased or rented DVD and make it available for unauthorized distribution. The quality of these rips can be as good as the original.

Beyond DVD release, the movie can be captured during an Internet distribution window or, later, during broadcast. We won't discuss this as a significant source of piracy because, as will be discussed, most of the damage that piracy can inflict occurs prior to this time.

From this discussion, we see that the source of all pre-release piracy is trusted insiders who have either stolen property from their employer or client or who have betrayed the trust of the movie studio. In a recent study, it has been estimated that approximately 77% of all movies available on a

P2P network are the result of insider piracy and that 95% of all available movies first appeared on the P2P network before the theatrical release date [1][2].

III. POTENTIAL IMPACT OF P2P PIRACY

A true assessment of the impact of P2P trading on revenues is well beyond the scope of this paper and would likely be fraught with controversy. For the purposes of this discussion we intend to suggest that post-release motion picture piracy is analogous to music piracy and that pre-release piracy has the potential to be more damaging than post-release piracy.

Looking at the potential impact of P2P piracy on the music industry, we make the generous assumption that all illegal downloads represent lost sales. In other words, we assume that any consumer who downloads a song will choose not to purchase the CD, but would have purchased it otherwise. This is known to be a high upper bound as many downloaders would not have purchased the CD anyway and many other downloaders choose to purchase the CD even after (or perhaps because of) downloading. The actual losses due to P2P trading are still a topic of disagreement.

We can make a similar upper bound approximation for post-release piracy of movies. We generously allow that any consumer who would otherwise have seen the movie in a cinema, purchased the movie as a PPV or VOD, and/or purchased or rented the DVD or video, will choose not to do so after downloading the movie.

Pre-release piracy has the potential for more significant impact on revenues. Again, we assume that a consumer who obtains a pre-release copy of the film will refrain from going to the theater. Most obviously, the earlier a pirated copy is available, the more release windows it can potentially impact. Pirated pre-release copies can impact every release window.

Some percentage of the folks who obtain a pirated copy of a film prior to the release date would have attended an opening weekend exhibition and this group can have a more significant impact on movie revenue. This is because opening weekend box office sales represent an important variable upon which future movie revenues are based. Many downstream revenues are preset as a function of opening weekend box office receipts: DVD and VHS duplication and distribution contracts, IFE, PPV, and VOD contracts. Thus, reduced sales on opening weekend can result in reduced revenue in all subsequent release windows.

IV. ADDRESSING PRE-RELEASE PIRACY

We conclude from the last two sections, that pre-release piracy can have a more significant impact on studio revenue, that the quality of pre-release copies can be very good, and that the source of these copies is primarily trusted insiders.

There may be debate regarding fair use and other copyright issues surrounding the trading of music or movie files on the Internet. All of the proponents of such trading

assume that the person offering the file for trade has obtained the content in a lawful way. Few would argue that stolen property could be legally traded. Thus, while the debate over legitimate customers' rights to trade continues, there should be no such debate when it comes to pre-release piracy.

In the case of production versions of the film, the pirate is often an employee of the production company itself. This employee is stealing intellectual property from the employer. Other times the pirate is an employee of a post-production facility that provides services to the content owner. This pirate is stealing the intellectual property of his clients.

In the case of screeners, the pirate is a trusted critic, Academy member, or sponsor who has agreed to insure that the received copy is not distributed. At the minimum, this pirate has violated that agreement.

There are two pirates in the case of advanced screenings. The theater operator who makes an unauthorized copy of the film is clearly stealing intellectual property and is violating the contract between the content owner and the theater owner. The invited guest who captures the film on camcorder is, at the least, violating the implicit or explicit conditions of the invitation.

In all but the last case, the pirate could be directly identified if tracking information could be associated with each individual copy. A tracking number could be added before each stage of production and into each copy that was sent out to third party service providers for processing. Screener copies could each be uniquely identified. Advanced screenings could be labeled with information identifying the theater to which it was distributed, the equipment on which it was shown, the date and time of showing, and perhaps information identifying the projectionist.

Such tracking information does not prevent piracy directly, but can be recovered from a pirated copy of a movie, thus revealing the person or organization responsible for the unauthorized release. As a forensic tool, tracking information gives the content owner information to help manage the piracy problem and serves as a deterrent to future piracy.

One technology that has been used for this tracking function is digital watermarking. Below we discuss this technology and the specific requirements of this application.

A. Digital Watermarking

Digital watermarking is the "practice of imperceptibly altering a work to embed a message about that work"[4]. Watermarking is a form of *data hiding* in which the alterations representing the message exploit the perceptual redundancies in the work rather than redundancies in the format. For digital motion imagery, the pixel values themselves are altered, changing the colors of various pixels slightly so that the statistics of resulting work are measurably altered.

By strict definition, the alteration must be imperceptible, both visually and audibly. However, in practice we refer to the *fidelity* of a watermark. Fidelity is the degree to which the

watermarked version of a work is perceptually similar to the original. Different applications require different degrees of fidelity.

Digital watermarks come in two flavors. *Robust* watermarks are designed to be recoverable from distorted versions of the watermarked work and *fragile* watermarks are designed to disappear with even the slightest modification. For tracking applications, we typically consider robust watermarks. Different applications require different degrees of robustness.

B. Tracking Requirements

Fidelity

The fidelity requirements vary for the different stages of tracking. During production, the highest level of fidelity is required. Multiple layers of watermarks will be embedded during production and all will be present in the final release version. The combined effect of these watermarks must be imperceptible. Alternatively, removable watermarks can be used. The development of watermarks that can be removed, even after editing, the addition of digital effects, color and brightness balancing, and other production and post-production processing is a problem currently being considered in the watermarking research community.

Advanced copies such as screeners, have a more modest fidelity requirement. In fact, the current practice is to add visible text overlays to these copies. The cinematographer will probably object to visible watermarks that reduce the chances of receiving an Oscar.

Advance screenings are perhaps more important than regular screenings as the audience contains the critics whose opinions impact the movies success. Watermarking of these copies requires very high fidelity. Again, it is common practice to add visible text overlays to these advance screenings.

Robustness

The robustness requirements for all of these tracking applications are the same. The tracking information must be recoverable from a pirate copy traded on a P2P network. P2P trading implies heavy compression and often resize. A common movie file found on P2P networks is 360×240 compressed with an MPEG 4 variant at a bitrate of 250 kbps.

In addition to resize and compression, camcorder capture introduces a number of distortions to which the watermark must be robust. These include geometric distortions, temporal distortion (change in frame rate), cropping (change in aspect ratio), occlusion (audience heads), as well as significant valumetric distortions (brightness, contrast, low-pass filtering, etc.)

Security

One important property of watermarking technologies is security against three classes of attack. Depending on the application, watermarks need to be secure against

unauthorized embedding, secure against unauthorized removal, and secure against unauthorized detection.

For these tracking applications, the watermark needs highest security against unauthorized embedding. When a forensic analysis identifies a specific person as the source of the pirated movie, the content owner must be able to insure that the recovered tracking information was actually embedded prior to distribution and not a forged watermark, embedded by the pirate. It must not be possible (or at least it must be computationally infeasible) for an adversary to frame someone by embedding a valid watermark.

The watermarks should also have high security against unauthorized removal. Clearly, with the ability to remove the watermark, an adversary could successfully pirate a movie without being caught. Removal of the watermark, if at all possible, should require significant expertise, computation, and manual effort.

Finally, the watermark should have high security against unauthorized detection. While it might not be immediately obvious why this would be required, consider a watermark removal attack called *sensitivity analysis* [5]. For most current watermarking schemes, it has been shown that an adversary with a detector can remove the watermark. Thus, a system that does not have security against unauthorized detection is also vulnerable to unauthorized removal.

V. WATERMARKING TECHNOLOGIES

A number of watermarking technologies have been developed over the past 10 years. Many fall into the category of *spread spectrum* techniques. Some of the general properties of spread spectrum watermarking are discussed below along with an assessment of their ability to meet the tracking requirements. The second category of watermarks discussed below is *low-frequency* watermarks. These have different properties and are better suited to meet the tracking requirements.

A. Spread Spectrum

Spread spectrum watermarking describes the class of watermarking techniques in which the watermark is a low amplitude, wide band signal. This signal may be “shaped” in time, space, and/or frequency based on perceptual model to insure imperceptibility. It is then added to or multiplied by the original work to obtain the watermarked work. Some common examples of spread spectrum watermarks are [6], [7], [8], and [9]. Spread spectrum watermarks can often be applied or approximated in compressed domain.

Most spread spectrum watermarks use correlation-based detection strategies. The most popular of these are linear correlation, normalized correlation, correlation coefficient, phase-only correlation, and Fisher Z statistic among others. See [4] for a detailed study of correlation-based watermarking techniques.

Just as with spread spectrum communications, wide bandwidth allows for very low amplitude watermark signals. This decreases the probability that a watermark will be perceptible. With the application of a good perceptual model, very high fidelity can be obtained.

Spread spectrum watermarks are also well known to be robust to additive white noise. These watermarks have been shown in practice to be robust to VHS recording and a number of standard video processes such as filtering, noise reduction, and color enhancement. As good as it is, spread spectrum watermarking is not robust to extreme low-pass filtering and is insufficient to survive camcorder capture. This robustness can be increased by increasing the amplitude of the watermark, but this is done at the expense of fidelity.

B. Low-Frequency Watermarks

While spread spectrum watermarking has become very popular, it is also accepted that low frequencies are the most likely to survive video processing. Most attempts at achieving increased robustness by incorporating low frequencies in the watermark have failed, as they were unable to achieve an acceptable level of fidelity.

A number of efforts have recognized that very low spatial or spatio-temporal frequencies can survive severe compression and camcorder capture. Three notable efforts are [10], [11] and [12], and [13].

In [10], the authors extract the mean luminance from each frame of the image sequence. This sequence of luminance values is then watermarked using a variant of the spread spectrum watermark described in [7]. Spatial and temporal perceptual models guide the watermark embedding. The watermark is 1D in space and spread spectrum in time. By relying on the mean luminance of each frame, it becomes robust to geometric distortions, however a “flickering” artifact is introduced and this method does not meet the fidelity requirements.

A second method, described in [11] and [12], starts with a spread spectrum signal and applies the processes of cellular automaton with voting rules and low-pass filtering to derive a 2D low-pass watermark pattern that is independent of the original content. Such a signal is extremely difficult to “shape” with a perceptual model and this method cannot meet the fidelity requirements of this tracking application.

In the work described in [13], the authors recognized that in this tracking application the detector can make use of the original work as well as any information available at the time of embedding. Such a detector is called an informed detector and this technique has significant advantages over non-informed or blind detectors (see [14] for a discussion of informed detection.)

In this method, a perceptual model is used to create a content-dependent, low-frequency watermark. The watermark pattern is a 3-D spatio-temporal volume comprised of a collection of local, low-frequency carriers. The perceptual

model and a secret key determine the sizes, shapes, and locations of the individual carriers. These carriers are modulated to reflect the message being embedded and then added to the original work. The detector uses the original work to adjust for geometric, temporal, and histogram distortions (see [15] for registration details [16]) and then uses it to remove the image “noise” from the watermark signal. Finally, the recovered signal is compared, via a correlation technique, to the watermark that was embedded to recover the modulation bits representing the watermark message.

This watermark has been shown to meet all of the fidelity, robustness, and security requirements of the tracking application. The high fidelity is due to the custom, content-dependent formation of the low frequency pattern. The high robustness is due to the fact that the watermark pattern is a narrow-band, low-frequency pattern and low-frequencies tend to survive all processing that a pirate might apply, including camcorder capture, resizing, and severe low bitrate compression. The security arises from the fact that the watermark pattern is dependent on both the original content and the watermark key.

VI. CONCLUSION

The most damaging motion picture piracy occurs when pre-release copies are stolen and distributed prior to the theatrical release of the movie. The quality of these pirate copies is generally quite high and the timing of the release has the most impact on revenue generated over the life of the film. The pirates in these cases are trusted insiders.

Watermarking is a technology that can robustly attach tracking information to each individual copy of a movie. The watermarked copy is perceptually indistinguishable from the original and the tracking information cannot be removed or forged. The tracking information can be recovered from a pirated copy of a movie to identify the source of the leak. This serves as a deterrent to future, pre-release, insider piracy.

Standard spread spectrum watermarking technologies are insufficient as they cannot simultaneously meet the stringent fidelity and robustness requirements of this tracking application. Low-frequency watermarks, however, do have the potential to achieve this goal as has been shown in [13].

- [1] S. Byers, L. Cranor, E. Cronin, D. Kormann, and P. McDaniel, “Analysis of Security Vulnerabilities in the Movie Production and Distribution Process,” *Proceedings of the 2003 ACM Workshop on Digital Rights Management*, October 2003.
- [2] P. McDaniel, S. Byers, D. Kormann, L. Cranor, and E. Cronin, “Exposing Digital Content Piracy: Approaches, Issues, and Experiences,” *Proceedings of the Thirty-Eighth Asilomar Conference on Signals, Systems, and Computers*, November 2004.
- [3] J. A. Bloom, “Security and Rights Management in Digital Cinema”, *Proceedings of IEEE International Conference on Multimedia and Expo, ICME'03*, Baltimore, July 2003.
- [4] I. J. Cox, M. L. Miller, and J. A. Bloom, *Digital Watermarking*, Morgan Kaufmann Publishers, Inc., San Francisco, 2001.
- [5] J.P.M.G. Linnartz and M. van Dijk, “Analysis of the Sensitivity Attack Against Electronic Watermarks in Images,” *Proceedings of Workshop on Information Hiding*, pp. 258-272, 1998.
- [6] I. J. Cox, J. Kilian, T. Leighton, and T. Shamon, “Secure Spread Spectrum Watermarking for Images, Audio and Video,” *Proc. International Conference on Image Processing, ICIP'96*, Vol III, pp. 243-246, 1996.
- [7] T. Kalker, G. Depovere, J. Haitsma and M. Maes, “A Video Watermarking System for Broadcast Monitoring,” *Security and Watermarking of Multimedia Contents*, SPIE Vol. 3657, pp 103-112, 1999.
- [8] C. Honsinger, “Data Embedding Using Phase Dispersion,” *IEEE Seminar on Secure Images and Image Authentication*, 3(9): 51-57, 2000.
- [9] F. Hartung and B. Girod, “Watermarking of Uncompressed and Compressed Video,” *Signal Processing*, 66(3): 283-301, 1998.
- [10] J. Haitsma and T. Kalker, “A Watermarking Scheme For Digital Cinema”, *Proceedings of the IEEE International Conference on Image Processing*, Vol. 2, pp. 487-489, 2001.
- [11] J. Fridrich, Digital Watermarking by Adding Random, Smooth Patterns, United States Patent 6,101,602, Issued Aug. 8, 2000.
- [12] J. Fridrich, “Methods for Data Hiding,” Technical Report, Center for Intelligent Systems & Department of Systems Science and Industrial Engineering, SUNY Binghamton, 1997.
- [13] J. Lubin, J. A. Bloom, and H. Cheng, "Robust, Content-Dependent, High-Fidelity Watermark for Tracking in Digital Cinema", *Security and Watermarking of Multimedia Contents V*, Ping Wah Wong, Edward J. Delp, Editors, Proceedings of SPIE Vol. 5020, 2003.
- [14] J. A. Bloom and M. L. Miller, “Informed Detection Revisited,” *Proceedings of the 2004 International Workshop on Digital Watermarking*, Seoul Korea, October 2004.
- [15] H. Cheng, “Temporal Video Registration,” *Proc. of IEEE Int'l Conf. on Acoustics, Speech and Signal Processing*, ICASSP'03, Vol. 3, pp. 489-492, 2003.
- [16] H. Cheng and M. Isnardi, “Spatial, Temporal and Histogram Video Registration for Digital Watermark Detection,” *Proc. of IEEE Int'l Conf. on Image Processing*, ICIP'03, Vol. 2, pp. 735-738, 2003.