Digital Image Watermarking in the "Real World"

Extended Abstract

Adrian Perrig Andrew Willmott
Computer Science Department
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh PA 15213

{Adrian_Perrig, Andrew_Willmott}@cs.cmu.edu

March 9, 1998

Abstract

Digital watermarking is a method of embedding identifying information in an image, in such a manner that it cannot easily be removed. An application of watermarking is copyright control, in which an image owner seeks to prevent illegal copying of the image. Commercial systems using such technology are now becoming available. In this paper we point out a number of the problems that watermarking faces as it moves into the "real world". We also examine weaknesses in the system of the current state-of-the-art methods that claim to provide digital image copyright protection. We will show how to overcome the shortcomings with respect to the real world and the entire digital image life cycle.

Keywords: Analysis of digital image watermarking systems, copyright protection, intellectual property protection.

1 Introduction

Since the widespread use of the Internet today, the term "slippery age" was used to express the short life span of data, the anonymity of users, the ease of fraud, etc. Namely the protection of *intellectual property* has become a major problem in the digital age. The ease of copying digital information without any loss of quality violates the *conservation of mass* property of traditional media, which inhibited wide global distribution in the past. On the Internet today it is possible

to duplicate digital information a million-fold and distribute it over the entire world in seconds. These issues worry creators of intellectual property to the point that they do not even consider to publish on the Internet.

To solve the problem of publishing digital images, researchers have come up with digital image watermarking. This method allows the owner of an original image to add an invisible watermark to the digital image before publishing it. The watermark serves to claim copyright on the image. The owner protects the watermark with a cryptographic secret key, preventing anybody not possessing the secret key from reading or even detecting the watermark. The watermark is also supposed to be robust against image tampering. Therefore anybody who wants to distribute the image further will also distribute the watermark with it, violating the copyright on the image. If the copyright holder can detect the fraud, he can prove ownership by showing that the image contains his proper private watermark.

This scheme works well in a "nice and tidy" clean-room research environment. But in the "real world" people play by different rules. The systems are not used as intended in the research environment and hidden back-doors are exploited. This report shows the problems that current digital image watermarking schemes face.

For the rest of this paper we will use the characters Alice, Bob and Mallory (as used in [Sch96]) to describe copyright protection scenarios. Alice in our case is the image creator and copyright holder and Bob is the Buyer of the image. In some other settings Alice and Bob always represent the righteous characters. Mallory on the other hand is the malicious attacker and tries to forge, steal, impersonate, eavesdrop, etc.

The convention for describing the watermarking procedure is that Alice *inserts* a watermark into the original image. To check the image for a watermark, she can *extract* it.

Relevance

The early work on digital image copyright protection has focused on the creation of a secure and robust watermark only. Many groups are working on this problem, for further information consult [CMYY96, Dig97, PJ96, Mar96, CKLS96, Rua97, TvSO95, ZK94]. These works are mainly concerned about the algorithmic watermark issues and they only touch the deployment problems marginally.

Later work has addressed some of the attacks on copyrighting systems. Cox and Linnartz have described in [CL97] the difficulties for implementing a rights management and copyright protection scheme for DVD's (digital versatile disk). They showed that it is possible to circumvent even sophisticated devices to get around the copyright

protection.

Anderson and Petitcolas have shown in [PAK98] ways to attack and break copyright protection schemes for digital images and audio. They have shown various real world problems but they have not taken the failure of a user into account. Many of the attacks have also been described in [Per97, HPR97].

This report takes the entire life cycle of a digital image into account - from the artist/creator to the end user. We will take a close look at how the user can negatively influence the security of a copyright system and how the software can minimize users mistakes. Further we are analyzing how well current state-of-the-art systems work considering real world attacks and various ways to circumvent protection schemes. Finally we will show ways to overcome shortcomings of today's systems.

2 "Real World" concerns

Unfortunately the "real world" differs significantly from research labs. Computer programmers or researchers assume that the end users will behave and think as "logically" as they do. For example image watermarking is used by a wide variety of users and should still be equally applicable and efficient for everybody. Unfortunately this is not the case for today's systems.

User Interface concerns

In the "real world" the users have only a limited understanding of the underlying mechanisms of image watermarking. They do not want to spend hours of training to use one function of their image processing software. Watermarking should be like a "black box" where the user enters his original image and by some magic the box outputs the watermarked image. No specific user understanding should be necessary.

Usually artists are the designers of images. The observation that artists do not like to degrade their work deliberately by inserting a watermark leads us to believe that they will either insert a weak watermark[1] or not insert anything at all. Security is not always a strong enough argument to convince artists to lower their image quality.

We can see that the user interface (UI) becomes very important: it must present the user with a clear model of the effects of watermarking and protect him from misuse. It needs to iron out users misconceptions about the watermarking technique. This observation is equally true for other security software. Whitten showed in [Whi98] that the security of the system should not rely on user understanding but be an implicit

[1] Today's watermarking schemes present a tradeoff between watermark strength and image degradation.

property of the software. Therefore the UI becomes a crucial point in any security system.

Legal issues

Let's step away from UI concerns and let's take a look at another way the real world challenges image copyrighting. The scenario is the following: in a country that does not adhere to the Berne convention [Cin95] on copyright protection a malicious person sets up his web-server distributing copyrighted images, music, etc. There is no way to prevent this person from his "illegal" distribution as the country does not provide the legal basis for prosecution. All forms of intellectual property share this common problem. The past has shown that such scenarios are not far fetched. In fact recent incidents prove that this attack is quite common.[2] The situation gets worse as the Internet expands its range with high speed connections to countries that are "traditionally" known for copyright infringement. If these countries do not change their laws, this problem can not be solved trivially by technical means.

The legal enforcement of copyright infringement in countries that enforce the Berne convention of copyright protection is not a simple task either. First, it is difficult to prove copyright infringement in court. The complication is that we cannot just say "Mallory distributed my copyrighted image on her web-server" but we need sound evidence for the fraud. For example an impartial witness could provide for a resolution. But again, things look differently in the real world. Web servers do not send non-repudiable responses. Therefore how can the witness really know where the data came from? In a possible scenario let us say that Mallory tries to convince Alice that Bob stole her image. But Bob is a good person and would not steal any images. So Mallory needs to trick Alice. One simple way to trick her is by DNS spoofing; when Alice accesses http://www.bob.com/image.gif on Mallory's machine, Mallory's DNS lookup function replies with a wrong address for www.bob.com, namely one of Mallory's servers. The server then delivers the "stolen" image to Alice's browser.

In the other case where Mallory really did steal Alice's image, Alice starts a law suit. But during this process, Mallory will surely remove the stolen image from the server. In the case where Alice first consults a notary to look at Mallory's web-site to confirm that the stolen image really is stored there, Mallory might refuse to send the image to notaries. This would certainly be difficult to achieve but it would be

[2] Examples include servers in Indonesia that distribute copyrighted music in high fidelity. Another example was given in [Wil98]: Warner Bros. was hunting bootleggers who were distributing Madonna's forthcoming album "Ray of Light" over the Internet.

technically feasible by excluding data delivery to certain domains.

Web spider issues

Web-spiders which scan the web for stolen images face the same problem as stated above: the web-server detects that the request originates from the spider and it will then not forward any illegal images to that site, or replace them with dummy images. Access controlled or pay-sites present another problem to the web-spider. It can not access the contents without paying or authenticating itself. Unfortunately we believe that the largest part of copyright infringement comes from access controlled sites. These problems present high barriers for any web-spider to overcome.

Various schemes have been proposed to do rights management [Int97, Cox96]. Rights management is where the information distributor can give usage access rights to the client. For example he could declare that only viewing, not printing of the image is possible. This sounds feasible at first, but again, things look differently in the real world. Even if the rights are enforced by using a tamper-proof smartcard in everybody's PC that checks the access rights prior to any action, the image can be stolen anyway. While the image is displayed on the screen, the information has to be present somewhere in the PC's memory. Therefore that memory can also be read by another program and saved. Watermarking faces a similar problem: the image already has to be present somewhere in memory before the watermark can be extracted. Similar concerns are discussed with respect to DVD rights management in [CL97].

Another reason things look differently in the real world is that the systems are used in ways not foreseen by the system designers. In watermarking we could imagine the following problem: Instead of showing the stolen image at once on a web-page, Mallory chops up the image into small blocks and creates many small images. In the web-page, the images are then arranged such that the viewer can see the original image again. Unfortunately, the individual images are too small to carry an extractable watermark. Only if multiple blocks were merged, the watermark could be extracted again. Therefore a web-spider could not detect the fraud, since it checks each image individually. This mosaic attack was also discovered simultaneously by Anderson and Petitcolas and is described in [PAK98].

Hacker circles

We will see in the next section which technical attacks are known up to today to remove watermarks. But even if Mallory wanted to remove a watermark and she knows about ways of using her image processing

toolkit with which she could potentially remove the watermark she would probably hesitate to do so. First, she knows that there is a large fine for copyright infringement and she is unsure if the attack would be successful. Second she is quite unsure of which transformations to apply - she did tests on her own watermarks, but are the commercial marks the same?

Hackers are globally well organized and already today there exist multiple programs to remove watermarks. You can see the following sites for watermark removal software [Xa, Xb, Kuh]. Therefore Mallory does not need to be a signal or image processing expert to remove the copyright with a high certainty. The power of such underground activity to create high quality watermark removal software should not be underestimated.

3 Technical problems of watermarking

The Watermarking Process

The technical challenge that watermarking presents is the successful concealment of the watermark signal in a much larger bandwidth medium. As [AP98] points out, usually we are more concerned with an attacker's ability to read or, worse, change or remove the watermark, than their ability to detect its presence. In the face of perfect compression methods, it is not clear that it is possible to conceal the presence of a watermark. For our watermark to be robust, the watermark information must be embedded in the target medium in such a way that removing this information irreparably damages the medium.

When considering static images, the commonly recognized transformations that a watermark should survive are:

- Rotation, scaling, translation, mirroring.
- Filtering (Gaussian blur, image sharpening, etc.)
- Adding noise to the image, adding jitter (duplicating and removing lines/columns of the image), cropping (removing the sides).
- Color remapping (color, quantization, adjustments in brightness or contrast)
- Lossy compression (JPEG, MPEG, fractal compression)

All of these transformations can preserve the value of the image to the user. After they have been applied, the image is still recognizably derived from the original image. This concept can be used as the foundation of a robust watermarking method [RP97]; if the image can be transformed into some space which is invariant to the value-preserving transforms listed above, the watermark can be applied in

that space before reversing the transformation to arrive back at the original image. This watermark will then also be invariant under these transformations.

An often-used method for covert channel communication over a higher-bandwidth medium is spread spectrum communications. This method can also be applied to watermarking images, as described in [CKLS96, RDB96]. It has the advantage of having a long history in both military and civilian digital communications, and is thus both well understood and of proven reliability. At the current time, it appears the most promising framework for embedding watermarks in either images or other digital media data streams.

Watermarks are most often used for copyright control. It is commonly recognized that no watermarking method can withstand all possible attacks. Thus, in the domain of network-accessible digital media, it may be better to regard a watermark as an aid to quickly comparing possible copies of a copyrighted work to the original, in order to locate copyright violations, than as an indicator of copyright in itself.

Attacks on Watermarks

The possible attacks against watermarks are wide and varied. Cox and Linnartz [CL97] and Anderson et al [PAK98] present a number of such attacks, including:

- Image modification attacks. These use image transformations such as those listed above.
- Bit-level attacks. If the attacker has access to a watermark presence detector, the contents and location of the watermark can be derived. This also makes it much easier to remove a watermark.
- Watermark-insertion attacks. If the attacker has access to a watermark insertion device, and the watermarking process is not a one-way function, it is possible to recover the original, unwatermarked image, by pre-distorting the copy, and rewatermarking it.
- Statistical averaging attacks. The attacker uses multiple watermarked images to estimate the watermark, and then subtracts this from the image. This is especially a problem with video since a large number of watermarked frames is available.
- Scrambling attacks. By inserting a scrambler before the watermark detector, and a de-scrambler after it, detection of the watermarking can be avoided.

As we can see from this list, even with a perfect watermarking method, there are various system-level attacks that can frustrate the secure use of watermarks in a copyright scheme.

Prioritizing Attack Resistance

As making a watermark resistant to a large number of image transformation attacks is a difficult task, it is important to prioritize these transforms. A watermark should aim to be most resistant to the most common attacks; as we have pointed out, complete watermark security is an unachievable goal. It is more realistic to aim for a high probability of recovery. A corollary of this is that a watermarking method must first of all be robust against those transformations used by legal users of the image. It is tempting to concentrate on the various complex attacks that malicious attackers might employ. However, this is pointless if the method is not properly resistant to image edits made by either the original owner or valid users of the image. We claim that, roughly in order of merit, a watermarking method must handle:

- Scaling, especially (filtered) down-sampling.
- GIF and JPEG compression. (Color quantization, and lossy compression.)
- Simple brightness, contrast, or gamma adjustment.
- Border cropping.

We will revisit this point in the next section.

Handling small images

Below a certain size it becomes difficult to embed and reliably recover a watermark from an image. Even the better watermarking methods have difficulty retaining watermarks in 100 by 100 pixel images that are lossily compressed. It is not clear that this problem can ever be solved by technical means, as it is stems from a simple lack of bandwidth. As the effective bandwidth of the picture is (roughly) proportional to its area, halving the dimensions of it reduces the available bandwidth by 4. Often small images are also color-reduced for efficient storage, further restricting the available bandwidth. This observation provides the foundation of the mosaic attack described in section 2.

It should be noted that as image size drops, the concern about copyright violations also drops. Small text extracts from larger works are allowed for personal use under the copyright act [Cin95]. In music, samples of a couple of seconds or less, especially if they have been distorted in some way, are generally accepted as "fair use". Longer, more recognizable samples require copyright clearance. [3]

Most of the literature on watermarks avoids mentioning this problem, because it seems obvious. However end users of watermarking are often not aware of the problems small images can cause, or if they are,

[3] This is not to say that such clearance is always gained.

are uncertain as to how large an image has to be before its watermark is secure. This makes it essential that a watermarking system has some mechanism for warning a user that a small-image watermark may not be robust.

4 A case study: Digimarc

The results presented here will be extended for the final draft.

As we have shown in section 2 the UI of security software is very important for the overall security of the system. The article [Kat97] showed that the Digimarc UI had various shortcomings. For example the dialog box where the user can choose the watermark's robustness, the user sees immediately the relation between image quality and watermark strength. It was observed that most users chose the best image quality, which resulted in the least watermark strength. Another problem was that the users inserted a watermark and then later applied image transformations, such as scaling and compression. For obvious reasons, this weakens the watermark. The user was not aware that watermarking should ideally always be the last step before publication. For these reasons, most of the images published did not contain the watermark at all or it was too weak to extract. Therefore Digimarc failed in providing the user with a usable implementation for image copyright protection. We will see in section 5 how the system can be improved to solve these problems.

Digimarc proposed to set up a web-spider [Dig97] which would scan web-servers and search for watermarked images. This approach fails to work for the same reasons as stated earlier. The web-server simply refuses to send any stolen image to the Digimarc web-spider. It would be quite complicated for Digimarc to set up a large number of different spiders to make this scheme to work. Another problem with this approach is that it will be difficult for Digimarc to prove that the data really came from that particular web-server and was not "injected" into the Digimarc database by a corrupt Digimarc employee. Yet another shortcoming of the Digimarc web-spider is that payment or access controlled sites are not checked. These sites are especially interesting for the image creator as he would like to make sure that his intellectual property is not sold anywhere else. Finally Digimarc only sends a monthly report to the customer. In the worst case the image could already have been distributed one month earlier. On the Internet one month is almost an eternity.

Another issue is that any user that wants the web-spider protection has only one choice: fully trust Digimarc and give them the user's private key. There is also no alternative for which watermarking method to use. Digimarc has a considerable market advantage as their system

comes directly bundled with Adobe Photoshop. The user is only left with the choice to "use it or leave it".

One of the greatest problems with commercial watermarking schemes is the difficulty in getting details on how these schemes work; for reasons of either competitive advantage or patent application, companies tend to keep their methods secret. This is "security through obscurity", a law called Kerkhöffs principle [Sch96], and means the companies are relying to some extent on the difficulty of discovering their algorithm to foil attacks, rather than on the strength of some cryptographic key. This has proved foolish in the past.

The Robustness of Digimarc Watermarks

We have made an informal investigation of the robustness of the Digimarc protection scheme in the face of common-place image operations in Photoshop. We used the example Digimarc image available from Digimarc's web site[4], which measures 215 by 142 pixels.

- **Compression** Average – watermark becomes unreadable when image is saved as low quality JPEG, even though the image quality is not noticeably reduced.
- **Color adjustment and quantization** Good – watermark is retained until the image is close to unrecognizable under various contrast, equalization and quantization operations.
- **Cropping, Translation** Good – much of the image must be overwritten before the watermark disappears.
- **Rotation** Poor – free rotations of 1 degree completely remove the watermark.
- **Scaling** Poor – scaling the image down by 15 pixels removes the mark, as does down-sizing by only 1 pixel and then applying high-quality JPEG compression.
- **Filtering** Good – resistance to Gaussian blurring and sharpening filters, and also to addition of noise, is high.

Rotation is a less common image operation, and thus Digimarc's extreme vulnerability to this may not be so important when considering non-hostile users. However, if the watermark can be removed by a small amount of rotation then many non-technical attackers will be able to remove the watermark easily from within Photoshop. The Digimarc method's vulnerability to scaling, on the other hand, is critical to non-hostile users.

In general we worry that the mismatch between technically impressive watermarking feats and valuable watermarking feats can confuse

[4] http://www.digimarc.com/cgi-bin/ci.pl?1+100011

end users. For instance, an impressive demonstration of Digimarc's capabilities could be given by scribbling over most of a watermarked image, and then demonstrating that the watermark remains intact. However, this scribbling destroys the value of the image, and thus there is no particular point in retaining the watermark.

Digimarc Software Distribution

We believe that one of the critical weaknesses of the Digimarc system is its provision of a watermark detection program[5]. The presence of attacks that can take advantage of this was pointed out in section 3. It is simple for a user, given a watermarked image, to apply various small transformations until the watermark is no longer detected.

This problem exists because Digimarc has two conflicting goals for its watermarking system. It is used both to detect copyright images on the web, by means of a search engine, and to provide viewers of the image with information about its author.

If the purpose of the watermark is solely to detect images that are being used without permission, there is no need for anonymous users to be able to detect the watermark's presence or its contents; indeed, security is strengthened if they cannot. Instead, the user of the watermarking software should only be able to search for watermarks they have embedded with their own key. This allows them to check their own images to ensure the watermark is still present after editing, without allowing attackers to do the same. Unfortunately, because Digimarc also wants to use the watermark to provide author identification, it must distribute free software to make watermark detection and retrieval possible for anyone, thus severely compromising security.

5 Improvements on the schemes

Now that we have seen all the issues arising when analyzing today's watermarking schemes, we will present possible improvements.

First we observed that watermarking is regarded as being critically important to solving the digital image copyright protection problem. But when copyrighting traditional paintings or photographs, watermarks are not used. So why should we rely so strongly on watermarks if they are complicated to handle and so easy to defeat? This forms the basis of a slightly different approach to copyright protection. A simple approach is based on a model analog to today's copyright office [Per97]. Every user registers the images to copyright at the copyright office. In case of dispute, the court of law can compare the disputed image with

[5]It is even possible to download a stand-alone watermark detector from the Digimarc web site [Dig97] if you do not have access to Photoshop.

the registered image. So far we are not using any watermarks. But watermarks are still useful for discovering fraud on the network because of the difficulty of automatic comparison of images. Watermarks can be extracted automatically when the secret key is known. We can therefore look at watermarking as a method for matching images efficiently rather than as a security device. We have therefore shifted the importance of watermarking to fraud detection.

It is most important that watermarks used for the purpose of image search in this way are kept separate from those used for the purpose of author identification. Using the same watermark for both purposes drastically weakens the security of the scheme. For image search watermarks, we do not need exact retrieval, as we can correlate the mark against those of the image(s) we are searching for[6], but we do need much greater robustness against watermark detection and removal. For public- information watermarks, the need for robustness is much less, but the need to retrieve the watermark intact is much greater.

The problems addressed with web-spiders can be solved by running the spider on a large ISP proxy server. For example the AOL proxy server could forward all the images that it has seen to the spider to check. With this method, access controlled sites could be checked as well as paying sites. Mallory could not refuse to send the image to the proxy server because it would affect a large number of customers that could not receive the image.

The problem of proving fraud to a court of law could be solved if each web-server would sign every response with his private key. We would therefore have a proof that the server sent the message to us. A payment protocol, such as Netbill, would also take care of this.

We stated earlier that the UI should protect the user from handling the watermarks erroneously. One possible improvement is to include a watermark strength indicator on the screen, so the user sees how the watermark reacts to image transformations. Ideally the watermark embedding is the last step before publishing the image. Therefore the software could delay the embedding until the user wants to save the image. Equally when loading the image, the watermark could be extracted and the user would always work with an unwatermarked image. The watermarking would therefore become a transparent operation to the user.

6 Conclusion

The real world presents a number of challenges to the successful use of watermarking for copyright control. Account must be taken of the

[6]Indeed, a search engine could return matches in order of probability of a correct match

natural limitations of watermarking, vulnerabilities in the specific watermarking method used, and weaknesses in the system that employs it. Even with both a secure watermarking method and a secure system that uses that method, there is no guarantee of successfully controlling copyright. A watermarking system is only as good as the legal system it is tied to; if third parties do not abide by that legal system, the copyright holder has no recourse to justice.

Often end users have unrealistic expectations of the strength of the watermarking process – they expect it to be infallible. In fact, with sufficiently small images it is impossible to guarantee that a watermark will be retained. As we have pointed out, to some extent these problems can be overcome by good user-interface design. However, watermarks should still be seen as aid to efficiently tracking copyright violations, rather than as a secure means of asserting copyright.

A real-world watermarking system should be robust in the face of image transformations, but must handle particularly common transformations (scaling and compression) best of all. The system must also be clear in its design goals: using a single watermark scheme for multiple purposes will lead to security compromises. Unfortunately the currently available commercial watermark systems fall some way short of these goals.

7 Acknowledgments

We would like to thank to Joseph Ó Ruanaidh and to Alexander Herrigel for their thoughtful feedback on an early draft and for insight they have provided in numerous discussions.

References

[AP98] Ross J. Anderson and Fabien A.P. Petitcolas. On The Limits of Steganography. In *IEEE Journal on Selected Areas in Comunications - Special Issue on Copyright & Privacy Protection*, 1998. http://www.cl.cam.ac.uk/users/fapp2/ papers/ On_The_Limits_Of_Steganography/.

[Cin95] Robert A. Cinque. Making Cyberspace Safe for Copyright: The Protection of Electronic Works in a Protocol to the Berne Convention. *18 FORDHAM INT'L L.J.*, 1995. citing Berne Convention art 7(6).

[CKLS96] I. Cox, J. Killian, T. Leighton, and T. Shamoon. Secure spread spectrum watermarking for images, audio and video. In *Proceedings of the IEEE Int. Conf. on Image Pro-*

cessing ICIP-96, pages 243–246, Lausanne, Switzerland, September 16-19 1996.

[CL97] Ingemar J. Cox and Jean-Paul M.G. Linnartz. Public watermarks and resistance to tampering. In *Proceedings of the IEEE International Conference on Image Processing*, 1997.

[CMYY96] Scott Craver, Nasir Memon, Boon-Lock Yeo, and Minerva Yeung. Can Invisible Watermarks Resolve Rightful Ownerships? Technical report, IBM Research Report, July 1996. http://www.watson.ibm.com:8080/ main-cgi-bin/search_paper.pl/ entry_ids=8214.

[Cox96] Brad Cox. *Superdistribution: objects as property on the electronic frontier*. Addison-Wesley, 1996.

[Dig97] Digimarc. Website. http://www.digimarc.com, 1997.

[HPR97] Alexander Herrigel, Adrian Perrig, and Joseph J.K. Ó Ruanaidh. A Copyright Protection Environment for Digital Images. In *VIS '97, Albert-Ludwigs Universität, Freiburg, Germany*, 1997.

[Int97] Intertrust. Website. http://www.intertrust.com, 1997.

[Kat97] Marty Katz. Digital Watermarks Often Fail on Web Images. *New York Times*, November 11th 1997.

[Kuh] Markus Kuhn. Stirmark. http://www.cl.cam.ac.uk/ ~mgk25/ download/. More information can be found at: http://www.cl.cam.ac.uk/ ~mgk25/stirmark.html.

[Mar96] Martin Kutter and Frédéric Jordan. Digital Signature of Color Images using Amplitude Modulation. Technical report, EPFL-LTS, 1996. http://ltswww.epfl.ch/ ~kutter/ publications/ amt.html.

[PAK98] Fabien A. P. Petitcolas, Ross J. Anderson, and Markus G. Kuhn. Attacks on Copyright Marking Systems. In *Second International Workshop on Information Hiding*, April 1998.

[Per97] Adrian Perrig. A Copyright Protection Environment for Digital Images. Technical report, Ecole Polytechnique Fédérale de Lausanne, EPFL - LSE , Switzerland, March 1997.

[PJ96] J. Puate and F. Jordan. Using fractal compression scheme to embed a digital signature into an image. http://ltswww.epfl.ch/ ~kutter/ publications/fvt.html, November 1996.

[RDB96] J. J. K. Ó Ruanaidh, W. J. Dowling, and F. M. Boland. Phase watermarking of digital images. In *Proceedings of the IEEE Int. Conf. on Image Processing ICIP-96*, pages 239–242, Lausanne, Switzerland, September 16-19 1996.

[RP97] J. J. K. Ó Ruanaidh and Thierry Pun. Rotation, Scale and Translation Invariant Digital Image Watermarking. Submitted to Signal Processing, February 1997.

[Rua97] Ó Ruanaidh. Phase watermarking of digital images. In *Proceedings of the IEEE Int. Conf. on Image Processing ICIP-97*, Lausanne, Switzerland, October 1997.

[Sch96] Bruce Schneier. *Applied Cryptography*. John Wiley & Sons, 1996.

[TvSO95] A.Z. Tirkel, R.G. van Schnydel, and C.F. Osborne. A TWO-DIMENSIONAL DIGITAL WATERMARK. Technical report, Scientific Technology, Australia, 1995.

[Whi98] Alma Whitten. Making Security Usable. PhD proposal at Carnegie Mellon University, January 1998. Work in progress.

[Wil98] David L. Wilson. Copyright vs. the right to copy. *San Jose Mercury News*, February 28 1998. http://www.mercurycenter.com/ business/ center/ copy030198.htm.

[Xa] Mallory X. Digimarc crack. http://www.chez.com/ pdufour/ zip/ digimarc.zip.

[Xb] Mallory X. Unzign: Is your watermark secure? http://www.altern.com/watermark.

[ZK94] J. Zhao and E. Koch. Embedding robust labels into images for copyright protection. Technical report, Fraunhofer Institute for Computer Graphics, Darmstadt, Germany, 1994.