

Digital Watermarking Security

Jonathan Blake and Shahram Latifi

University of Nevada, Las Vegas NV 89154-1022, USA

E-mail: jonathan.blake@unlv.edu

ABSTRACT

As creative works (e.g. books, films, music, photographs) become increasingly available in digital formats in a highly connected world, it also becomes increasingly difficult to secure intellectual property rights. Digital watermarking is one potential technology to aid intellectual property owners in controlling and tracking the use of their works. Surveys the state of digital watermarking research and examines the attacks that the technology faces and how it fares against them. Digital watermarking is an inherently difficult design problem subject to many constraints. The technology currently faces an uphill battle to be secure against relatively simple attacks.

Keywords: Watermark, security, intellectual property

1. INTRODUCTION

Watermarking was originally a technique to embed a visible figure within the structure of a sheet of paper. Such a watermark is formed by thinning the paper where the figure should appear. The mark would be hidden from view under typical viewing conditions but become visible when light is shone through the paper. Watermarks were used to authenticate paper currency and encode information about the publisher of a work, among other uses.

Digital watermarking seeks to perform an analogous function in digital media rather than in paper. It inseparably embeds a hidden, secondary signal (the watermark or mark) within a primary signal (the cover data or cover signal). The primary signal is often intended for human consumption (e.g. audio or image data). The secondary signal provides additional information about the primary signal. This additional information can be used for a wide variety of purposes: copy prevention, copyright protection, authentication, copy control, device control, etc. Watermarks differ from headers and other out-of-band data channels by virtue of their being embedded within the primary signal itself.

Watermarking has become increasingly important in securing intellectual property rights. By embedding copyright information within the work itself rather than in headers that can be easily lost in format conversions or malicious attack, digital watermarking holds out the hope that theft can be discovered and minimised. Furthermore, watermarking provides an additional layer of security beyond conventional cryptographic protections. Once a work has been decrypted, for example, cryptographic methods no longer protect it. An attacker is free to use the content at will. If the work were watermarked,

however, it may be possible to prevent unauthorised use or track down the attacker.

Some uses of digital watermarking require no security because there is no incentive to disrupt the watermark. Most uses, on the other hand, require security against active attacks of various kinds. Much research has gone into creating digital watermark systems which are secure against malicious attack.

Securing digital watermarks has proven inherently difficult for a wide variety of reasons. Many uses (e.g. protecting digital video in consumer devices from illicit copying) require watermarks and watermark detectors to be in the control of potential adversaries (any consumer could be an attacker) for an indefinite period. This gives the attacker time to examine the function of the detector and manipulate the watermarked data. Even tamper resistant hardware can provide only a little added security if an attacker has a practically unlimited supply of fresh copies of the hardware as can be expected in the case of consumer electronics where attackers, such as organised crime, have the financial incentive and means to carry out the attack. Securing watermarks against all the various possible attacks under these conditions is difficult¹.

This paper provides a brief introduction to watermarking techniques, various categories of attacks, and a sample of known attacks.

2. BACKGROUND

Watermarking systems are composed of, at the most fundamental level, a watermark embedder (which attempts to incorporate the watermark into the cover signal) and a watermark detector (which attempts to detect a watermark in a received signal). At the core of

each is a watermark encoder and decoder respectively. The encoder translates the input signal into a watermark signal. In the case of an image watermarking system, for example, the input signal which will be embedded may be translated into a pixel array which is scaled to fit the dimensions of the cover image. The input could then be added in the spatial domain to the cover data, pixel for pixel. The encoding and embedding may take place in the cover data's native domain or in a transformation domain (e.g. FFT, wavelet, DCT, etc.)². The decoder then attempts to find the watermark in the signal it receives and reconstruct the message originally sent to the embedder.

Watermarking systems are further categorised by the information used by the embedder and detector. Figure 1 depicts a watermarking system with an informed detector, the original cover data is available to the detector and is subtracted from its input signal. The residual signal is then decoded to determine if a watermark is present. If so, the contents of its message is decoded. The availability of the cover data greatly simplifies the job of the detector, but this also limits the application of this technology. It could be used to prove ownership of disputed intellectual property, for example, where the owner can be expected to provide the secret key and the cover data. Informed detection would be useless in consumer electronics where it would be infeasible to hold the original of every watermarked work in case it is ever needed for comparison.

The embedder and detector usually also receive a secret key which changes the way the watermark is embedded in the cover data. In the case of image watermarking, a system may use a standard reference pixel array to encode a particular message. A copy

protection system may have, for example, one reference watermark to encode copying permitted and another to encode copying forbidden. If these reference marks were publicly known—which is often the case with standardised watermarking schemes—then it would be trivial to remove or mask the watermark if the reference marks were used as-is. Instead, secret keys change how the reference mark is embedded. This is analogous to how spread spectrum communication uses keys to resist jamming. Since the same key is used at the embedder and detector, this is also analogous to symmetric key cryptography.

Between the embedder and detector, the watermarked signal is transmitted over a noisy channel. This noise can represent any distortion, random electromagnetic noise, quantisation, compression, attenuation, malicious tampering, etc. Robust watermarking systems must be designed to function in spite of the noise that it can expect to encounter in its intended application. Robustness measures the likelihood that a watermark will be detectable and decodable after it has passed through a noisy channel.

To make watermarking feasible in applications where the detector can't expect the cover data to be available, blind detection as depicted in Fig. 2 must be developed. A blind detector must attempt to detect a watermark based solely on the standard reference marks, the secret key, and the signal it receives. It does not have the cover data available. This is a considerably more complex problem than informed detection. Blind detection is usually accomplished by correlating the input signal to the reference mark as modified by the secret key.

There are three basic methods of correlation used in watermarking systems: linear correlation, normalised correlation, and correlation using a coefficient³. Other

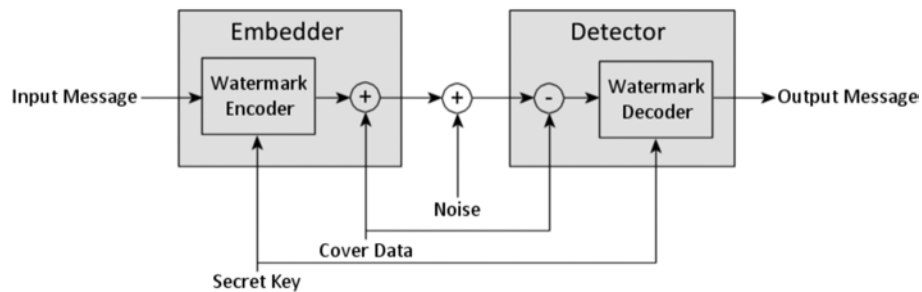


Figure 1. Watermarking system with Informed Detector.

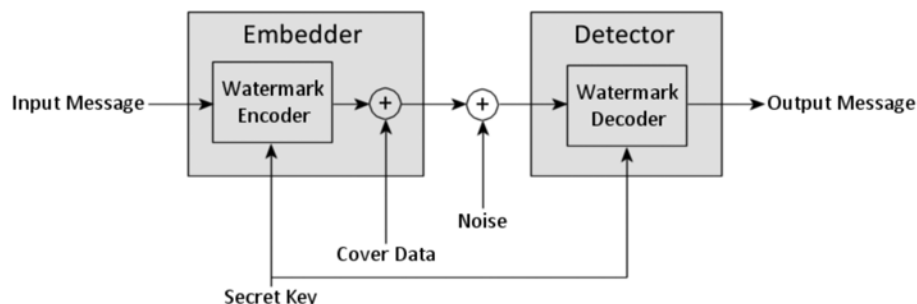


Figure 2. Watermarking system with Blind Detector.

correlation metrics have been proposed but most are equivalent to these three.

$$m_{lc} = \frac{1}{n} \sum_i c[i] \cdot w[i] \tag{1}$$

The first metric, linear correlation, is the simplest. It is defined in Eqn. (1) where c is a vector representing the input signal, w is a vector representing the reference watermark to be tested for, and n is the number of vector elements in c . This produces a scalar value which is roughly proportional to how likely it is that the watermark is present. A watermarking system which employs linear correlation would compare this value to a threshold value. If the correlation metric is above the threshold, the detector assumes that w is present and proceeds to decode it. Otherwise, w is assumed to be absent. An appropriate threshold value is chosen to balance the likelihood of false positives against that of false negatives, depending on the application based on experimental results.

$$m_{nc} = \frac{1}{n} \sum_i \frac{c[i]}{|c|} \cdot \frac{w[i]}{|w|} \tag{2}$$

Linear correlation is vulnerable to changes in the amplitude of c . This might happen if the contrast of an image is adjusted or if a signal is attenuated in transit. Normalising c and w to unit amplitude before correlating them as in Eqn. (2) can help overcome this vulnerability. This metric is normalised correlation.

$$m_{cc} = \frac{1}{n} \sum_i \frac{(c[i] - \bar{c})}{|c - \bar{c}|} \cdot \frac{(w[i] - \bar{w})}{|w - \bar{w}|} \tag{3}$$

Normalised correlation is still vulnerable to changes in the DC component of c such as a change in brightness in a watermarked image. This can also be overcome by subtracting the mean of each vector before computing the normalised correlation. This metric is called the correlation coefficient, defined in Eqn. 3.

One design goal of watermarking systems is imperceptibility. Since the cover data is primarily meant for human consumption, the watermark should not unduly interfere with that purpose. In some applications, such as medical imaging, imperceptibility of the distortion created by embedding a watermark is critical. Fidelity measures how perceptually similar a watermarked signal

is to the cover data. There are various metrics of fidelity ranging from simple calculations like the mean squared error to more sophisticated models of human perception as informed by experimental observations.

The effectiveness of an embedder/detector pair is measured as the likelihood that a detector will correctly detect and decode an embedded watermark immediately after the embedder stage, i.e. without any distortion before the detector stage. A system may not be 100 per cent effective by this measurement. Imperceptibility competes against the ideals of robustness and effectiveness. In general, the more effective and robust a watermark is, the more perceptible it is. A balance must be struck.

Embedding strength is another possible embedder input parameter which scales the mark to be embedded to make the system 100 per cent effective (in the ideal case) while keeping the watermark imperceptible. If the embedding strength is too high, the fidelity to the original cover data is damaged; if it is too low, the detector may fail to detect and decode the mark. In simple systems, this parameter can be preset, possibly based on experimental data, but this will result in some perceptible watermarks and some undetectable watermarks.

One strategy to improve the flexibility of the embedding strength parameter is to have a detector within the embedder. The detector offers feedback about detectability of the mark so the embedder can iteratively adjust the strength parameter for a particular cover work until the watermark is detectable.

Another strategy is to give the watermark encoder the cover data as another input as shown in Fig. 3. The encoder uses the characteristics of the input to adjust the embedding strength. In the frequency domain for example, an encoder may embed more information in high frequencies where the changes are less perceptible. Even better results can be achieved by incorporating more sophisticated models of human perception. This allows the embedder to fine tune the balance between perceptibility and robustness.

3. TYPES OF ATTACKS

It can be difficult to categorise attacks in order to study them more effectively. For example, one categorisation of attacks gives a list of what it considers four inherently different attacking concepts: removal,

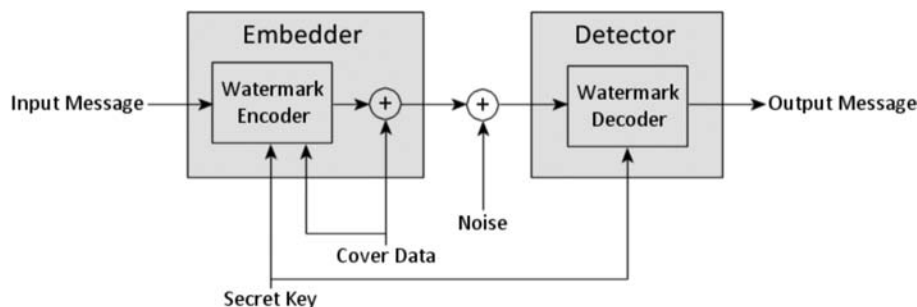


Figure 3. Watermarking system with informed embedder.

geometrical, cryptographic, and protocol⁴. Removal attacks aim at removing a watermark. Geometrical attacks seek to make the watermark undetectable by distorting the cover data. The distinction between these two types of attacks seems superfluous. Both have the effect of making the watermark undetectable. Cryptographic attacks are ambiguously defined by their similarity to attacks against cryptographic systems. Finally, protocol attacks are said to attack the concept of the watermarking application. It is unclear what this is intended to mean. It seems to be a catch-all category since they place their own novel attack (i.e. copying a watermark from one work to another) within this category. Overall, this is an unsatisfying categorisation. Better categorisations can be obtained by asking simple questions about the attacker.

3.1 What Does the Attacker Know?

Another useful way to categorise a potential attack is by what information the attacker knows which can help in carrying out an attack.

Perhaps the most optimistic is the assumption that the attacker knows nothing specific about a particular watermarking system. The attacker may have a work that may or may not be watermarked. This assumption also requires that all users of the watermarking system can be perfectly trusted to never become an attacker and to never make a mistake in preserving the confidentiality of the system. The attacker may know about watermarking and its weaknesses in general. A system that relies on these assumptions is only minimally secure.

It seems reasonable to assume that the attacker can gain access to a known watermarked work, perhaps the work that is under attack. This allows the attacker to perform a rudimentary level of analysis. If an attacker has more than one known watermarked work, a collusion attack (which will be discussed later) becomes possible.

Knowing the watermarking algorithm can also help an attacker. It is safest to assume that an attacker can learn the algorithm². In fact, the cryptographic community virtually always shares the details of cryptographic algorithms. This allows a wide audience to attempt to find weaknesses and perhaps improve the algorithms. The digital watermarking community follows suit. The assumption that an attacker knows a watermarking algorithm is also why secret keys are used. If the security of the system relies on the secrecy of the algorithm, then if that secrecy is ever breached, an entirely new system must be created. On the other hand, if the security of a system relies on the secrecy of a key alone, then a breach of key security is confined to data watermarked with that key.

Having a detector can also aid an attacker. Various attacks become possible when an attacker can slightly modify a watermarked work and can repeatedly consult the detector to see if the watermark has been removed. These oracle attacks can be quite effective even in the absence of knowledge of how the algorithm works. This level of access must be assumed in any consumer device.

3.2 What Does the Attacker Want?

There are three broad goals that attackers of watermarking systems may have: unauthorised detection, unauthorised embedding, and unauthorised removal. A fourth goal, unauthorised modification, is actually a combination of removal and embedding. Beyond this, an attacker may attempt to exploit a weakness in the way a watermarking system is deployed.

Detection may be a protected action in some watermarking applications. If a business charges its customers for the service of detecting unauthorised use of their intellectual property by embedding special watermarks prior to its distribution, for example, the business needs to prevent its competitors and customers from detecting the watermarks. The company wants to be the sole provider of detection. Detection can be further broken down into three discrete, progressively powerful actions: detection of the presence of a watermark, distinguishing one watermark from another, and decoding the message of the watermark.

Unauthorised embedding can frustrate systems that seek to authenticate the ownership of the cover data. Aside from encoding and embedding a new message, an adversary may also try to copy a legitimate watermark from one cover work to another⁴.

Unauthorised removal is probably the most typical form of attack against watermarking systems. This attack may more properly be thought of as rendering the watermark undetectable. It is not necessary to actually restore the original cover data to a near pristine state. All that is truly required is that the watermark is undetectable and that the cover data is still in an usable state. For example, one rudimentary form of watermarking is to replace the least significant bits (LSB) of the cover data with the message to be transmitted. Changes in the LSB are generally imperceptible. An attacker aiming to remove the watermark would not need to restore the LSB (probably impossible without side information); the attacker would only need to randomise the LSB or replace them with a constant value. Such an attack would make the watermark undetectable yet preserve the usefulness of the distorted watermarked cover data.

A further distinction within removal attacks is drawn. Masking attacks merely make the watermark undetectable to a given detector. For example, if a detector is incapable of detecting an image that has been flipped around the vertical axis, doing so masks the watermark for that detector. A slightly more sophisticated detector may also check the flipped image to prevent this attack. For this detector, the watermark has been preserved and remains detectable. In contrast, elimination attacks truly remove a watermark to the point that no detector no matter how sophisticated will be able to detect it.

4. EXAMPLES

Digital watermark technology has elicited many creative attacks. Only a few of the many attacks which have been published are included here. Others can be found in literature⁵.

4.1 Distortions

The goal of a watermark removal attack is to change watermarked data in such a way that it is still usable (e.g. looks or sounds good to a human being) but frustrates the watermarking system. A wide variety of distortions may fall within this category.

Most watermarking systems rely on synchronised data. The detection of a watermark relies on the correlation of the data to be tested with a reference mark as altered by a secret key. Any distortion that interferes with this correlation can cause the watermarking detector to fail. For example, shifting an image by one pixel, an audio track by one sample, or a video by one frame may be enough to foil a watermark detector. To counter this simple attack, a detector may do a search for the watermark within close proximity to the expected location or attempt detection after a number of common transformations. One drawback to this approach is that each additional test increases the likelihood of a false positive, i.e. the detector reporting a watermark where there is none. This must be taken into account when designing the system, usually by increasing the detection threshold.

There are many examples of simple, rather ordinary geometrical distortions that can fool a detector. In the case of two dimensional images, these distortions include cropping, small rotations, skewing, small aspect ratio changes, and so on. StirMark is a research tool originally created to simulate the distortion created by scanning a printed image⁶.

StirMark basically simulates a resampling process, i.e. it introduces the same kind of errors into an image that you would expect if you printed an image on a high-quality printer and then scan the image again with a high-quality scanner. The algorithm applies a minor geometric distortion, i.e., the image is slightly stretched, sheared, shifted, and/or rotated by an unnoticeable random amount and then resampled using Nyquist ... interpolation. In addition, a transfer function that introduces a small and smoothly distributed error into all sample values is applied, which acts like a small non-linear analog/digital converter imperfection typically found in scanners and display devices.

StirMark proved surprisingly effective against contemporary watermarking systems⁷. It has been used extensively in watermarking literature as a benchmark for robustness. StirMark continues to be an important benchmark in watermarking research due to the difficulty of designing watermarking systems which are robust against even simple geometrical distortions.

Another form of distortion attack is signal processing such as noise removal. Watermarking systems often add information to the high frequencies of a work because the alteration is less perceptible there. This creates a vulnerability that can be exploited to remove the watermark. Passing the watermarked data through a low-pass filter may be sufficient to remove the watermark.

4.2 Mosaic Attack

A mosaic attack segments the watermarked work into smaller slices that are too small to carry the watermark⁸. These slices are then presented in a unified way that makes it look like a single work. For example, an image can be split into smaller images and assembled on a web page to appear as one image. Many image watermarking systems work on blocks of pixels rather than on the whole image at once, so these slices would need to be smaller than the block size for this attack to work.

In general, a successful removal attack may be carried out by scrambling the watermarked cover data prior to the detector and descrambling it after the detector has failed to detect the watermark. The mosaic attack is a specific case of the scrambling/descrambling attack. It can be imagined that an attacker could manufacture and sell a hardware scrambler that would intercept a media signal prior to a detector chip and a hardware descrambler that would reassemble the media signal prior to output to a human being.

4.3 Collusion Attack

A collusion attack involves gathering more than one instance of watermarked works in order to gain added information about the watermark itself. A collusion attack can take two forms: either using several different works marked with the same watermark, or using several copies of the same work marked with different watermarks.

If several works can be gathered that have been marked with the same watermark (e.g. using a photographer's portfolio which has been watermarked with copyright information), these works can be averaged together. If some portion of the watermark is invariant across these works, the averaged result will correlate highly with the watermark. Each additional image will bring the watermark data into sharper focus. An attacker can use this averaged data as an approximate watermark to remove the watermark, or even to embed the watermark in other works. This removal is targeted and impacts the quality of the watermarked data in a mild way, coming relatively close to the original cover data.

The second form of the collusion attack requires the attacker to gather several of the same work marked with different watermarks. This may arise, for example, in copy control systems whose aim is to track authorised copies. In these systems, if an authorised copy is leaked to the public or to other unauthorised parties, the individualised watermark can be used to determine which authorised copy was leaked. If several authorised copies can be gathered together and averaged, the individual watermarks should be averaged out leaving an approximation of the original. If there is an invariant portion of the watermark, detection of the watermark may still be possible, but individualised information will be difficult to extract thereby frustrating a portion of the purpose of the watermark. It is possible to design systems that resist this attack⁹.

4.4 Ambiguity Attack

An ambiguity attack does not remove a watermark. Instead, it casts doubt on the ownership of a watermarked work¹¹. Watermarks are sometimes used to lay claim to a distributed work. An intellectual property owner can embed a watermark into a work and thereafter distribute it. If a question of ownership arises, the owner can provide their secret key to a third party (e.g. a court of law) to demonstrate that the work bears their watermark. The ambiguity attack makes the work appear to bear more than one watermark: the true watermark and one in the possession of the attacker.

In an informed detection system, this attack is as simple as creating a fake watermark and subtracting it from the distributed work in order to form a false original. If the true mark and the false mark are uncorrelated, then the true original will be correlated to the fake watermark and therefore also appear to be watermarked with the false mark. The owner and the attacker have seemingly equal claims on the work.

In blind detection systems, the attacker must find a false mark that is highly correlated with the distributed work. Since this false mark is very probably orthogonal to the true mark, its correlation to the watermarked work is due to the characteristics of the original work. Again, both the watermarked work and the original will appear to be watermarked with the false mark.

There has been some efforts to create non-invertible watermarking schemes which make it infeasible to find a second watermark for a given work¹⁰. This can be accomplished by using a cryptographic signature of the work as an input to the embedder alongside the secret key. This is simple to implement in an informed detector where the original is available to recompute the signature. A standard cryptographic hash function is sufficient to create the signature. This is not so simple in blind detectors. The signature must still be computable after any distortion or attack has altered the image. This is non-trivial.

4.5 Sensitivity Analysis Attack

If an attacker has unlimited access to a detector (which is often the case in consumer electronics), then it becomes possible to mount a sensitivity analysis attack, a rather sophisticated attack. There are three stages to such an attack. The first stage is to find a distorted version of the watermarked work which lies close to the detection region. This can be accomplished by incrementally altering the watermarked work until the detector no longer detects the watermark. This process need not guarantee that the distorted work will be usable as-is. Its role is solely to determine a point on or near the boundary of the detection region. The second step is to determine the normal to the boundary. This can be done by iteratively adding random vectors of increasing magnitude to the distorted work until the detector reports a watermark. With enough such vectors, an estimation of the normal can be made. The final step is to subtract

a scaled version of the normal from the watermarked work. The scaling of the normal can be adjusted until it is barely outside the detection region. Through this process, an unwatermarked work can be found that is as close as possible in vector space to the watermarked work. Assumed is that proximity in vector space is strongly related to being perceptually proximate, and that the detection region has a uniform normal vector.

One approach to making this attack less computationally feasible is to provide a random result if the detection metric falls within a range of the threshold. This increases the number of trials necessary to locate the boundary of the detection region, to find its normal, and to find a distorted version of the watermarked work that is reliably reported by the detector as unwatermarked. This approach does not eliminate the possibility of a sensitivity analysis attack, but it can make such an attack more costly. The drawback is that it necessarily increases the number of false positives and negatives.

Another approach is to alter the shape of the detection region boundary in such a way that it has many normals.

4.6 Gradient Descent Attack

Beyond a simple detector, if an attacker has not a detector that reports the value of the detection metric, it is possible to perform a gradient descent attack. The idea is similar to a sensitivity analysis attack in that it attempts to find the watermarked work closest to the watermarked work in vector space. It is possible using the detection metric to determine a gradient. It is assumed that descending the gradient from the watermarked work is the shortest path out of the detection region.

4.7 Histogram Attack

A fixed depth image watermarking system (i.e. all elements of the watermark vector have the same absolute value) is vulnerable to the histogram attack is introduced¹¹. Images with many distinct peaks in their histogram are especially vulnerable to this attack (e.g. images which have had their colour values reduced to only a subset of the full range of values possible). When a fixed depth watermark is applied to an image, it will transform single peaks into two separate peaks. From this information, the attacker is able to say that the pixels in the peak in the value closer to zero probably correspond to watermark pixels with the negative value while the pixels in the other peak probably correspond to watermark pixels with the positive value. The confidence of this estimate can be increased with each additional watermarked image which the attacker has access to. The attacker does not need a detector to mount this attack.

4.8 Copy Attack

The final type of attack covered here is the copy attack¹². It is saved for last because the first step is to perform a watermark removal attack such as several covered earlier that approximates the original cover

data. Once an approximate original has been obtained, the difference to the watermarked work can be taken. This difference approximates the value of the watermark. This can be added to another work in an attempt to watermark it, an unauthorised embedding.

One potential counter is to tie a watermark to a given work so that it can't be used elsewhere. This can be accomplished, for example, by using a signature of the work in computing the mark. This signature must not be altered by the inclusion of the watermark. Otherwise, the detector would be unable to compute the same signature as the embedder.

5. CONCLUSIONS

Securing watermarks is an inherently difficult design problem. The security of a mark is constrained by the data in which it is embedded and by the perceptual acuity of human observers. Many applications require that potential attackers have unlimited access to watermarked data and detectors. If even a single attacker is successful, the ease of perfectly copying digital data ensures that this single failure can benefit all potential attackers¹. Watermarks may catch the unwary attacker or present an inconvenience to a determined attacker, but it cannot presently stand on its own as a security measure. The field is still relatively young and may be able to produce secure watermarking systems in the future, but the cards are stacked against the technology.

REFERENCES

1. Arnold, M.; Schmucker, M. & Wolthusen, S.D. Techniques and applications of digital watermarking and content protection. Boston: Artech House, 2003, 9 p.
2. Cayre, F.; Fontaine, C. & Furon, T. Watermarking security: Theory and practice. *IEEE Trans. Signal Proce.*, 2005, **53**(10) pt II, 3976-987.
3. Cox, I.; Miller, M. & Bloom, J. Digital watermarking. San Francisco: Morgan Kaufmann Publishers, 2002, pp. 75-82.
4. Kutter, M.; Voloshynovskiy, S. & Herrigel, A. The watermark copy attack. *In Proceedings of the SPIE*, 2000, **3971**, 371-80.
5. Cox, I. & Linnartz, J. Some general methods for tampering with watermarks, *IEEE J. Selected Areas Comm.*, 1998, **16**(4), 587-93.
6. Kuhn, M. StirMark – Image-watermarking robustness test. 10 November 1997. <http://www.cl.cam.ac.uk/~mgk25/>

stirmark.html. [Accessed: 26 October 2008].

7. Petitcolas, F.; Anderson, R. & Kuhn, M. Attacks on copyright marking systems. In Lecture Notes in Computer Science, *edited by* David Aucsmith, Portland, Oregon, USA, 14-17 April, 1998, **1525**, pp. 218-238.
8. Petitcolas, F.; Anderson, R. & Kuhn, M. Information hiding: A survey. *In Proceedings of the IEEE*, 1999, **87**(7), 1062-078.
9. He, S. & Wu, M. Improving collusion resistance of error correcting code based multimedia fingerprinting. *In IEEE International Conference on Acoustics, Speech, and Signal Proceedings*. 2005. **2**, pp. ii/1029-ii/1032.
10. Craver, S.; Memon, N.; Yeo, B. & Yeung, M. Resolving rightful ownerships with invisible watermarking techniques: Limitations, attacks, and implications. *IEEE J. Selected Areas Comm.*, 1998, **16**(4), 573-86.
11. Maes, M. Twin peaks: The histogram attack on fixed depth image watermarks. *In Proceedings of the 2nd International Workshop on Information Hiding*, 1998, pp. 290-305.
12. Kutter, M.; Voloshynovskiy, S. & Herrigel, A. The watermark copy attack. *In Proceedings of the SPIE*, 2000, **3971**, 371-80.

Contributors



Mr Jonathan Blake received the BSc (Computer Engg.) from the University of Nevada, Las Vegas in 2003 where he is currently a graduate student in the Electrical Engineering program and also works in the Department of Institutional Analysis and Planning as a Data Warehouse Developer.

Dr Shahram Latifi received the MSc (Electrical Engg.) from Fanni, Teheran University, Iran in 1980. He received the MSc and PhD (Electrical and Computer Engg.) from Louisiana State University, Baton Rouge, in 1986 and 1989, respectively. He is currently a Professor of Electrical Engineering at the University of Nevada, Las Vegas. His research areas includes: Image processing, biosurveillance, biometrics, document analysis, computer networks, fault tolerant computing, parallel processing, and data compression.