EURASIP Journal on Applied Signal Processing 2002:2, 126–132

The First 50 Years of Electronic Watermarking

Ingemar J. Cox

NEC Research Institute, 4 Independence Way, Princeton, NJ 08540, USA
Email: ingemar@research.nj.nec.com

Matt L. Miller

NEC Research Institute, 4 Independence Way, Princeton, NJ 08540, USA
Email: mlm@research.nj.nec.com

Received 8 October 2001 and in revised form 28 October 2001

Electronic watermarking can be traced back as far as 1954. The last 10 years has seen considerable interest in digital watermarking, due, in large part, to concerns about illegal piracy of copyrighted content. In this paper, we consider the following questions: is the interest warranted? What are the commercial applications of the technology? What scientific progress has been made in the last 10 years? What are the most exciting areas for research? And where might the next 10 years take us? In our opinion, the interest in watermarking is appropriate. However, we expect that copyright applications will be overshadowed by applications such as broadcast monitoring, authentication, and tracking content distributed within corporations. We further see a variety of applications emerging that add value to media, such as annotation and linking content to the Web. These latter applications may turn out to be the most compelling. Considerable progress has been made toward enabling these applications—perceptual modelling, security threats and countermeasures, and the development of a bag of tricks for efficient implementations. Further progress is needed in methods for handling geometric and temporal distortions. We expect other exciting developments to arise from research in informed watermarking.

Keywords and phrases: digital watermarking, data hiding, steganography.

1. INTRODUCTION

In 1954, Emil Hembrooke of the Muzac Corporation filed a patent entitled "Identification of sound and like signals" [1] in which is described a method for imperceptibly embedding an identification code into music for the purpose of proving ownership. The patent states "The present invention makes possible the positive identification of the origin of a musical presentation and thereby constitutes an effective means of preventing such piracy, that is, it can be likened to a watermark in paper." Electronic watermarking had been invented![1]

Since that time, a number of watermarking technologies have been developed and deployed for a variety of applications. Interest in embedded signaling continued throughout the next 35 years. For example, systems were developed for advertisement verification and device control both of which are discussed in the next section. However, electronic watermarking (particularly digital watermarking) did not receive substantial interest as a research topic until the 1990's. In the first half of that decade, interest in the topic expanded rapidly and today entire conference proceedings are devoted to the subject.

This increase in interest was motivated by copyright concerns that became acute with advances in computer technology and the development of the Web. These technologies enable the perfect copying and distribution of copyrighted material to almost anywhere in the world at almost no cost. To address these concerns, a number of industry technology groups were established, perhaps the best known being the Copy Protection Technical Working Group (CPTWG) and the Strategic Digital Music Initiative (SDMI). The former is concerned with digital video content stored on DVD discs and the latter with digital music.

These industry groups recognized that cryptography can only protect the *distribution* of content and that once a customer decrypts it, all protection is lost. Watermarking can complement cryptography, providing protection after decryption, even after the content has entered the analog world. Nevertheless, initial expectations of watermarking were probably too high, particularly with respect to intentional efforts

[1] To the best of our knowledge, this is the earliest reference to electronic watermarking. We do cite a patent dated 1953 [2] later as an example of device control. However, the patent description is ambiguous as to whether this is really watermarking or not. If readers are aware of earlier technology, please let us know.

to remove a watermark from content. However, subsequent responses to requests for proposals met specifications and the slow adoption of the technology is, in our opinion, due primarily to the diverging business interests of the three industry groups—content owners and manufacturers of consumer electronic equipment and computers—that must reach a concensus.

This leads us to the main question of the present paper: is the current business and academic interest in watermarking warranted?

From a business perspective, the question is whether watermarking can provide economic solutions to real problems. Current business interest is focused on a number of applications that broadly fall into the categories of security and device control. From a security perspective, there has been criticism that many proposed watermark security solutions are "weak," that is, it is relatively straightforward to circumvent the security system. While this is true, there are many business applications where "weak" security is preferable to no security. We therefore expect that businesses will deploy a number of security applications based on watermarking. In addition, many device control applications have no security requirement, since there is no motivation to remove the watermark. Device control, particularly as it pertains to the linking of traditional media to the Web, is receiving increased attention from businesses and we expect that this interest will increase. Business usages of watermarking are discussed in more detail in Section 2.

From an academic perspective, the question is whether watermarking introduces new and interesting problems for basic and applied research. Watermarking is an interdisciplinary study that draws experts from communications, cryptography and audio and image processing. Interesting new problems have been posed in each of these disciplines based on the unique requirements of watermarking applications. Commercial implementations of watermarking must meet difficult and often conflicting economic and engineering constraints. These problems are addressed in more detail in Section 3.

Our opinion is that current interest in watermarking is warranted, although expectations in the early 1990's were often too high. This raises the final question of this paper: what are the most exciting areas for research and where might the next 10 years take us? We address these questions in Section 4.

2. COMMERCIAL APPLICATIONS

Is watermarking important commercially? To answer this question, we begin by noting that a number of companies have employed watermarking for several years—decades in some cases. We regard this as empirical evidence that watermarking is, indeed, commercially viable. We then address the question more analytically, examining the practicality of some proposed watermarking applications, in light of current research in the field. We conclude that, although businesses may need to lower their expectations of performance, watermarks can serve most of these functions economically.

2.1. *Early uses of watermarking*

The applications of watermarking are well known and can be broadly classified as copyright control (owner identification, proof of ownership, transaction tracking, and copy control) broadcast monitoring and device control. What is less well-known is that watermarks have been deployed for some of these applications for several decades.

Owner identification appears to have been pioneered by the Muzak Corporation [1]. Their system, which used a notch filter to block, with varying duration, the audio signal at 1 kHz, encoded identification information using Morse code. The system remained in use until the early 1980's, when a change in Muzak's business model ended their interest in identifying music they owned [3]. More recently, each DiVX DVD player manufactured by the now defunct DiVX Corporation, contained watermark embedding circuitry that supported transaction tracking that was intended to deter piracy.

Advertisement monitoring and audience measurement companies have also used embedded signaling for some time. Both Nielsen Media Research, now part of VNU, and Competitive Media Reporting (CMR), now part of Taylor Nelson Sofres, employ watermarking to provide advertisement verification services and these systems have probably been in use for about 10 years. More recently, Verance Corporation has introduced a service to monitor television and radio broadcast media using their audio watermarking technology.

A number of companies have experimented with embedded signalling for device control purposes. In a 1962 patent assigned to Lynch Carrier Systems Inc., Noller [4] described a "inband signalling system" designed to control telephony equipment. In an even earlier patent assigned to Musicast Inc., Tomberlin et al. [2] proposed to distribute music to businesses by partnering with existing radio broadcasters. Their patent describes embedding a low frequency 30 Hz control signal at the point of transmission which will allow receivers to remove advertisements. Baer of the Sanders Associates Inc. was issued a patent in 1976 [5] for a video watermark intended for interactive television applications. In a 1981 patent assigned to Dolby Labs [6], Dolby describes "A sub-audible in-band tone system …for identifying an FM stereophonic radio broadcast which is specially encoded, as with dynamic range improvement encoding or quadraphonic encoding, …[and] which can control a visual display and switch in appropriate signal decoding circuitry when the tone is detected." A few years later, in 1989, Interactive Systems Inc. was awarded a patent [7] for a "Method and apparatus for in-band, video broadcasting of commands to interactive devices." An early application of this technology was in the synchronization of children's toys with live-broadcast or recorded video. Interactive Systems, has since become VEIL[2] Interactive Technologies. This company offers watermarking solutions for a number of different applications including interactive television, interactive toys, and advertisement monitoring.

[2]VEIL stands for Video Encoded Invisible Light.

2.2. Potential applications

In addition to the ongoing watermarking activities of Nielsen, CMR, and VEIL Interactive Systems, several new applications have sprung up in the 1990's. Whether these applications prove economically viable remains to be seen, but we can offer some educated guesses based on how well the current state of technology satisfies the applications' requirements. We discuss, in turn, the applications of transaction tracking (also known as fingerprinting), proof of ownership, copy control, legacy system enhancement, and a range of applications we refer to broadly as database linking.

2.2.1 Transaction tracking

In transaction tracking, or fingerprinting, a unique watermark is embedded into each copy of a Work. Typically, the watermark identifies the legal recipient of the copy, and can be used to trace the source of illegally redistributed content. Large-scale use of watermarks for transaction tracking, such as that implemented by DiVX,[3] is known to be vulnerable to collusion attacks, which usually require fewer than 20 copies to be effective. In such an attack, an adversary obtains several copies of a single Work, each with a different watermark, and uses them to obtain an approximation of the original, unwatermarked Work. Most existing or envisioned watermarks can be removed using fewer than 20 copies [8, 9, 10]. Thus, an adversary with 20 DiVX DVD players could produce watermark-free copies. Nevertheless, the system might still be worthwhile, since it would catch adversaries who lacked the dilligence or knowledge to perform these attacks, and this might prevent enough piracy to justify the system's cost.

On the other hand, smaller-scale transaction tracking applications, in which collusion attacks are unlikely, can probably be implemented with a very high degree of security. For example, if a Hollywood studio wishes to distribute movie dailies to a few key personnel, it is extremely unlikely that even two executives would collude in leaking these movie clips to the press. By using the original clip during the detection process (informed detection), a studio could design a watermark that is very difficult to remove.

2.2.2 Proof of ownership

Muzak's original interest in watermarking was to distinguish between theirs and similar recordings. The most ambitious form of such an application, which has received much attention in the watermarking literature, is the use of watermarks to actually *prove* ownership in a court of law.

In 1996, Craver et al. pointed out that there is an inherent problem in using watermarks for proof of ownership [11]. Specifically, with many watermarking methods, it is possible for adversaries to make it appear as though all distributed copies of a Work contain *their* watermarks, even though those marks were never actually embedded. However, the original paper suggested a solution to this problem, involving a cryptographic link between the watermark and the original Work, and we have seen no weakness in this solution. We therefore believe that, with a properly designed system, it is technically possible to prove ownership with watermarks. Business and legal issues appear to be the only hurdles to adoption of such a technology.

2.2.3 Copy control

If every recording device contained a watermark detector, watermarks could be used to prevent copying of copyrighted material. Watermarking for copy-control has been the subject of much R&D effort through the latter half of the 90's.[4]

There are two main areas of difficulty in implementing a watermarking copy-control system—one technical, the other political. The technical problem is that everyone must be able to detect the watermarks and within this context, current technology can only provide weak security.[5] Nevertheless weak protection against copying can still be economic. For example, it is very easy to circumvent the Macrovision system for preventing copying on VHS tapes (which is not based on watermarking). In fact, several legitimate pieces of video equipment remove Macrovision protection as a side effect. But Macrovision still prevents a great deal of casual copying, and studios have continued using it for several years. Thus, even if the copy protection provided by watermarking is weak, it may still be worthwhile.

The more serious problem in implementing a watermarking copy-control system is the political problem of persuading manufacturers to include watermark detectors in their recording devices. These detectors add cost yet do not necessarily add any value to the equipment. In fact, they *reduce* the value, since many consumers would like to be able to make illegal recordings. Thus, equipment manufacturers must be forced to include detectors, by a combination of laws and contractual obligations. The political wrangling that results, together with conflicts over patent rights, are probably greater impediments to the deployment of these systems than any technical problems.

2.2.4 Authentication

Authentication is well understood. A digital signature can be embedded as a watermark in a Work. And in fact, Epson offers a camera systems that does just this. An advantage of this arrangement is for legacy systems. There has been concern because embedding a signature alters the Work. However, the recent introduction of erasible watermarks [14] should dispell this concern.

[3] Perhaps the most ambitious implementation of transaction tracking was deployed by the DiVX Corporation in the late 1990's. Each DiVX-enabled DVD player embedded a unique watermark into video that it played. If the video was subsequently pirated and redistributed, the DiVX Corporation could use the watermark to identify the exact player used, and, thereby identify the source of the pirated Work.

[4] As noted earlier, two on-going, high-profile projects to deploy such copy-control systems have been undertaken. The Copy Protection Technical Working Group (CPTWG) has worked on a system for protecting video on DVD since 1995, and the Secure Digital Music Initiative (SDMI) has worked on an audio system since 1999.

[5] This is because the general availability of detectors permits adversaries to apply a sensitivity attack [12, 13].

2.2.5 *Legacy system enhancement and database linking*

Watermarking may also play a valuable role in enhancing the functionality of legacy systems while maintaining compatibility with deployed devices. For example, Schreiber et al. [15] proposed "a compatible high-definition television system using the noise-margin method of hiding enhancement information." Although this high-definition television system was not adopted, similar proposals have more recently been made for digital radio [16, 17].

Recently, Digimarc has pioneered a class of device control applications that link traditional print media to associated websites in a product called MediaBridge. Philips has also demonstrated an audio watermarking technology for music [18]. When the music is played, the audible signal can be digitized using the microphone present in many PDA's and the PDA can decode the watermark and thereby identifying the song. If the PDA has a wireless Web connection, it can then link the song to an associated site that, for example, may provide additional information or offer the song for purchase. Similar technology has also been demonstrated by Microsoft [19].

Work in these areas is still in its infancy. However, it is expected that security will not become an issue. Rather, robustness, fidelity, and payload requirements are the key issues and we believe that these requirements can or will be met.

The 1990's also say a new business development, the creation of companies that promoted watermarking as their core competence. This is as opposed to previous companies who exploited watermarking technology but promoted a product or service that watermarking was a part of. Whether these companies remain focused on developing the core watermarking technology or ultimately develop an application market, it is clear that watermarking has, is and will continue to be used.

3. RESEARCH PROGRESS

Is watermarking a worthwhile topic of research? To answer this, we need to ask whether watermarking is leading to interesting problems in basic research and whether engineering progress is leading to practical solutions.

3.1. *Basic research*

Very early work on watermarking was essentially heuristic, in part, because watermarking was not recognized as a distinct technology. This began to change in the late 1980's and early 1990's when a number of published papers described a variety of different watermarking algorithms. A more rigorous understanding of watermarking then began to be developed, beginning in the mid-1990's.

Perhaps the most significant progress has been in the development of increasingly sophisticated models of watermarking. In the early 1990's it became common to model watermarking as a communications channel in which the cover Work and any subsequent distortions between the time of embedding and detection were treated as noise. The constraint of imperceptibility was met by imposing a global power constraint at the embedder.

In these early systems the added watermark signal is independent of the cover Work and we refer to this as blind embedding. Similarly, blind detection refers to the detection of a watermark signal in a cover Work, the detection being independent of the unwatermarked Work.

In 1999, contemporaneous results from [20, 21, 22] recognized that watermarking is more accurately modelled as communications with side information [23]. The resulting watermark algorithms are referred to as informed embedding and/or informed encoding [24]. This is because the added watermark pattern is a function of the cover Work.

This model was further refined with the introduction in [25, 26, 27] of Costa's paper, "Writing on Dirty Paper" [28] to the watermarking community. Costa examined the capacity of a channel with two additive white Gaussian noise sources, the first of which is known. In Costa's analogy, the first noise source represent dirty paper. The watermark embedder writes a message on the dirty paper using only a limited quantity of ink. Then, during transmission, more unknown noise is added to the paper before its receipt at the detector, which has no knowledge of either the first or second noise source. Costa's surpring result is that the channel capacity is independent of the first noise source. This result has profound implications for watermarking where the cover Work can be thought of as the first, known noise source. It implies that, with the right coding, the capacity of a watermarking system may be independent of the cover Work even when blind detection is utilized. This work has since been extended [27] to more closely approximate the case for watermarking.

Watermarking must not only transmit a message, but it must also maintain the fidelity of the underlying cover Work while surviving common distortions that the cover Work may undergo. These fidelity and robustness constraints often conflict. As noted previously, early watermarking systems applied a global power constraint to satisfy fidelity constraints. In 1995, it was recognized that the fidelity constraint required a perceptual model that allowed the embedded watermark signal to be locally varied in response to the local properties of the corresponding cover Work [29]. Many watermarking systems have been developed that employ a variety of perceptual models and they are generally superior to algorithms with no such models [30, 31, 32]. These perceptually-based watermark embedders were early forms of informed embedding, since the added watermark signal is dependent on the cover Work.

The watermark communications channel can often be considered to exist in a hostile environment. For example, when watermarking is employed for copyright purposes, there is often a strong incentive to remove the watermark. The last 10 years has seen significant progress in the development of attacks and counter attacks. Researchers have documented many different attacks that an adversary might apply, for example, collusion attacks [8, 33], ambiguity attacks [34], copy attacks [35], sensitivity and gradient descent attacks [12, 13]. In addition, solutions to some of these threats, such as for ambiguity and copy attacks, have also been

proposed [34, 36]. This effort has provided valuable insights into what the threats are, under what conditions these threats can be neutralized and the limitations of current systems.

Spread spectrum communications was introduced at the same time as perceptual modelling in order to deal with the conflicting fidelity and robustness requirements [37, 38]. Spread spectrum communications spreads a narrow band signal over a much wider frequency band such that the signal-to-noise ratio in any single frequency is very low. However, with precise knowledge of the spreading function, the receiver is able to extract the transmitted signal, summing up the signals in each of the frequencies such that the detector signal-to-noise ratio is strong. These characteristics allow weak watermark signals to be embedded that, in many cases, can be reliably detected. Spread spectrum communications is also difficult for an adversary to detect or jam and this is a further advantage of the technology.

Around 1998, more rigorous quantitative measures of performance were introduced based on traditional false alarm and bit error rate techniques [39]. Theoretical progress was coincident with the development of more sophisticated models. Traditional false alarm and bit error rate techniques have been applied to watermarking [39]. In addition, more accurate noise models were developed, particularly for quantization noise. The effect of quantization noise on watermarking is important because cover Works are often heavily quantized as part of lossy compression. The effect of quantization was rigorously modeled in [40] in which it was recognized that dither modulation was analogous to watermarking.

3.2. *Applied research*

This conceptual and theoretical progress paralleled significant engineering progress by small and large companies as well as universities. Much of this effort has focused on meeting fidelity, robustness and economic constraints.

Steady progress has been made, particularly with respect to the problem of geometric and temporal distortions. Several different strategies have been pursued that can be categorized as exhaustive search, explicit synchronization/registration, autocorrelation [41], invariants [42, 43], and implicit synchronization [44, 45, 46, 47]. While no breakthroughs are expected, a number of design choices are now available.

There has also been significant experimentation with a variety of different marking spaces. For example, frequency decompositions such as DCT, FFT, wavelet and Fourier-Mellin transforms. While there is no clear superiority of one space over another, considerable expertise has been developed for embedding watermarks in MPEG and JPEG encoded content. This work facilitates the design of very inexpensive watermark detectors that are suitable for large-scale deployment, for example, for DVD copy control applications.

4. THE FUTURE

Some of the advances discussed above are in their infancy, and much interesting work remains to be done. In some cases, we believe that significant results may be imminent, which makes an area exciting. In other cases, we do not see any breakthroughs on the horizon, but significant results would increase the suitability of watermarks for a wider variety of applications, and are therefore worth further study.

We believe that informed watermarking offers significant near-term improvements. While proposed codes for informed embedding are computational efficient they are not robust to valumetric scaling. A solution was briefly proposed in [48], but further investigation is needed to realize computationally efficient and robust codes.

Handling geometric/temporal distortions in a blind detector remains a difficult problem. However, a number of different approaches have been investigated and incremental progress is being made. A breakthrough is probably not imminent, but progress in this area would lead to significantly more robust systems.

While many papers have illustrated the use of a variety of fidelity models, there has been very little work [49] on how to optimally embed a watermark with fidelity and robustness constraints. We expect this to become a fruitful new area of research.

Not all watermarks need to be secure. This is especially true of applications for which there is no adversary, for example, linking media to the Web. And even weak security has value in many business environments. Nevertheless, it remains an open question whether a watermark system can be designed that permits public detection of the watermark while preventing an adversary from removing the watermark. The authors of the present paper are divided about whether it is even theoretically possible to do this. Sensitivity analysis and gradient descent attacks appear to threaten any watermarking system in which the detector is publicly available. A number of researchers have attempted to design secure, public watermarking systems, but all appear susceptible to attack. It would be interesting to know whether such a system is even possible.

If the past is any prediction of the future, then it is clear that watermarking technology will continue to be used by businesses. It is also reasonable to expect that legacy systems will be enhanced through the use of embedded signalling in order to maintain backward compatibility. The linking of traditional media to the Web is still in its infancy and it remains uncertain whether consumers will value services that facilitate commerce and discovery. So we conclude this paper with an exercise for the reader. Imagine that all content is watermarked with a technology that is open, free and can be read by anyone. As such, any and all content is identifiable by consumer devices. What services might these devices provide?

ACKNOWLEDGEMENT

Portions reprinted, with permission, from 2001 IEEE Fourth Workshop on Multimedia Signal Processing, 225–230, © 2001 IEEE [50].

REFERENCES

[1] E. F. Hembrooke, "Identification of sound and like signals," United States Patent, 3,004,104, 1961.

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.