



Chapitre d'actes

2001

Published version

Open Access

This is the published version of the publication, made available in accordance with the publisher's policy.

Second generation benchmarking and application oriented evaluation

Pereira, Shelby; Voloshynovskyy, Svyatoslav; Madueno Ruiz, Maribel Maria Isabel; Marchand-Maillet, Stéphane; Pun, Thierry

How to cite

PEREIRA, Shelby et al. Second generation benchmarking and application oriented evaluation. In: Information Hiding : 4th international workshop, IH 2001, Pittsburgh, PA, USA, April 25-27, 2001 : proceedings. Pittsburgh (USA). [s.l.] : [s.n.], 2001. p. 340–353. doi: 10.1007/3-540-45496-9_25

This publication URL: <https://archive-ouverte.unige.ch/unige:47871>

Publication DOI: [10.1007/3-540-45496-9_25](https://doi.org/10.1007/3-540-45496-9_25)

Second generation benchmarking and application oriented evaluation

S. Pereira, S. Voloshynovskiy, M. Madueno, S. Marchand-Maillet and T. Pun

University of Geneva - CUI,
24 rue General Dufour, CH 1211
Geneva 4, Switzerland,
{Shelby.Pereira,svolos,Maribel.Madueno,Thierry.Pun}@cui.unige.ch
WWW home page: <http://vision.unige.ch/>

Abstract. Digital image watermarking techniques for copyright protection have become increasingly robust. The best algorithms perform well against the now standard benchmark tests included in the Stirmark package. However the stirmark tests are limited since in general they do not properly model the watermarking process and consequently are limited in their potential to removing the best watermarks. Here we propose a second generation benchmark for image watermarking which includes attacks which take into account powerful prior information about the watermark and the watermarking algorithms. We follow the model of the Stirmark benchmark and propose the 8 following categories of tests: denoising (ML and MAP), wavelet compression, watermark copy attack, active desynchronization, denoising, geometrical attacks, and denoising followed by perceptual remodulation. The attacks in this benchmark are intended to complement those in the Stirmark benchmark with the ultimate goal of providing some standard tools for the comparison of watermarking algorithms. In addition, we take the important step of presenting results as a function of application. This is an important contribution since it is unlikely that one technology will be suitable for all applications.

1 Introduction

Digital watermarking has emerged as an appropriate tool for the protection of author's rights. It is now well accepted that an effective watermarking scheme must successfully deal with the triple requirement of *imperceptibility* (visibility) - *robustness* - *capacity* [20]. *Imperceptibility* requires that the marked data and the original data should be perceptually undistinguishable. *Robustness* refers to the fact that the embedded information should be reliably decodable after alterations of the marked data. Often the level of robustness is dictated by the application. *Capacity* requires to the amount of information that is being embedded in the watermark. In typical applications we require between 60 and 100 bits. This is necessary so as to uniquely associate images with buyers and sellers.

In addition to these requirements, the issue of algorithm complexity is also of importance. In some applications for example, it is necessary that the algorithms lend themselves to a hardware implementation. In other applications such as video watermarking, real-time embedding and detection may be essential. To further complicate the issue, the requirement on complexity may depend on the protocols used to distribute the media.

Given the relatively complex tradeoffs involved in designing a watermarking system, the question of how to perform fair comparisons between different algorithms naturally arises. A lack of systematic benchmarking of existing methods however creates confusion amongst content providers and watermarking technology suppliers. The benchmarking tool Stirmark [18, 19] integrates a number of image processing operations or geometrical transformations aimed at removing watermarks from a stego image. The design of this tool does not take into account the statistical properties of the images and watermarks in the design of attacks. As a result, pirates can design more efficient attacks that are not currently included in the benchmarking tools. This could lead to a tremendous difference between what existing benchmarks test and real world attacks. Another problem with the Stirmark benchmarking tool is that it does not take into account the fact that different applications require different levels of robustness. This is an important issue and is currently a subject of investigation within the Certimark European Project [2].

In [24], an important step was taken in designing attacks which include important priors on the image and the watermarking algorithm. The paper uses estimation based techniques to derive the optimal attacks for a given watermark distribution. Furthermore, a new method based on Watson’s metric [25] was proposed for determining visual quality of a watermark image. It is the aim of this paper to further explore such attacks, and formulate a second generation benchmark which includes a comprehensive set of attacks including important attacks which take into account embedding strategies used by many current algorithms. This paper falls within the scope of the current Certimark European project whose central aim is to provide a proper means for evaluating watermarking technologies. The aim of this benchmark is not to invalidate the benchmark already proposed by Petitcolas, but rather to present a number of attacks which have not been considered.

In addition to the new attacks, we address the important issue of weighting attacks as a function of applications. This is an important new direction for benchmarking. In the original Stirmark benchmark, scores were reported by averaging over groups of attacks. We propose to modify this scheme so that the importance of attacks is weighted according to application.

The paper is structured as follows. In section 2 we review and categorize existing attacks. In section 3 we review the second generation benchmark proposed in [24]. It is within the scope of this second generation benchmark that we will include our attacks. In sections 4 to 6 we present new removal attacks, new geometrical attacks, and new protocol attacks respectively. Finally in section 7 we present our results followed by a conclusion in section 8.

2 State-of-art Watermarking attacks

We will adopt the attack classification scheme detailed in our previous paper [24]. We review the main points here. The wide class of existing attacks can be divided into four main categories: interference and removal attacks, geometrical attacks, cryptographic attacks and protocol attacks. Figure 1 summarizes the different attacks.

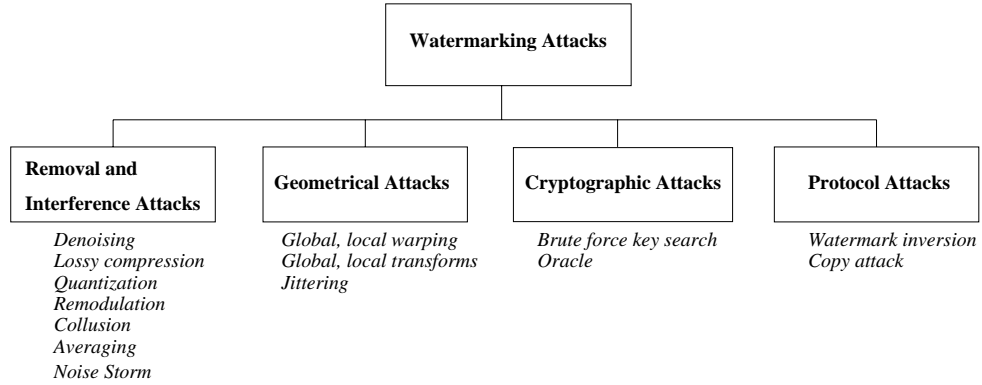


Fig. 1. Classification of watermarking attacks

2.1 Interference and Removal Attacks

In [24], interference and removal attacks are discussed in detail. The main idea consists of assuming that the watermark is additive noise relative to the original image. The interference attacks are those which further add noise to the watermarked image. This noise may have any of a number of different statistical distributions such as Gaussian or Laplacian. The removal attacks exploit the linear additive model in order to derive optimal estimators used for denoising and consequently removing of the watermark. In other cases both the removal attacks and the interference attacks can be combined such as in the denoising with perceptual remodulation attacks. These attacks are further detailed in section 4.

2.2 Geometrical attacks

In contrast to the removal attacks, geometrical attacks intend not to remove the embedded watermark itself, but to distort it through spatial alterations of the stego data. The attacks are usually such that the watermark detector loses synchronization with the embedded information. The most well know integrated

software versions of these attacks are Unzign and Stirmark. Unzign [21] introduces local pixel jittering and is very efficient in attacking spatial domain watermarking schemes. Stirmark [17] introduces both global geometrical and local distortions. The global distortions are rotation, scaling, change of aspect ratio, translation and shearing that belong to the class of general affine transformations. The line/column removal and cropping/translation are also integrated in Stirmark. Most recent watermarking methods survive after these attacks due to the usage of special synchronization techniques. Robustness to the global geometrical distortions rely on the use of either a transform invariant domain [11], or an additional template [12, 3, 13], or an Autocorrelation Function (ACF) of the watermark itself [8],[24].

If robustness to global affine transformations is a solved problem, the local random alterations integrated in Stirmark still remains an open problem almost for all techniques. The so called random bending attack exploits the fact that the human visual system is not sensitive against shifts and local affine modifications. Therefore, pixels are locally shifted, scaled and rotated without significant visual distortions.

2.3 Cryptographic attacks

Cryptographic attacks are very similar to the attacks used in cryptography. There are the brute force attacks which aim at finding secret information through an exhaustive search. Since many watermarking schemes use a secret key it is very important to use keys with a secure length. Another attack in this category is the so called Oracle attack [16, 5] which can be used to create a non-watermarked image when a watermark detector device is available.

2.4 Protocol attacks

The protocol attacks aim at attacking the concept of the watermarking application. The first protocol attack was proposed by Craver et al [6]. They introduce the framework of invertible watermark and show that for copyright protection applications watermarks need to be non-invertible. The idea of inversion consists of the fact that an attacker who has a copy of the stego data can claim that the data contains also the attacker's watermark by *subtracting* his own watermark. This can create a situation of ambiguity with respect to the real ownership of the data. The requirement of non-invertibility on the watermarking technology implies that it should not be possible to extract a watermark from non-watermarked image. As a solution to this problem, the authors propose to make watermarks signal-dependent by using a one-way function.

The copy attack [9] also belongs to the group of the protocol attacks. The goal of the attack is not to destroy the watermark or impair its detection, but consists rather in the prediction of the watermark from the cover image, like in the case of the remodulation attack, followed by copying the predicted watermark on the target data. The estimated watermark is then adapted to the local

features of the stego data to satisfy its imperceptability. The process of copying the watermark requires neither algorithmic knowledge of the watermarking technology nor the watermarking key. However, in the published version of this attack it was assumed that the watermarking algorithm exploits linear additive techniques. The derivation of the optimal MAP estimate for multiplicative watermarks or generally non-additive techniques is required to cover methods like SysCop of MediaSec [10], Barni [1] and Pereira [14, 15] that are mostly used in the transform domains.

Although the above classification makes it possible to have a clear separation between the different classes of attacks, it is necessary to note that very often a malicious attacker applies not only a single attack at the moment, but rather a combination of two or more attacks. Such a possibility is predicted in the Stirmark benchmark where practically all geometrical transformations are accompanied by lossy compression.

3 Review of Proposed Attacks Included in Initial Second Generation Benchmark

In [24], we proposed a second generation benchmark which includes a variety of estimation based attacks as well as other attacks which take into account priors about the watermarking embedding strategy. We briefly review the proposed benchmark before addressing the new attacks. While the Stirmark benchmark heavily weights geometric transformations and contains non-adaptive attacks, the benchmark we proposed includes models of the image and watermark in order to produce more adapted attacks.

3.1 Review of Benchmark Proposal

The benchmark consists of six categories of attacks where for each attacked image a 1 is assigned if the watermark is decoded and 0 if not. A detailed discussion of all attacks is provided in [24] and will not be repeated here. The categories are the following where we note in parentheses the abbreviations we use later for reporting results:

1. Denoising (DEN): We perform three types of denoising, Wiener filtering, soft thresholding and hard thresholding.
2. Denoising followed by perceptual remodulation (DPR).
3. Hard Thresholding followed by Stirmark random bending (DRB).
4. Copy Attack (CA): We estimate the watermark using Wiener filtering and copy it onto another image.
5. Template removal followed by small rotation (TR). Here, DFT peaks are removed since these are frequently used for synchronization as in [1, 3, 13].
6. Wavelet Compression (WC): In this section we compress the image using bitrates [2, 1, 0.9, 0.8, 0.7, 0.6, 0.5, 0.4, 0.3, 0.2, 0.1]. The finer sampling at low bitrates allows us to better localize at which point the algorithms break down.

In some applications bitrates in the range of 0.1-0.2 are frequently encountered. We note that this corresponds roughly to a JPEG quality factor of 10% however the artifacts are much less problematic since the blocking effects do not occur with wavelet compression.

It is within this basic framework that we wish to add several new attacks to provide a more complete evaluation of watermarking algorithms.

4 New Removal Attacks

Having reviewed the attacks contained in our second generation benchmark, we now consider a number of new attacks which are contained in the current implementation. None of the attacks which we describe in this section have been included in the Stirmark benchmarking tool. We present the attacks in accordance with the classification scheme presented in figure 1.

In this section we present 7 attacks which all fall into the category of removal attacks. First we consider 3 new Maximum Likelihood (ML) estimation based attacks, we then consider a new MAP attack which includes more powerful prior information about the watermark, and finally we apply the MAP estimators to the denoising with perceptual remodulation (DPR) attack and also consider a second possible DPR attack. Finally we add the collusion attack.

Maximum Likelihood Estimation Attacks: Here we consider three new attacks which are based on the assumption that the watermark is additive. That is, the watermarking process can be modeled by:

$$\mathbf{y} = \mathbf{x} + \mathbf{w} \quad (1)$$

where \mathbf{x} is the cover image, \mathbf{w} is the watermark, \mathbf{y} is the stego image. The ML is then given by

$$\hat{\mathbf{x}} = \arg \max_{\mathbf{x} \in \mathbb{R}^N} \{p_{\mathbf{w}}(\mathbf{y} | \mathbf{x})\}. \quad (2)$$

The Stirmark 3.1 program contains the local mean and local median filters which correspond to the case of a Gaussian and Laplacian watermark respectively. The Laplacian model is particularly useful in modelling impulsive watermarks. This is important with respect to watermarks generated by increasing a pixel to encode a 1 and decreasing it to encode a 0.

Here we propose to add the local trimmed mean filter and the midpoint filter. In the theory of robust statistics, the mixture model of the Gaussian and Laplacian distributions is used. The closed solution in this case is the local trimmed mean filter that uses order statistics such as the median filter but produces the trimmed version of the mean centered about the median point. The size of the window used for the mean computation is determined by the percentage of the impulse outliers given by parameter ϵ hence the name ϵ -contaminated. The midpoint filter corresponds to the ML estimate for a uniform distribution of the watermark. In practice the midpoint filter consists of replacing the middle point in the local window with $\frac{max-min}{2}$.

MAP based attack

We first note that the Stirmark program contains the local mean filter which corresponds to the optimal estimate of a Gaussian uncorrelated watermark. Here we wish to exploit the knowledge that almost all watermarking schemes use low pass watermarks. In fact it is now well known that in order for a watermark to be robust it must have low pass characteristics. Ultimately a compromise must be made since low pass watermarks also tend to be more visible. This compromise was first discussed by Cox [4].

Denoising assuming low pass watermark: In order to improve our estimate of the watermark in situations where the watermark is low pass, we choose a non-stationary Gaussian model for the image $x \sim N(\bar{x}, R_x)$ with local mean \bar{x} and covariance matrix R_x , and a Gaussian model for the watermark noise with correlation $n \sim N(0, R_n)$. Assuming image and noise are conditionally i.i.d. one obtains the well known Wiener filter as the solution. We note that in the case of watermark removal, the noise covariance matrix will typically only include correlation at a distance of at most 3 pixels.

Denoising followed by perceptual remodulation: Voloshynovskiy [23] proposed a generalized two stage attack based on denoising/compression and on spatial watermark prediction using an MAP estimate of the watermark followed by perceptual remodulation to create the least favorable noise distribution for the watermark decoder. Here we propose two improvements on the basic attack in which in both cases we add prior knowledge about the watermarking embedding strategy.

The first attack consists of first applying the MAP estimate from the previous section which includes the assumption of a correlated watermark and then applying perceptual remodulation. We recall from [24] that perceptual remodulation consists of flipping the estimated sign and adding it back to the image after applying perceptual masking. We also recall that we do not modify all the pixels, but typically only a certain percentage (at least 30% in practice).

The second attack in this category consists of exploiting the weakness of correlated watermarks. Rather than remodulating isolated pixels, we remodulate local groups of pixels together. If indeed the watermark was low pass, this attack should be much more effective since the remodulation is low pass.

5 New Geometrical Attacks

The current Stirmark benchmark contains a wide range of geometrical attacks. Here we propose the following new attacks: projective transforms, non-uniform line removal, and collage attacks [8], as well as an attack on periodical watermarks.

Projective Transforms: The geometrical transformations included in the Stirmark program operate in a two dimensional space. However, modern image processing software, such as Corel Draw or Adobe Photoshop, typically include a variety of three dimensional transformations. Consequently we propose to add

projective transformations to the benchmark. The image is considered to be in a three dimensional space. Geometrical transformations including rotations, shearing, scale changes and other general transformations are applied in 3D. Finally the image is projected back onto a two dimensional space. This latter step is typically done via a perspective or parallel projection. A comprehensive review of projective geometry is given by Faugeras [7].

Collage Attacks: In this section we consider the collage attack proposed by Kutter [8]. We create a collage of several portions of images in which one portion contains a watermark. A successful watermark detector should be able to recover the watermark. In addition we consider situations where the watermarked image is rotated and/or scaled prior to being included in a collage.

Non-uniform Line Removal: In the Stirmark benchmark, the line removal attack is included however only periodical removal of lines is considered. Unfortunately this has the drawback that the resulting distortion can be approximated by a rescaling. Here we propose to add non-uniform line removal. That is lines are removed at different intervals both in the horizontal and vertical directions. In this case, the distortion can no longer be modelled as a rescaling.

6 Protocol Attacks

We include in this section the copy attack proposed by Kutter [9]. The basic idea is to estimate the watermark from one image and then adding it to another image after applying perceptual masking. Wiener filtering is used to estimate the watermark.

7 Results

The new benchmarking tool is called Checkmark and is available from our website <http://watermarking.unige.ch/>. The program is written in Matlab and calls the Stirmark program version 3.1.79 in order to generate the original attacks proposed by Petitcolas. In addition the package contains an XML description of the relationship of a given application to a set of attacks. Consequently, rather than proposing an overall score, it is easy to generate the results as a function of application. The use of XML yields a flexible way of adding new attacks and applications and weighting the results as a function of an application. Another the advantage of XML is that the results can then be easily parsed and easily converted from one format to another.

In order to generate the results, we provide sample scripts for several detectors. The process consists of calling the detector and generating an XML file for the results of the detector relative to all attacks. Also included in this XML file are the quality metrics of the resulting watermarked images as compared to the originals. In order to generate a final version of the results, an XSL style sheet is used to parse the XML and generate HTML web pages where the results are organized. Results obtained for various technologies as a function of application are also displayed at our website. A complete description of all parameter setting

is included in the Checkmark software package and will not be repeated here due to lack of space. Our main purpose here is to outline the general approach taken in classifying attacks. All results are displayed at the website, and all relevant files have been made available.

In the first version of the program we have considered the following applications: general copyright protection and banknote protection. In the case of general copyright protection all attacks are applicable however it is clear that not all are equally important. For example it is not that important to be able to resist quality factor of 10 with respect to JPEG. Such low quality compression degrades the image and in most applications it is not important. As a preliminary step in assigning weights, we first categorize compression levels as low, medium, and high both for the wavelet and JPEG compression levels. With respect to JPEG low quality is between 10-30, medium between 40-60 and high from 70-100. A similar weighting is made with respect to wavelet compression. For the application of Copyright protection, the attacks in the low compression category are given a weight of only half that of the medium and high compression rates.

Similarly for geometric attacks we classify the geometric changes as slight, medium and large. Slight changes are those which are more or less invisible. For example a rotation of 1 degree or a scale change of 10% would fall into this category. The bulk of the geometric changes are classified as medium while changes such as a scaling of 200% or 50% are classified as large. For the copyright protection application, slight and medium changes receive the highest weightings and count for twice the weighting of the large changes which would typically be less frequently encountered in most applications.

In the case of the denoising attacks, the same categorization into three groups is made. Here we observe that larger window sizes produce larger distortions, and typically more blurring. Consequently the attacks are classified as a function of distortion as low, medium and high corresponding to window sizes of 3, 5, and 7. In order to produce a final score for the application, we first perform the weighted averages within each category and then perform the overall average over the categories. We do note however that in analyzing results, we should still look at each category of attacks, since one number does not adequately describe the strengths and weaknesses of a technology with respect to an application.

The second application considered is banknote protection. The scenario for banknote protection is described in [22] where some specialized attacks are described. The central idea is to have a watermark detector embedded in all devices (cameras, photocopiers and scanner) to prevent people from scanning in a banknote and copying it. In this context only a small subset of attacks are applicable. All attacks which are applicable have the same weighting. We consider the applicable attacks below:

- Geometrical attacks: All geometric attacks which do not substantially modify the size of the image are applicable. For example a person might rotate the banknote prior to scanning. We also note that the projective transformations are important since this corresponds to a person displaying the banknote at an inclination to a camera. Even the random bending attack is important

here since it can be used to model the distortions associated with cheaper web cameras.

- copy attack: This attack is relevant since if the watermarked banknote is successfully scanned the watermark may be copied to other banknotes.
- cropping: This is extremely important as described in [22] since the attacker can cut the banknote into pieces prior to scanning and then recombine the digitized image.
- lowpass filtering and denoising attacks: These attacks tend to blur the image and can be used to model the blurring associated with a defocused camera as described in [22].

At this time only two applications have been considered. The website will be updated as new applications and attacks are included. Furthermore we welcome submission of XML result files generated from the Checkmark tool. This will allow us to generate a centralized repository of results for easy access which will surely help the research community by facilitating the task of determining the strenghts and weaknesses of algorithms.

8 Conclusion

In this article we have added a number of new attacks to the ones contained in the second generation benchmark proposed in [24]. Better understanding of the mechanisms of possible attacks will lead to the development of more efficient and robust watermarking techniques and as such our results present an important step in this direction. Furthermore as active participants in the current Certimark European project, the main purpose of this paper is to present a new benchmarking tool which can be used as an evaluation tool for image watermarking algorithms. We have also addressed the issue of application oriented benchmarking. A myriad of applications have appeared in watermarking and it is now clear that all applications have there own requirements. As a first step we have considered the applications of copyright protection and banknote protection and generated XML files which describe these applications. Work is currently under way to add other applications as well as new attacks. The use of the XML interface greatly facilitates this task and is one of the main contributions of this work.

ACKNOWLEDGMENTS

We thank Frederic Deguillaume and Alexander Herrigel for many fruitful discussions. This work has been partly financed by the European CERTIMARK project which deals with benchmarking of watermarking algorithms. This work has also been financed by DCT-Digital Copyright Technologies, Switzerland.

References

1. M. Barni, F. Bartolini, A. De Rosa, and A. Piva. A new decoder for the optimum recovery of non-additive watermarks. *IEEE Transactions on Image Processing*, submitted 2000.
2. Certimark european project, 2000-2002.
3. Digimarc Corporation. <http://www.digimarc.com/>. January 1997.
4. 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 Processing ICIP-96*, pages 243–246, Lausanne, Switzerland, 1996.
5. I. J. Cox and J.-P. M. G. Linnartz. Some general methods for tampering with watermarks. *IEEE Journal on Selected Areas in Communications*, 16(4):587–593, May 1998.
6. S. Craver, N. Memon, B. L. Yeo, and M. M. Yeung. Can invisible watermark resolve rightful ownerships? In *Fifth Conference on Storage and Retrieval for Image and Video Database*, volume 3022, pages 310–321, San Jose, CA, USA, February 1997.
7. O. Faugeras. *Three-Dimensional Computer Vision*. The MIT Press, Cambridge Massachusetts, 1993.
8. M. Kutter. *Digital image watermarking: hiding information in images*. PhD thesis, EPFL, Lausanne, Switzerland, August 1999.
9. M. Kutter, S. Voloshynovskiy, and A. Herrigel. Watermark copy attack. In Ping Wah Wong and Edward J. Delp, editors, *IS&T/SPIE's 12th Annual Symposium, Electronic Imaging 2000: Security and Watermarking of Multimedia Content II*, volume 3971 of *SPIE Proceedings*, San Jose, California USA, 23–28 jan 2000.
10. MediaSec. <http://www.mediasec.com/products/download/>. March 2000.
11. J. Ó Ruanaidh and T. Pun. Rotation, scale and translation invariant spread spectrum digital image watermarking. *Signal Processing*, 66(3):303–317, 1998.
12. S. Pereira, J. J. K. Ó Ruanaidh, F. Deguillaume, G. Csurka, and T. Pun. Template based recovery of Fourier-based watermarks using Log-polar and Log-log maps. In *Int. Conference on Multimedia Computing and Systems, Special Session on Multimedia Data Security and Watermarking*, June 1999.
13. S. Pereira and T. Pun. Fast robust template matching for affine resistant watermarks. In *3rd International Information Hiding Workshop*, Dresden, Germany, September 1999.
14. S. Pereira, S. Voloshynovskiy, and T. Pun. Effective channel coding for DCT watermarks. In *International Conference on Image Processing (ICIP'2000)*, Vancouver, Canada, September 2000.
15. Shelby Pereira and Thierry Pun. A framework for optimal adaptive DCT watermarks using linear programming. In *Tenth European Signal Processing Conference (EUSIPCO'2000)*, Tampere, Finland, sep 5–8 2000.
16. A. Perrig. A copyright protection environment for digital images. Diploma dissertation. *Ecole Polytechnique Federal de Lausanne, Lausanne, Switzerland*, February, 1997.
17. F. A. P. Petitcolas. <http://www.cl.cam.ac.uk/fapp2/watermarking/stirmark/>. In *Stirmark3.1(79)*, 1999.
18. F. A. P. Petitcolas and R. J. Anderson. Attacks on copyright marking systems. In *2nd International Information Hiding Workshop*, pages 219–239, Portland, Oregon, USA, April 1998.
19. F. A. P. Petitcolas and R. J. Anderson. Evaluation of copyright marking systems. In *IEEE Multimedia Systems (ICMCS'99)*, volume 1, pages 574–579, Florence, Italy, June 1999.

20. C. I. Podilchuk and W. Zeng. Perceptual watermarking of still images. In *Proc. Electronic Imaging*, volume 3016, San Jose, CA, USA, February 1996.
21. Unzign watermark removal software, July 1997.
22. S. Voloshynovskiy, A. Herrigel, and T. Pun. Blur/deblur attack against document protection systems based on digital watermarking. In *Information Hiding Workshop 2001*, Pittsburg, USA, April 2001.
23. S. Voloshynovskiy, S. Pereira, A. Herrigel, N. Baumgärtner, and T. Pun. A generalized watermark attack based on stochastic watermark estimation and perceptual remodulation. In Ping Wah Wong and Edward J. Delp, editors, *IS&T/SPIE's 12th Annual Symposium, Electronic Imaging 2000: Security and Watermarking of Multimedia Content II*, volume 3971 of *SPIE Proceedings*, San Jose, California USA, 23–28 January 2000. (Paper EI 3971-34).
24. S. Voloshynovskiy, S. Pereira, V. Iquise, and T. Pun. Attack modelling: Towards a second generation watermarking benchmark. *Signal Processing*, June 2001.
25. A. B. Watson. DCT quantization matrices visually optimized for individual images. In *Proc. SPIE: Human vision, Visual Processing and Digital Display IV*, volume 1913, pages 202–216. SPIE, 1993.