

UNIVERSITÉ DE NICE-SOPHIA ANTIPOLIS – UFR Sciences  
École Doctorale des Sciences et Technologies de l'Information et de la Communication

## THÈSE

pour obtenir le titre de  
**Docteur en Sciences**  
de l'Université de Nice-Sophia Antipolis

Discipline: Automatique Traitement du Signal et des Images

présentée et soutenue par  
**Gwenaël DOËRR**

## SECURITY ISSUE AND COLLUSION ATTACKS IN VIDEO WATERMARKING

Thèse dirigée par Jean-Luc DUGELAY  
soutenue le 10 Juin 2005

### Jury:

Prof.	Pierre DUHAMEL	Supélec	Président
Prof.	Ingemar COX	University College of London	Rapporteur
Prof.	Mauro BARNI	Università di Siena	Rapporteur
Dr.	Darko KIROVSKI	Microsoft Research	Examineur
Prof.	Jean-Luc DUGELAY	Institut Eurécom	Examineur



SECURITY ISSUE AND COLLUSION  
ATTACKS IN VIDEO WATERMARKING

GWENAËL DOËRR

10 JUIN 2005



# CONTENTS

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Digital Watermarking . . . . .	1
1.1.1	Blind Watermarking . . . . .	3
1.1.2	Informed Watermarking . . . . .	5
1.1.3	Further Reading . . . . .	6
1.2	Outline of the Thesis . . . . .	6
<hr/>		
<b>I</b>	<b>A State-of-the-Art Overview</b>	<b>9</b>
<b>2</b>	<b>Applications</b>	<b>11</b>
2.1	Steganography . . . . .	11
2.2	Data Hiding . . . . .	13
2.2.1	Labelling for Data Retrieval . . . . .	13
2.2.2	Data Hiding for Data Compression . . . . .	14
2.2.3	Data Hiding for Error Recovery . . . . .	14
2.3	IPR Protection . . . . .	15
2.3.1	Proof of Ownership . . . . .	16
2.3.2	Access Control . . . . .	16
2.3.3	Broadcast Monitoring . . . . .	19
2.3.4	Fingerprinting . . . . .	20
2.3.5	Authentication . . . . .	22
<b>3</b>	<b>Major Trends</b>	<b>25</b>
3.1	From Still Images to Video Watermarking . . . . .	26
3.2	Integration of the Temporal Dimension . . . . .	29
3.2.1	Video as a Monodimensional Signal . . . . .	29
3.2.2	Video as a Temporal Signal . . . . .	31
3.2.3	Video as a 3D Signal . . . . .	32

3.3	Exploiting the Video Compression Format . . . . .	32
3.3.1	Modify Coefficients in the Transformed Domain . . . . .	33
3.3.2	Modify Motion Information . . . . .	34
3.3.3	Modify VLC Codewords . . . . .	35
<b>4</b>	<b>Challenges</b>	<b>39</b>
4.1	Various Non Hostile Video Processing Primitives . . . . .	39
4.2	Temporal Synchronization . . . . .	42
4.2.1	Key Scheduling . . . . .	43
4.2.2	Frame Dependent Watermarking . . . . .	44
4.3	Distortion Evaluation . . . . .	45
4.4	Real-time Watermarking . . . . .	47
4.4.1	Low-Complexity Algorithms . . . . .	48
4.4.2	Reduce the Rate of Expensive Computations . . . . .	49
4.4.3	Preprocessing to Facilitate Watermark Embedding . . . . .	49
4.5	Security Issue and Collusion Attacks . . . . .	51
4.5.1	Trust in a Hostile Environment . . . . .	51
4.5.2	Security versus Robustness . . . . .	53
4.5.3	Security in the Real World . . . . .	55
4.5.4	Collusion Attacks . . . . .	56
<hr/>		
<b>II</b>	<b>Security Issue and Collusion Attacks</b>	<b>59</b>
<b>5</b>	<b>Eavesdropping the Watermarking Channel</b>	<b>61</b>
5.1	Baseline Framework . . . . .	61
5.1.1	Frame-by-Frame Watermarking . . . . .	62
5.1.2	Weaknesses Against Simple Collusion Attacks . . . . .	64
5.2	Switching Between Orthogonal Watermarks . . . . .	67
5.2.1	SS-N System . . . . .	67
5.2.2	Enhanced Security . . . . .	68
5.2.3	Experimental Results . . . . .	69
5.2.4	Potential Weaknesses . . . . .	70
5.3	Embedding Strength Modulation . . . . .	75
5.3.1	Novel Perspective . . . . .	75
5.3.2	SS- $\alpha$ System . . . . .	77
5.3.3	Watermarking Subspace Estimation Draining (WSED) . . . . .	80
5.4	Scalar Costa Schemes (SCS) . . . . .	82
5.5	Discussion . . . . .	84

<b>6</b>	<b>Jamming the Watermarking Channel</b>	<b>87</b>
6.1	Temporal Frame Averaging after Registration . . . . .	88
6.1.1	Attack Description . . . . .	88
6.1.2	TFAR Evaluation . . . . .	90
6.2	Block Replacement Attack . . . . .	92
6.2.1	Block Restoration . . . . .	92
6.2.2	Block Swapping . . . . .	93
6.2.3	Blocks Combination . . . . .	95
6.2.4	Block Projection . . . . .	99
6.2.5	BRA Evaluation . . . . .	101
6.3	Discussion . . . . .	105

---

### **III Signal Coherent Watermarking** **107**

<b>7</b>	<b>Motion Compensated Watermarking</b>	<b>109</b>
7.1	Watermarking Exploiting Video Mosaicing (SS-Reg) . . . . .	110
7.1.1	Watermark Embedding . . . . .	110
7.1.2	Watermark Detection . . . . .	112
7.2	System Analysis . . . . .	113
7.2.1	Enhanced Security . . . . .	113
7.2.2	Video Capacity . . . . .	114
7.2.3	Watermark Visibility . . . . .	115
7.3	Further Studies . . . . .	116
7.3.1	Detection Performances Improvement . . . . .	116
7.3.2	Local Resynchronization . . . . .	118
7.4	Discussion . . . . .	124
<b>8</b>	<b>Similarities Inheritance</b>	<b>127</b>
8.1	Linear Watermarking with Neighborhood Characteristics . . . . .	127
8.2	A Practical Implementation Using Gabor Features . . . . .	131
8.3	Relationship with Multiplicative Watermarking Schemes in the Fre- quency Domain . . . . .	134
8.4	Investigations . . . . .	136
8.4.1	Protocol . . . . .	136
8.4.2	Influence of the Number of GEFs . . . . .	137
8.4.3	Signal Coherent Watermarks Comparison . . . . .	139
8.5	Discussion . . . . .	140

---

<b>9 Conclusion</b>	<b>141</b>
9.1 Summary and Contributions . . . . .	141
9.2 Tracks for Future Work . . . . .	145
<b>A Résumé Français (Version Longue)</b>	<b>147</b>
A.1 Introduction . . . . .	147
A.2 Problématique de la sécurité . . . . .	149
A.2.1 Confiance dans un environnement hostile . . . . .	149
A.2.2 Robustesse et Sécurité . . . . .	150
A.2.3 Attaques par collusion . . . . .	152
A.3 Collusion en vidéo . . . . .	154
A.3.1 Estimer une structure redondante . . . . .	154
A.3.2 Combiner différents tatouages . . . . .	157
A.4 Tatouage cohérent avec le signal . . . . .	159
A.4.1 Gérer le mouvement de la caméra . . . . .	160
A.4.2 Hériter des autosimilarités . . . . .	162
A.5 Conclusion . . . . .	163
<b>B Euremark</b>	<b>165</b>
B.1 Watermark embedding . . . . .	165
B.1.1 Cover generation . . . . .	165
B.1.2 Watermark formatting . . . . .	166
B.1.3 Modulation . . . . .	167
B.2 Watermark extraction . . . . .	167
B.2.1 Cover extraction . . . . .	167
B.2.2 Payload extraction . . . . .	168



# LIST OF FIGURES

1.1	Number of publications on INSPEC. . . . .	2
1.2	Generic watermarking scheme. . . . .	3
1.3	Trade-off in digital watermarking. . . . .	4
1.4	Informed watermarking scheme. . . . .	6
2.1	DVD copy-protection system. . . . .	18
2.2	Alternative watermarking strategies for video streaming. . . . .	21
2.3	Original and tampered video frames. . . . .	23
3.1	JAWS embedding procedure. . . . .	27
3.2	Example of SPOMF detection. . . . .	28
3.3	Line scan of a video stream. . . . .	30
3.4	Independent Component Analysis of the <i>piano</i> sequence. . . . .	31
3.5	DEW embedding procedure. . . . .	34
3.6	VLC embedding process. . . . .	37
4.1	Example of distortion created by a handled camera (exaggerated). . . . .	40
4.2	Alternative key schedules. . . . .	44
4.3	Human Visual System and spatiotemporal signals. . . . .	46
4.4	Embedding strategies with MPEG video streams. . . . .	48
4.5	Geometric interpretation of different embedding strategies. . . . .	50
4.6	Watermarking attacks breakdown. . . . .	54
4.7	Collusion in watermarking. . . . .	57
5.1	SS system. . . . .	62
5.2	SS-1 system. . . . .	64
5.3	Visual illustration of basic collusion attacks. . . . .	65
5.4	SS- $N$ system. . . . .	68
5.5	Performances against basic collusion attacks. . . . .	71
5.6	Watermarks distribution on the unit sphere. . . . .	76
5.7	SS- $\alpha$ system. . . . .	79
5.8	SCS probability density function. . . . .	83

6.1	Temporal Frame Averaging after Registration (TFAR). . . . .	89
6.2	Impact of TFAR on SS and SS-1 systems. . . . .	91
6.3	Error concealment techniques. . . . .	93
6.4	Block swapping attack. . . . .	95
6.5	Influence of the number of blocks used for combination. . . . .	96
6.6	Performances of BRA. . . . .	104
7.1	Embedding procedure for camera motion coherent watermarking. . . . .	111
7.2	Impact of TFAR on the SS-Reg system. . . . .	114
7.3	Impact of the chip rate on detection performances. . . . .	117
7.4	Alternative strategies for watermark resampling. . . . .	118
7.5	Estimated optical flow with BM and EGM. . . . .	120
7.6	Influence of the multi-scales framework on the optical flow. . . . .	122
7.7	Impact of TFAR on the SS-Reg system. . . . .	123
7.8	3D object watermarking through texture mapping. . . . .	125
8.1	Geometrical point of view of the linear form $\varphi$ . . . . .	129
8.2	Evolution of $f_{\varphi \mu}$ when the feature space dimension increases. . . . .	130
8.3	GEF pair visualization. . . . .	132
8.4	Graphical representation of the Gabor filter bank. . . . .	133
8.5	Relationship between FFT and DCT. . . . .	135
8.6	Number of GEFs vs. performances against BRA. . . . .	138
8.7	Impact of BRA on SS, Gabor, FFT and DCT systems. . . . .	139
A.1	Classification robustesse/sécurité. . . . .	151
A.2	Collusion en tatouage numérique. . . . .	153
A.3	Attaque par Estimation du Tatouage et Remodulation (ETR). . . . .	156
A.4	Moyennage Temporel après Recalage (MTR). . . . .	159
A.5	Attaque par Remplacement de Bloc (ARB). . . . .	160
A.6	Tatouage cohérent avec le mouvement de la caméra. . . . .	161
B.1	Self-similarities. . . . .	166

# LIST OF TABLES

2.1	Video watermarking applications. . . . .	12
3.1	Pros and cons of the different approaches to video watermarking.	25
3.2	Example of lc-VLCs. . . . .	36
4.1	Examples of non-hostile video processing primitives. . . . .	42
5.1	Description of the videos used for experiments. . . . .	70
5.2	Impact of WECR on the SS- $N$ system. . . . .	75
6.1	Block swapping attack. . . . .	94
6.2	Fixed number of blocks combination attack. . . . .	97
6.3	Adaptive number of blocks combination attack. . . . .	98
6.4	Block projection on a PCA-defined subspace attack. . . . .	101
6.5	Luminance quantization matrix used in JPEG. . . . .	103



# ACKNOWLEDGEMENTS

To begin with, I would like to thank all the members of my jury. It is indeed a great honor for me to have such worldwide recognized experts in my jury. First of all, I am really thankful to Professor Duhamel from Supélec (France) to have accepted to be the president of this jury. Next, I would like to acknowledge the efforts from both reviewers, Professor Ingemar Cox from University College London (UK) and Professor Mauro Barni from Università di Siena (Italy), whose comments on the initial manuscript of this thesis have enabled to significantly enhance its clarity and quality. In particular, I would like to thank Professor Cox, and also Mr. Matt Miller, for transmitting to me their taste for good and innovative research during my six month internship at the NEC Research Institute, Inc. (USA). In fact, an important part of my thesis was funded by a scholarship of the French Ministry of Research, which I was granted because of the quality of the work I have done at NEC. Finally, I am also grateful to the last, but not the least, member of the jury: Dr. Darko Kirovski from Microsoft Research (USA).

Next, I would like to thank my Ph.D. supervisor Professor Jean-Luc Dugelay. He accepted to host me in his laboratory within the Eurécom Institute after I have completed my work on trellis dirty paper watermarks during my internship in Princeton (USA). Furthermore, he provided additional funding during my Ph.D. thesis and has always been available to discuss the different issues which showed up, even late in the evening. At this point, I would also like to thank all the people which have been involved more or less closely with my work. In particular, I thank Professor Henri Nicolas from IRISA (France) for his great help on motion compensation and video mosaicing. Once again, I thank Dr. Darko Kirovski as well as my summer trainee Mr. Lucas Grangé for our fruitful collaboration on Block Replacement Attacks. I am also thankful to Dr. Teddy Furon from IRISA (France) for exchanging our points of view on security in digital watermarking and collusion attacks. Finally, I want to thank all the researchers that I have met during conferences: your comments and feedback have been key elements to keep this thesis going forward. I also enjoyed your enthusiasm during *after-session hours*. In particular, I really have good memories of the karaoke contest in Seoul during IWDW 2003.

Now is time to thank all my fellow students and colleagues. First on the list, the watermarkers at the Eurécom Institute: Dr. Stéphane Roche and his fractal based watermarking system, Dr. Christian Rey and his blind registration algorithm, Dr. Emmanuel Garcia and his protection of 3D objects through texture watermarking and the fresh new Jihane Bennour working on 3D mesh watermarking. Next, I want to thank other people who happened to share my office during those three years and who had to cope with my mood of the day: Dr. Florent Perronnin, Caroline Mallauran and even the Pakistanis terrorist Dr. Ejaz Khan. I am also very thankful to the other staff members in the multimedia communications department. Many of your jokes enlightened my bad days where nothing was running fine. Some of them belong to the *old team*: Dr. Ana Christina Andrés del Valle and her Spanish enthusiasm, Dr. Ithéri Yahiaoui, Dr. Philippe de Cuetos and Dr. Kyung Tak Lee. I really appreciated how you welcomed me. There are also a few ones who joined the group at almost the same time as me: Fabrice Souvannavong and his insane but unforgettable expeditions in snow rackets (best wishes for your upcoming wedding) and the never really happy Italian Fabio Valente. Recently, a whole bunch of new Ph.D. candidates has arrived: Luca Brayda, Eric Galmar, Federico Matta, Olivier Villon, Vivek Tyagi and the last two ones, Teodora Erbes and Joakim Jiten, who share the same weird passion for breaking their leg (or something else) to adopt some kind of silly walk. Best luck for your future research. Finally, needless to say that without Benoit Huet and his twisted spirit, many coffee breaks would have been less funny.

Even if the multimedia communications department has been isolated in an auxiliary building, I have to admit that there are still some nice people in the main one. I have been waiting for three years now that the Moroccan team with Younes Souilmi, Issam Toufik and Hicham Anouar sets up some kind of trekking expedition in the desert and/or the Atlas. I enjoyed the discussion about nothing and everything in the Austin Power fashioned forum with Maxime Guillaud, Mari Kobayashi and Dr. Navid Nikaein. Finally I would like to send a dedication to the key hidden staff people. First, the secretaries Christine Mangiapan, Dephine Gouaty and Anne Duflos for taking care of all the administrative stuff in our everyday life. Next, the librarian Laurence Porte for always finding my requested references. And of course the IT team without whom, whatever has been said, there would be neither an acceptable network setup, nor a decent computing cluster.

My social life may have been somewhat limited to the Eurécom Insitute for the last three years but it would not be fair to neglect the support given by many friends scattered all over the country. First to come are Fanny Guillemain and her twin sisters, François Daoust and Pierre Dejoue for the pleasant canoe rides during summer vacation. I am also thankful to Vincent Fallas and Julien Brisset for those few *hot* nights in Paris. I also thank all the people from “Les Doctoriales

2004". It was a great week in Fréjus despite the fact that Nàiade never came to life. And of course, I have to finally acknowledge all my family relatives for their never-ending support.

Well I am sure that I am forgetting some of you folks. But I have to stop those acknowledgements at some point and it sounds to be a good point to stop. Therefore, whoever I have forgotten, do not take it personally... I also thank you!





# ABSTRACT

Ten years after its infancy, digital watermarking is still considered as a *young technology*. Despite the fact that it has been introduced for security-related applications such as copyright protection, almost no study has been conducted to assert the survival of embedded watermarks in a hostile environment. In this thesis, it will be shown that this lack of evaluation has led to critical security pitfalls against statistical analysis, also referred to as collusion attacks. Such attacks typically consider several watermarked documents and combine them to produce unwatermarked content. This threat is all the more relevant when digital video is considered since each individual video frame can be regarded as a single watermarked document by itself. Next, several countermeasures are introduced to combat the highlighted weaknesses. In particular, motion compensated watermarking and signal coherent watermarking will be investigated to produce watermarks which exhibit the same spatio-temporal self-similarities as the host video signal.

**Keywords:** Digital watermarking, digital video, security, collusion, motion compensated watermarking, signal coherent watermarking.



# RÉSUMÉ

Dix ans après son apparition, le tatouage numérique est encore considéré comme une *technologie jeune*. En dépit du fait qu'il ait été introduit pour des applications ayant trait à la sécurité telles que la protection de droit d'auteur, quasiment aucune étude n'a été conduite afin d'évaluer la survie des tatouages insérés dans un environnement hostile. Dans cette thèse, il sera montré qu'un tel manque d'évaluation a aboutit à de graves failles de sécurité face à des analyses statistiques, que l'on appelle aussi des attaques par collusion. De telles attaques considèrent typiquement plusieurs documents tatoués et les combinent afin de produire des contenus non-tatoués. Ce danger est d'autant plus important lorsque de la vidéo numérique est considérée du fait que chaque trame vidéo peut être vue individuellement comme un document tatoué. Ensuite, différentes ripostes sont introduites pour combattre les faiblesses préalablement isolées. En particulier, le tatouage compensant le mouvement et le tatouage cohérent avec le signal seront étudiés afin d'obtenir un signal de tatouage qui présente les mêmes autosimilarités spatio-temporelles que le signal vidéo hôte<sup>1</sup>.

**Mots clefs:** Tatouage numérique, vidéo numérique, sécurité, collusion, tatouage compensant le mouvement, tatouage cohérent avec le signal.

---

<sup>1</sup>Une version longue du résumé en français de cette thèse est disponible en Annexe A.



---


## Introduction

---

If you hold a common banknote up to the light, a watermarked drawing appears. This watermark is invisible during normal use and carries some information about the object in which it is embedded. The watermarks of two different kind of banknotes are indeed different. This watermark is directly inserted into the paper during the papermaking process. This very old technique is known to prevent common methods of counterfeiting. In the past few years, the use and distribution of digital multimedia data has exploded. Because it appeared that traditional protection mechanisms were no longer sufficient, content owners requested new means for copyright protection. The previous paper watermark philosophy has been transposed to digital data. Digital watermarking, the art of hiding information in a robust and invisible manner, was born. The recent interest regarding digital watermarking is demonstrated in Figure 1.1, which reports the increasing number of scientific papers dealing with this topic. Today, entire scientific conferences are dedicated to digital watermarking e.g. “SPIE: Security, Steganography and Watermarking of Multimedia Content”. Moreover, even if it is a relatively new technology, some industries have already commercialized watermarking products e.g. the widespread Digimarc.

### 1.1 Digital Watermarking

The end of the previous millennium has seen the transition from the analog to the digital world. Nowadays, audio CDs, Internet and DVDs are more and more widespread. However film and music content owners are still reluctant to release



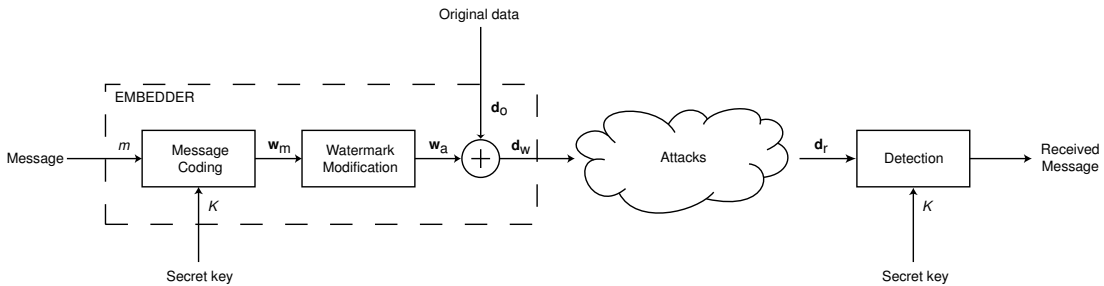
**Figure 1.1:** Number of publications registered within the INSPEC database which contain some specific keywords, February 2004.

digital content. This is mainly due to the fact that if digital content is left unprotected, it can be copied rapidly, perfectly, at large scale, without any limitation on the number of copies and distributed easily e.g. via Internet. Protection of digital content has relied for a long time on encryption but it appeared that encryption alone is not sufficient enough to protect digital data all along its lifetime. Sooner or later, digital content has to be decrypted to be eventually presented to the human consumer. At this very moment, the protection offered by encryption no longer exists and a user may duplicate and/or manipulate it.

Digital watermarking has consequently been introduced as a complementary protection technology. The basic idea consists in hiding information imperceptibly into digital content. This watermarked signal should survive most common signal processing primitives and even malicious ones if possible. The hidden information is inherently tied to digital content and protects it when encryption has disappeared. It is important to understand that digital watermarking does not replace encryption. They are two complementary techniques. On one hand, encryption prevents an unauthorized user from accessing digital content in clear during its transport. On the other hand, digital watermarking leaves an underlying invisible piece of evidence in digital data if a user, who had access to the data in clear after decryption, starts using digital data illegally (reproduction,

alteration).

Depending on what information is available during the extraction process, two separate classes of watermark detectors have been defined. If the detector has access to the original data additionally to the watermarked data, the watermark detector is called non-blind. However this kind of algorithm is less and less represented nowadays. Keeping an original version of each released digital data is indeed a very strong constraint for digital content owners in terms of storage capacity. As a result, most of the watermark detectors are actually considered as blind: the detector has only access to the watermarked data in order to extract the hidden message.




**Figure 1.2:** Generic watermarking scheme.

### 1.1.1 Blind Watermarking

Figure 1.2 depicts a simple watermarking scheme with blind detection. The goal is to embed the message  $m$  into some original data  $\mathbf{d}_o$ . The first step consists in encoding the message to be hidden with a secret key  $K$ . Typically the message is over sampled in order to match the dimension of the original data and is XORed with a pseudo-random noise generated thanks to a pseudo-random number generator which takes the secret key  $K$  as an input seed. Next, the generated watermark signal  $\mathbf{w}_m$  is modified e.g. it is scaled by a given watermarking strength. The final step simply adds the obtained watermark  $\mathbf{w}_a$  to the original data in order to obtain the watermarked data  $\mathbf{d}_w$ . This watermark embedding could be performed in whatever desired domain (spatial, Fast Fourier Transform (FFT), Discrete Cosine Transform (DCT), Fourier-Mellin). Watermarked data is then transmitted and is likely to be submitted to various signal processing operations (lossy compression, noise addition, filtering) which can be seen as attacks altering the watermark signal. If at some moment, someone wants to check if a watermark has been embedded with the secret key  $K$  in some received digital data  $\mathbf{d}_r$ , the data is simply sent through a detector. The majority of the existing detection algorithms can be seen as the computation of a correlation score between received data  $\mathbf{d}_r$  and the generated watermark  $\mathbf{w}_m$ . This correlation score

is then compared to a threshold in order to assert the presence of the watermark or not.



**Figure 1.3:** Trade-off in digital watermarking.

There exists a complex trade-off in digital watermarking between three parameters: data payload, fidelity and robustness. It is illustrated in Figure 1.3 and further presented below.

**Payload.** Data payload can be defined by the number of bits that can be hidden in digital data, which is inherently tied to the number of alternative messages that can be embedded thanks to the watermarking algorithm. It should be noted that, most of the time, data payload depends on the size of the host data. The more host samples are available, the more bits can be hidden. The capacity is consequently often given in terms of bits per sample.

**Fidelity.** Watermarking digital content can be seen as an insertion of some watermark signal in the original content and this signal is bound to introduce some distortion. As in lossy compression, one of the requirements in digital watermarking is that this distortion should remain imperceptible. In other words, a human observer should not be able to detect if some digital data has been watermarked or not. The watermarking process should not introduce suspicious perceptible artifacts. The fidelity can also be seen as the perceptual similarity between watermarked and unwatermarked data.

**Robustness.** The robustness of a watermarking scheme can be defined as the ability of the detector to extract the hidden watermark from some altered watermarked data. The alteration can be malicious or not i.e. the alteration can result from a common processing (filtering, lossy compression, noise addition) or from an attack attempting to remove the watermark (StirMark [173], dewatermarking attack [158]). As a result, the robustness is evaluated via the survival of the watermark after attacks.

It is quite easy to see that those three parameters are conflicting. One may want to increase the watermarking strength to increase the robustness but this results



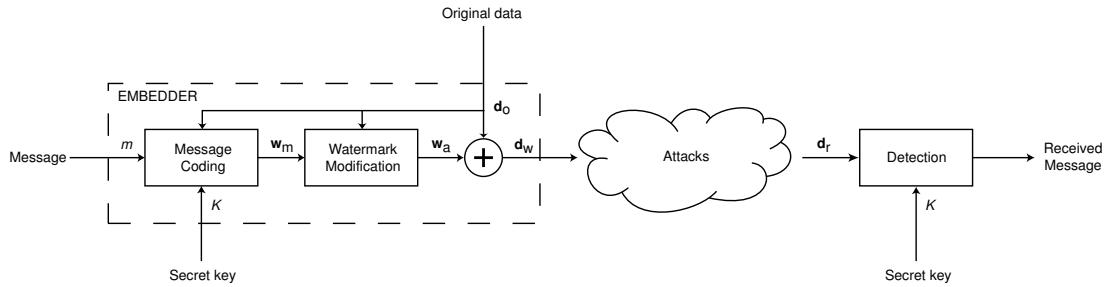
in a more perceptible watermark on the other hand. Similarly, one can increase the data payload by decreasing the number of samples allocated to each hidden bit but this is counterbalanced by a loss of robustness.

As a result, a trade-off has to be found and it is often tied to the targeted application. It is useless to design a high capacity algorithm if there are only a few different messages to be hidden in practice. This is typically the case in a copy control application where two bits are enough to encode the three messages copy-always, copy-once and copy-never. Most of the time, the watermark signal should have a low energy so that the induced distortion remains imperceptible. However in a high degrading environment, it is sometimes necessary to embed a strong watermark so that it survives the transmission. Finally some applications do not require the watermark to be robust. In fact the weakness of a fragile watermark can even be exploited to ensure the integrity of digital data [159]. If no watermark is found, digital data is not considered legitimate and is discarded. There is not consequently one optimal watermarking algorithm. Each watermarking scheme is based on a different trade-off and one has to be cautious when benchmarking various algorithms. It should be ensured that the methods under investigation are evaluated under similar conditions [109]. In other words, to perform a fair performance comparison in terms of robustness, the evaluated watermarking algorithm should have roughly the same capacity and introduce approximately the same visual distortion.

### 1.1.2 Informed Watermarking

The last few years have seen the emergence of a new trend in the watermarking community. The watermarking process is now seen as the transmission of a signal through a noisy channel. Original data is then seen as interfering noise which reduces significantly the amount of reliably communicable watermark information. In this new perspective, Chen and Wornell noticed a precious paper written by Costa [30]. He showed that, if a message is sent through a channel corrupted by two successive additive white Gaussian noise sources and if the transmitter knows the first noise source, then this first noise source *has no effect on the channel capacity*. From a watermarking point of view, the message can be seen as the watermark, the first known noise source as the original data and the second unknown noise source as the attacks. Even if Costa's model is substantially different from a real watermarking system, it means that side information at the embedder enables to reduce interference from the original data. This implication has received further support from subsequent theoretical work.

In Figure 1.2, the embedder can be seen as blind. Information contained in the original data is not exploited during the message coding and watermark modification steps. Costa's work encourages designing new algorithms based on Figure 1.4 where side information is taken into account during those two steps.



**Figure 1.4:** Informed watermarking scheme.

Informed watermarking can be done during message coding (informed coding) and/or watermark modification (informed embedding). With informed coding, for a given message, a pool of different alternative watermarks is available and the embedder chooses the one for which the interference introduced by the original data will be minimized. With informed embedding, the goal is to optimally modify the watermark so that the detector extracts the expected message. A typical example is to perceptually shape the watermark accordingly to the original data so that fidelity is increased while robustness is maintained.

### 1.1.3 Further Reading

Presenting the whole theory behind digital watermarking is far beyond the scope of this thesis and the interested reader is invited to read the various books devoted to the subject. An introducing overview of digital watermarking can be found in [95]. Further details are developed in [35] where the authors even provide samples of source code. This should be enough to have an overview of the domain. Nevertheless, many other complementary books [87, 196, 5, 168] also deal with digital watermarking and it may be useful to browse through their contents to consider different points of view. Finally, for an in depth discussion on informed watermarking, the reader is redirected toward Egger's Ph.D. thesis [66].

## 1.2 Outline of the Thesis

Digital watermarking has first been extensively studied for still images. Today however, many new watermarking schemes are proposed for other types of digital multimedia data, so called as new objects: audio, video, text, 3D meshes... This thesis is completely devoted to digital video watermarking and is divided in three main parts.

The first part gives an overview of video watermarking. Chapter 2 exhibits the potential benefit induced by introducing digital watermarking in applications

where video material is used. Next, the major trends in video watermarking are highlighted in Chapter 3 and a few reference algorithms are rapidly presented. Then the main challenges to be taken up in the context of video watermarking are reviewed in Chapter 4. One of these issues is then chosen and further developed in the remainder of the thesis.

As a result, the second part focuses on security evaluation using collusion attacks. In particular, two alternative strategies are presented. Chapter 5 identifies security pitfalls when the embedded watermarks show off some kind of redundant structure. In this case, with multiple observations, an attacker is able to learn some information about this structure and to exploit this knowledge to defeat the system. Alternatively, Chapter 6 also shows that completely uncorrelated watermarks can also be removed.

The last part introduces a couple of countermeasures to combat the different jamming attacks which have been described in the previous part. In Chapter 7, motion compensated watermarking enables to survive temporal filtering along the motion axis. Additionally, signal coherent watermarks are designed in Chapter 8 so that block replacement attacks no longer remove embedded watermarks. Eventually, the different contributions of this thesis are summarized in chapter 9 and tracks for future work are also indicated.



# Part I

## A State-of-the-Art Overview



---

# Applications

---

If the increasing interest concerning digital watermarking during the last decade is most likely due to the increase in concern over copyright protection of digital content, it is also emphasized by its commercial potential. The following section is consequently completely dedicated to the presentation of various applications in which digital watermarking can bring a valuable support in the context of video. Digital video watermarking may indeed be used in many various applications and some of them are far from the original copyright enforcement context. The applications presented in this section have been gathered in Table 2. This is not an exhaustive list and many applications are still to be imagined. In Section 2.1, digital watermarking is presented as a possible mean to establish a covert communication channel. Next, several applications are described in Section 2.2 when the robustness constraint has almost been raised. The only remaining concern is to transmit additional information to enable new services. Finally, Section 2.3 deals with Intellectual Property Right (IPR) protection. This was indeed the primary purpose which has motivated the introduction of digital watermarking.

## 2.1 Steganography

The term steganography is derived from the Greek words *στεγανος* (steganos) which means “covered” and *γραφειν* (graphein) which means “writing”. It is the art of concealed communications i.e. the goal is to keep the very existence of the message secret. From time immemorial, human beings have always invented new means of transmitting information in a secret fashion. Some of the earliest

**Table 2.1:** Video watermarking: applications and associated purpose.

Applications	Purpose of the embedded watermark
Steganography	Convey secret information
Labelling	Bind semantic meaning to the host content
Data compression	Transmit enhancement features (color, audio)
Error recovery	Convey additional information to enable error control
Proof of ownership	Identify the video copyright holder
Access control	Prevent unauthorized playback and copying
Broadcast monitoring	Identify the broadcasted video items
Fingerprinting	Identify the source of leakage in a content distribution network
Authentication	Ensure that the original video has not been tampered with

examples trace back to the ancient Greece and have been described in details by Herodotus [81]. At this time, common practices consisted of etching messages in wooden tablets and covering them with wax, or tattooing a shaved messenger's head, letting his hair grow back, then shaving it again when he arrived at his contact point. Innovations to produce such covert communications usually occur during time of war. For instance, during World War II, a German spy sent the following *innocent sounding* message [91]:

Apparently neutral's protest is thoroughly discounted and ignored.  
Isman hard hit. Blockade issue affects pretext for embargo on byprod-  
ucts, ejecting suets and vegetable oils.

Now, taking the second letter in each word reveals the hidden sensitive information:

Pershing sails from NY June 1.

Digital watermarking can also be exploited to convey a secret message within a video stream. However, in this perspective, a few specific points have to be raised. First of all, the embedding process should not leave any footprint. Indeed covert communication schemes are often modeled using the so-called *prisoner problem* i.e. a prisoner who wants to communicate with a party outside his cell [169]. To avoid any illegal communication, the warden examines all the messages sent by the prisoner and punishes him every time that a covert message is revealed (even if he is not able to read and/or understand this message). Once casted within a statistical framework, this problem can be analyzed and solved using information-theoretical tools. The main point here is that someone without the secret key



should not be able to detect whether a video stream is watermarked or not using for example steganalysis tools [23]. Another key aspect about steganography is that the host signal can be chosen depending on the message to be transmitted. Nevertheless, the statistical invisibility property required for steganography is very difficult to achieve and further research in this area is in progress. Furthermore, despite all the rumors alleging that Al-Qaeda exploited steganography in its video messages to manage its terrorism activities, no formal evidence of such claims has been made public yet.

## 2.2 Data Hiding

Data hiding basically refers to applications where digital watermarking is exploited to transmit additional information and thus enable new services for the customers. In such a framework, the robustness constraint can be raised or at least lightened. Indeed, since the embedded watermark is exploited to provide an additional service, customers are likely not to process the video stream at the risk of losing this new functionality. As a result, the watermark is only required to survive a few operations which are clearly identified depending on the application. On the other hand, a large number of bits need to be embedded so that a valuable service can be obtained. Three alternative data-hiding applications are rapidly reviewed in the next subsections to highlight the potential interest of such applications.

### 2.2.1 Labelling for Data Retrieval

Now that we have entered the digital era, more and more multimedia contents are available in a digital form. In fact, there are so many contents to be managed that storage databases have become huge to the point where accessing useful information is an issue. To overcome this problem, content-based retrieval is receiving an increasing interest. The basic idea consists in letting the user indicate to the system the type of content he/she is looking for. This description can be done according to several criteria e.g. the multimedia type (text, audio, image, video), information about creation (author, place, date) or even better a rough description at a semantic level. For example, a typical semantic request could be *find all the videos depicting animal life in the jungle*. Unfortunately, automated data analysis at a semantic level is a very difficult task. A possible way to get around this setback is to manually attach to each multimedia item a description of its semantic content. Such a hand-made labeling process is of course very time-consuming and one would ensure that each label is indissolubly tied with the object that it describes. In this perspective, digital watermarking sounds to be a perfect match to enforce this labeling functionality. When a multimedia

item moves from an archive to a new one, possibly passing through the analog domain and changing the storage compression format, the watermark encoding the description at a semantic level survives. Such an approach can even reach a higher level if the MPEG-4 video compression standard is considered [105]. In this upcoming format, video content is described using small units called *video objects*. Examples of such objects could be movie characters, cars, animals, etc. The promising asset of video objects is that they virtually enable to compose any new scene using audio-visual objects taken from other videos. As a result, if a label is embedded in a video object, when it is copy-edited to create a novel video sequence, the label automatically travels with the object thus avoiding the necessity of labeling it again.

### 2.2.2 Data Hiding for Data Compression

Embedding useful data directly into the video stream can spare much storage space. A typical video stream is made up of two different parallel streams: the audio and video streams. Those two streams need to be synchronized during playback for pleasant viewing, which is difficult to maintain during cropping operations. Hiding the audio stream into the video one [141] will implicitly provide efficient and robust synchronization, while significantly reducing the required storage need or available bandwidth. In the same fashion, color information can also be embedded within a grey scales video stream to be used as an enhancement layer on the decoder side [74]. The Picture-in-Picture technology is present in many television sets and uses separate data streams to superimpose a small video window over the full-size video displayed on the television set. Here again, digital watermarking enables to embed the secondary video stream into the carrier one [181]. As a result, during playback, the watermark is extracted and the embedded video is displayed in a window within the host video. The asset of such an approach is that only one stream needs to be transmitted. This strategy can be further extended so that a user can switch to the *PG* version of an *R* rated movie, with alternative dialogs and scenes replacing inappropriate content.

### 2.2.3 Data Hiding for Error Recovery

The attentive reader may have noticed that video watermarking and video coding are two conflicting technologies. A perfect video codec should remove any extra redundant information. In other words, two visually similar videos should have the same compressed representation. If one day, such an optimal video codec is designed, then video watermarking will disappear since unwatermarked and watermarked data would have the same compressed representation. Digital watermarking can be consequently seen as the exploitation of the imperfections of the compression algorithms to hide information. However recent research has shown

that digital watermarking can benefit to the coding community. The video coding process can be sequenced in two steps. During *source coding*, any redundant information is removed to obtain the most possible compressed representation of the data while keeping its original visual quality. This compressed representation is then submitted to *channel coding*, where extra redundant information is added for error correction. Channel coding is necessary since errors are likely to occur during the transmission, e.g. in a wireless environment.

Digital watermarking can be introduced as an alternative solution for introducing error correcting information after source coding, without inducing any overhead [9]. Experiments have demonstrated the feasibility of such a strategy and results are even reported showing that digital watermarking can exhibit better performances than traditional error correction mechanisms [162]. The basic idea is to embed in the bitstream redundant information of important features, e.g. edge information or motion vectors, for future error concealment purposes. Many variations on this theme have been published in the literature: motion information about the current frame can be embedded in the next frame [172], block information (block type, major edge direction) can be hidden in distant blocks [201], a coarse representation of each video frame can be inserted in the frequency domain [19, 128], etc. The major challenge in this perspective of data hiding for error recovery is to jointly optimize all the parameters i.e. source coding, channel coding and watermarking parameters.

## 2.3 IPR Protection

In the mid 90's, the digital world took the homes by storms. Personal computers became more and more popular, digital devices were interconnected on high speed networks, efficient softwares enabled file sharing and multimedia items editing... It was a complete revolution. But it also raises many concerns regarding IPR protection. The Internet was indeed so *free* that the notion of intellectual property was almost obsolete. Copyrighted items were exchanged freely, thus resulting in a drastic loss of revenues for the majors from the music and cinema industries. Tampering digital data was so easy that trust on the Internet almost vanished. As a result, researchers have investigated how to restore the value of intellectual property in the digital world. Although many solutions rely on cryptography, digital watermarking has also been proposed as a mean to ensure IPR protection and a few applications will be reviewed in the next subsections to illustrate this approach.

### 2.3.1 Proof of Ownership

Copyright protection is historically the very first targeted applications for digital watermarking. The underlying strategy consists in embedding a watermark, identifying the copyright owner, in digital multimedia data. If an illegal copy is found, the copyright owner can prove his/her paternity thanks to the embedded watermark and can sue the illegal user in court. This perfect scenario is however likely to be disturbed by malicious users in the real world [36]. If an attacker adds a second watermark into a video clip, both the original owner and the attacker can claim ownership and therefore defeat the purpose of using watermarking. Using the original video clip during the verification procedure happens to prevent the multiple ownership problems in some cases. However, this problem still holds if the watermarking algorithm is invertible because it allows the attacker to produce his/her own counterfeited original video clip. In this case, both the original owner and the attacker have an original video clip which contains the watermark of the other one. As a result, nobody can claim ownership! This situation is referred to as the *deadlock* problem in the watermarking community. Watermark algorithms are consequently required to be non-invertible to provide copyright protection services and they are often backed up by an elaborated protocol with a trusted third party. Copyright protection has been investigated for video watermarking [155] even if this is not the most targeted application.

Instead of protecting the whole video stream, copyright owners might rather want to protect only a part of the video content. The commercial value in a video is indeed often concentrated in a small number of video objects e.g. the face of an actor. Moreover, future video formats will distinguish the different objects in a video. This will be the case with the upcoming MPEG-4 format. Recent research has consequently investigated digital watermarking of video objects [154]. Watermarking video objects prevents unauthorized reuse in other video clips. However video objects are likely to be submitted to various video editing such as scaling, rotation, shifting and flipping. As a result, special care must be taken regarding the resilience of the watermark against such operations. This can be quite easily obtained thanks to a geometrical normalization [11], according to the moments and axes of the video object, prior to embedding and extraction.

### 2.3.2 Access Control

The Digital Versatile Disk (DVD) and DVD players appeared on the consumer market in late 1996. This new technology was enthusiastically welcomed since DVD players provide a very high-quality video signal. However, the advantages of digital video are counterbalanced by an increased risk of illegal copying. In contrast to traditional VHS tape copying, each copy of digital video data is a perfect reproduction. This raised the concern of copyright owners and Hollywood

studios requested that several levels of copy protection should be investigated before any device with digital video recording capabilities could be introduced.

The Copy Protection Technical Working Group (CPTWG) has consequently been created to work on copy protection issues in DVD. A standard has not been defined yet. However a system, which could become the future specification for DVD copy protection, has been defined [12]. The three first components are already built in consumer devices and the other three are still under development.

**The Content Scrambling System (CSS).** This method has been developed by Matsushita and scrambles MPEG-2 video. A pair of keys is required for descrambling: one is unique to the disk and the other is specific to the MPEG file being descrambled. Scrambled content is not viewable.

**The Analog Protection System (APS).** Macrovision developed this system to modify NTSC/PAL video signals. The resulting video signals can be displayed on televisions but cannot be recorded on VCR's. However, the data on a disk are not NTSC/PAL encoded and APS has to be applied after encoding in the DVD player. Some bits are consequently stored in the MPEG stream header and give the information of whether and how APS should be applied.

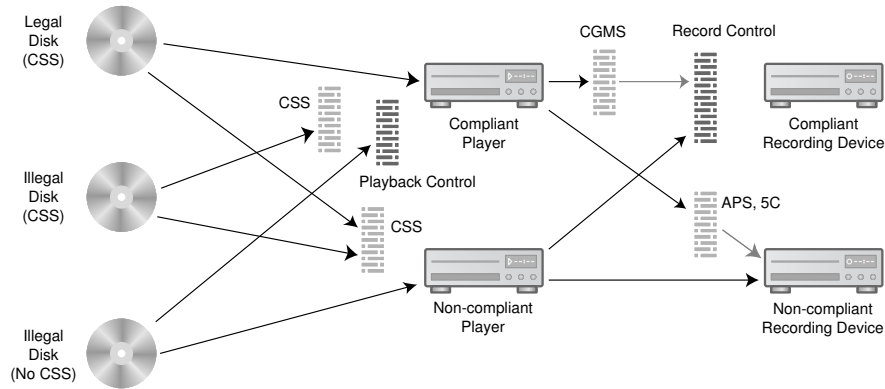
**The Copy Generation Management System (CGMS).** This is a pair of bits stored in the header of an MPEG stream encoding one of the three possible rules for copying: *copy-always*, *copy-never* and *copy-once*. The copy-once case is included so that time-shifting is allowed i.e. a copy of broadcast media is made for later viewing.

**5C.** A coalition of five companies designs this mechanism. It allows several compliant devices, connected to the same digital video bus e.g. IEEE1394 (firewire), to exchange keys in an authenticated manner so that encrypted data can be sent over the bus. Noncompliant devices do not have access to the keys and cannot decrypt the data.

**Watermarking.** The main purpose of watermarking is to provide a more secure solution than storing bits in the MPEG stream header. In DVD, digital watermarking is primarily intended for the CGMS bits and secondary for the APS bits.

**Physical identifiers.** The idea is to design secure physical media identifiers to be able to distinguish between original media and copies.

Figure 2.1 shows how those mechanisms have been put together in the DVD so that copy protection is enforced. The additional performance brought by watermarking is emphasized by the dark walls.



**Figure 2.1:** DVD copy-protection system [12].

Everything starts when Hollywood studios release a new copyrighted DVD with CGMS bits encoding the message *copy-never*. Both CSS keys are stored on the lead-in area of the DVD. This area is only read by compliant players. This prevents factory-pressed legal disks from being displayed by noncompliant players. Moreover bit-for-bit illegal copies will contain CSS scrambled content, but not the keys. As a result, such illegal copies cannot be displayed by any player, compliant or not. If the output signal given by compliant players is digital, CGMS bits prevent copying in the compliant world while 5C will avoid any communication with any noncompliant devices. However, to date, analog monitors are still widespread and even compliant players output an analog signal for compatibility. Since CGMS bits do not survive digital to analog conversion, watermarking is introduced to avoid copying in the compliant world. Unfortunately, in the noncompliant world, APS only disables copying of analog NTSC/PAL signals on VHS tapes. Disks without CSS or CGMS can then be easily generated e.g. thanks to a simple PC with a video capture card.

Now illegal disks containing unscrambled content without CSS or CGMS are available. They may have been generated as described previously. But they can also be generated directly from an original legal disk since CSS was cracked in 1999 [149]. The remaining CGMS bits can then be trivially stripped from the MPEG stream. Such illegal copies can of course be displayed by noncompliant players but watermarking has to be introduced to prevent those copies to enter the compliant world. Compliant players will detect the *copy-never* watermark embedded in *unscrambled* DVD-ROM and will refuse playback. The video signal given by a noncompliant player can be recorded by noncompliant recording devices. However watermarking prevents copying with compliant devices. The whole protection system results in two hermetically separated worlds. A consumer should have both types of players to display legal and illegal disks. The expense of such a strategy will help to “keep honest people honest”.

It is important for DVD recorders to support the *copy-once* case to allow time shifting. When the recorder detects the *copy-once* message, it should modify the stream so that the hidden message becomes *copy-never*. This can be easily done in the case of stored bits in the MPEG header but it is less straightforward when using watermarking. Two proposals are investigated. The first one consists in superimposing a second watermark when a *copy-once* watermark is detected. The two watermarks together will then encode the message *copy-never*. The second proposal avoids remarking and exploits the ticket concept [125]. The idea is to use two hidden signals: an embedded watermark  $W$  and a physical ticket  $T$ . There exists a relationship between the two signals which can be written  $F^n(T) = W$ , where  $F(.)$  is a one way hash function and  $n$  is the number of allowed passages through compliant devices. The ticket is decremented each time the data go through a compliant player or recorder. In other words, the ticket is modified according to the relation  $T' = F(T)$ . During playback, the ticket in transit can be embedded in MPEG `user_data` bits or in the blanking intervals of the NTSC/PAL standard. During recording, the ticket can be physically marked in the *wobble*<sup>1</sup> in the lead-in of optical disks.

### 2.3.3 Broadcast Monitoring

Many valuable products are distributed over the television network. News items, such as those sold by companies like Reuters or Associated Press, can be worth over 100,000 USD. In France, during the final of the 2002 FIFA World Cup Korea Japan<sup>TM</sup>, advertisers had to pay 100,000 Euros to broadcast a thirty seconds commercial break shot on television. The same commercial would even have been billed 220,000 Euros if the French national team had played during the final. Owners of copyrighted videos want to get their royalties each time their property is broadcasted. The whole television market is worth many billions of dollars and Intellectual Property Rights violations are likely to occur. As a result, a broadcast surveillance system has to be built to monitor all broadcasted channels. This will help verifying that content owners get paid correctly and that advertisers get what they have paid for. Such a mechanism will prevent confidence tricks such as the one discovered in Japan in 1997 when two TV stations were convicted of overbooking air time [98].

The most naive approach of broadcast monitoring consists of a pool of human observers watching the broadcasts and recording whatever they see. However, this very simple method is far from being optimal. Human employees are expensive and are not foolproof. As a result, research has been conducted to find a way of automating broadcast monitoring. The first approach, referred to as *passive*

---

<sup>1</sup>The wobble is a radial deviation of the position of pits and lands relative to the ideal spiral. Noncompliant recorders will not insert a ticket and the illegal disk will not enter the compliant world.

*monitoring*, basically makes a computer simulate a human observer: it monitors the broadcasts and compares the received signals with a database of known videos. This approach is non intrusive and does not require cooperation from advertisers or broadcasters. However such a system has two major drawbacks. First, it relies on the comparison between received signals against a large database, which is non trivial in practice. Pertinent signatures, clearly identifying each video, have to be defined and an efficient search for nearest neighbors in a large database has to be designed. This results in a system that is not fully reliable. This may be accurate for acquiring competitive market research data i.e. when a company wants to know how much its competitors spend in advertising. On the contrary, a small error rate (5%) is dramatic for verification services because of the large amount of money at stake. The second con is that the reference database is likely to be large and the storage and management costs might become rapidly prohibitive.

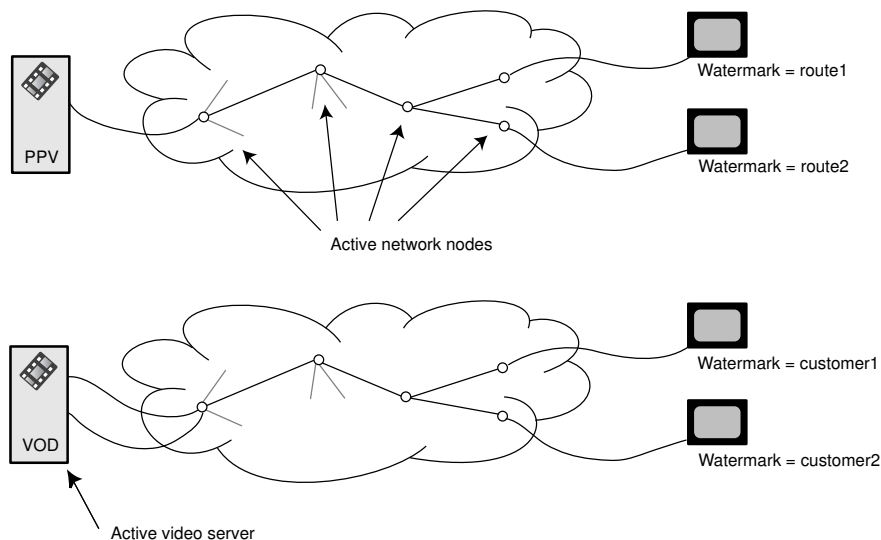
To reach the accuracy required for verification services, a new kind of systems, referred as *active monitoring*, has been designed. The underlying idea is to transmit computer-recognizable identification information along with the data. Such identification information is straightforward to decode reliably and to interpret correctly. This approach is known to be simpler to implement than passive monitoring. First implementations of active monitoring placed the identification information in a separate area of the broadcast signal e.g. the Vertical Blanking Interval (VBI) of an analog NTSC/PAL video signal. However dissimulating identification data into other data is exactly the purpose of digital watermarking. Even if watermark embedding is more complicated than storing information in some unused part of a video stream, digital watermarking can be considered as a robust way to implement active monitoring. The European project VIVA (Visual Identity Verification Auditor) proved the feasibility of such a system [43]. The participants used a real-time watermarking scheme which provides active monitoring services over a satellite link. The complexity of the detection algorithm is moderate enough to allow simultaneous monitoring of many channels.

### 2.3.4 Fingerprinting

The explosion of the Internet has created a new way of acquiring copyrighted content. When a user wants to obtain a new video clip or a new movie, the simplest strategy is to log on Internet and to use one of the popular peer-to-peer systems e.g. Napster, KaZaA, Morpheus, eMule. Multimedia digital contents, stored throughout the world on thousands of computers logged on at the same moment, will instantly get accessible. As a result, European engineering students often download and watch the most recent Hollywood movies a long time before they are released in their own country. The situation is even worse in audio with



the exchange of MP3<sup>2</sup> files. As a result, copyright owners lose a large amount of royalties [116]. Legal action has been taken to ban such distributed systems but, when Napster has been sentenced guilty, two other systems appeared. The basic problem does not come from peer-to-peer systems. It would be a great tool if only legal data was transiting on such distributed networks. The problem is that a traitor has made available copyrighted material without any kind of permission. The basic idea would consequently be to be able to identify the traitor when an illegal copy is found to sue him/her in court. This can be done by embedding an indelible and invisible watermark identifying the customer.



**Figure 2.2:** Alternative watermarking strategies for video streaming.

### Tracing malicious customers

In a near future, the way people are looking at TV will be significantly modified. Video streaming is indeed likely to become more and more widespread. It is consequently necessary to find a way of protecting digital video content and digital watermarking seems to be a potential candidate [123]. Pay-Per-View (PPV) and Video-on-Demand (VoD) are two real-life applications of video streaming. In both applications, digital watermarking can be used to enforce a fingerprinting policy. The customer ID is embedded into the delivered video data to trace back any user breaking his/her license agreement. The main difference resides in the watermarking strategy as depicted in Figure 2.2. Embedding the watermark on the customer side has been suggested [75] but it should be avoided if possible

<sup>2</sup>The MPEG-1 audio layer 3 (MP3) is a popular audio format for transmitting audio files across the Internet.

to prevent reverse engineering. In a PPV environment, a video server multicasts some videos and customers have only to connect to the server to obtain the video. The video server is passive. At a given moment, it delivers the same video stream to multiple users. To enforce fingerprinting, a proposed method [17] is to have each network element (router, node or whatever) embed a piece of watermark as the video stream is relayed. The resulting watermark will contain a trace of the route followed by the video stream. Such a strategy requires support from network providers, who might not be forthcoming about it. In a VoD framework, the video server is active. It receives a request from a customer and sends the requested video. It is a multi-unicast strategy. This time, the video server can insert a watermark identifying the customer since each connection is dedicated to only one customer. The main challenge is then to scale the system to many users.

### **Digital cinema**

Another fingerprinting application has been considered with the apparition of a new kind of piracy. Nowadays illegal copying of brand new movies projected onto cinema screen by means of a handheld video camera has become a common practice. The most memorable example is surely when, one week after its US release, the very anticipated “Starwars Episode I: The Phantom Menace” was available on the Internet in a low quality version, with visible head shadows of audience members. Although the quality of such copies is usually very low, their economical impact can be enormous. Moreover, the upcoming digital cinema format to be introduced in theatres raises some concern. With higher visual quality, the threat becomes larger and Hollywood studios want to oblige cinema owners to forbid the presence of video cameras in their premises. Once again, digital watermarking could provide a solution [76]. A watermark can be embedded during show time identifying the cinema, the presentation date and time. If an illegal copy created with a video camera is found, the watermark is extracted and the cinema to blame is identified. After many blames, the cinema is sanctioned with a ban on the availability of content. However, even if this emerging application has held the attention of several industries [186, 130], a recent industry survey has stated that there is presently no effective response to the challenge of creating secure watermarks [127] and that further research is required.

#### **2.3.5 Authentication**

Large amounts of video data are distributed throughout the Internet every day. More and more video cameras are installed in public facilities for surveillance purpose. However, popular video editing softwares permit today to easily tamper with video content, as shown in Figure 2.3, and video content is no more

reliable. For example, in some countries, a video shot from a surveillance camera cannot be used as a piece of evidence in a courtroom because it is not considered trustworthy enough. When someone is emailed a somewhat unusual video, it is quite impossible to determine if it is an original or a hoax. Authentication techniques are consequently needed to ensure authenticity of video content. Methods have to be designed for verifying the originality of video content and preventing forgery. When a customer purchases video content via electronic commerce, he wants to be sure that it comes from the alleged producer and that no one has tampered with the content. The very first research efforts for data authentication used cryptography. The major drawback of such an approach is that it provides a *complete verification*. In other words, the data is considered as untouchable and the data for authentication has to be exactly the same one as the original one. But this strong constraint might be too restricting. One might prefer to allow some distortions on the digital data if the original content has not been significantly modified. This is typically the case in wireless environment where some noise is added to the data. This approach is referred as *content verification*.



**Figure 2.3:** Original and tampered video frames.

Researchers have investigated the use of digital watermarking to verify the integrity of digital video content. A basic approach consists in regularly embedding an incremental timestamp in the frames of the video [139]. As a result, frame cuts, foreign frame insertion, frame swapping, and frame rate alteration can be easily detected. This approach is very efficient for detecting temporal alteration of the video stream. However, it might fail in detecting alterations of the content itself e.g. a character is completely removed from a movie. Investigations have consequently been conducted to prevent modifications of the video content itself. One proposal [46] embeds the edge map of each frame in the video stream. During the verification process, if the video content has been modified, there will be a mismatch between the extracted edge map from the verified video and the

watermarked edge map. The detector will consequently report content tampering. Another proposal exploits the idea that a movie is made up of one audio and one video stream and that both need to be protected against unauthorized tampering. The fundamental idea is then to combine video and audio watermarking [45] to obtain an efficient authenticating system. Features of both streams are embedded one into another. Modification from either the sound track, or the video track, is immediately spotted by the detector, since the extracted and watermarked features will differ.

### Video surveillance

Recent advances have significantly enlarged the scope and enhance the quality of automatic video surveillance. Researchers are constantly trying to improve this technology with respect to continuous and effective monitoring, cost reduction and reliable control of distant or high-risk sites. However, in practice, several issues have to be considered and the integrity of video surveillance data in front of a court of law is one of them. Indeed, digital videos have virtually no value because of the numerous public editing tools available. Furthermore, proving the true origin of data (who, when, where) should also be possible. A straightforward cryptographic solution [10] consists in *printing* source information, such as the date, the time and/or the place, in each video frame and then in computing a digest of each frame by means of a proper hash function. Those digests are then encrypted with an asymmetric-key scheme and are transmitted along with the video stream. Nevertheless, alternative approaches are still desirable to overcome possible drawbacks of this cryptographic approach. For example, the image digest is tied to an image format and thus constrains the possibilities of authenticating the video stream. Another point, which has already been raised previously, is that no distinction is made between malicious and innocuous modifications: the digest changes dramatically.

Those weaknesses have motivated the design of watermarking based solutions for video surveillance sequences authentication [10, 161]. There are two main approaches. First, fragile or semi-fragile watermarks can be exploited. Whereas such watermarks survive innocuous manipulations such as moderate lossy compression, they are dramatically altered when tampering occurs e.g. replacing the guilty face with *background* material. Another solution consists in hiding a summary of the video stream, whatever it is, using a robust watermark. Any tampering is then detected by comparing the video stream with the underlying summary. Therefore, digital watermarking can be regarded as a really candidate technology to ensure authenticity. However, this emerging technology is relatively immature with respect to cryptography and further studies are required to answer open questions such as the ultimate level of security and robustness achievable through watermarking.

# 3

---

## Major Trends

---

Although almost 20% of the scientific literature considers video content according to Figure 1.1, digital video watermarking still remains a somewhat unexplored area of research which basically benefits from the results obtained for still images. Therefore, this chapter will not detail an exhaustive list of algorithms for video watermarking. The goal is rather to isolate some major trends which can be of interest to give a global view of the scientific production on the topic. Of course, a few selected reference algorithms will be examined in details for illustration purpose. In video watermarking, the research effort is indeed divided into three main directions. In Section 3.1, video content is considered as a succession of still images. As a result, a simple and straightforward approach is to reuse existing watermarking schemes for still images in a frame-by-frame fashion. Alternatively, as presented in Section 3.2, one can try to integrate the temporal dimension in the watermarking procedure. In practice, this can be implemented very simply

**Table 3.1:** Pros and cons of the different approaches to video watermarking.

	<b>Pros</b>	<b>Cons</b>
<i>Image <math>\rightarrow</math> video</i>	Inherit from all the results for still images	Computationally intensive
<i>Temporal dimension</i>	Video-driven algorithms which often permit higher robustness	Can be computationally intensive
<i>Compression standard</i>	Simple algorithms which make real-time achievable	Watermark inherently tied to the video format

by considering video content as a collection of signal samples or, very much more elaborate, by exploiting for instance 3D signal transforms. Finally, the last approach relies on the observation that video content is usually compressed with a specific video compression standard for storage/transmission convenience. Thus, Section 3.3 describes a few ways of exploiting such standards to obtain very efficient video watermarking schemes. Of course, each one of those strategies has its own pros and cons with respect to complexity, robustness performances, etc. They have been reminded in Table 3.1.

### 3.1 From Still Images to Video Watermarking

In its very first years, digital watermarking has been extensively investigated almost exclusively for still images. Many interesting results and algorithms were found and when new areas, such as video, were researched, the basic concern was to try to reuse the previously found results. As a result, the watermarking community first considered digital video content as a succession of still images and adapted existing watermarking schemes for still images to the video in a frame-by-frame fashion. Exactly the same phenomenon occurred when the coding community switched from image coding to video coding. The first proposed algorithm for video coding was indeed Moving JPEG (M-JPEG), which simply compresses each frame of the video with the image compression standard JPEG. The simplest way of extending a watermarking scheme for still images is to embed the same watermark in the frames of the video at a regular rate. On the detector side, the presence of the watermark is checked in every frame. If the video has been watermarked, a regular pulse should be observed in the response of the detector [6]. However, the main drawback of such a scheme is that it has no payload. The detector only tells if a given watermark is present or not but it does not extract any hidden binary message. Video content is much larger in size than a single still image. Since one should be able to hide more bits in a larger host signal, high payload watermarks for video could be expected. There are alternative ways to achieve this goal. To begin with, the redundantly embedded watermark signal can encode the same payload i.e. the same message is embedded in all the frames of the video. This is typically the case with the algorithm JAWS which is further detailed hereafter [93]. Alternatively, one can also embed an independent multi-bits watermark in each frame of the video to exploit the whole available bandwidth [44]. It should be noticed in this latter case that the gain in embedding capacity is counterbalanced by a loss of robustness since each bit is spread on fewer samples and an increased sensibility against desynchronization.


### JAWS: Just Another Watermarking System

The JAWS algorithm has been designed at the end of the 90's by researchers from Philips Research [93]. It was originally proposed for broadcast monitoring and has also been one of the leading candidates for watermarking in DVD. The embedding process is depicted in Figure 3.1. First of all, an  $M \times M$  normally distributed reference pattern  $\mathbf{p}_r$  is generated with a secret key. In a second step, a reference watermark  $\mathbf{w}_r$  is created according to the following equation:

$$\mathbf{w}_r = \mathbf{p}_r - \text{shift}(\mathbf{p}_r, m) \quad (3.1)$$

where the  $\text{shift}(\cdot)$  function returns a cyclically shifted version of the reference pattern  $\mathbf{p}_r$  and  $m$  is some binary message to be hidden. A major characteristic of JAWS is that, the message is encoded by the shift between the two reference patterns. This reference watermark  $\mathbf{w}_r$  is then tiled, possibly with truncation, to obtain the full-size watermark  $\mathbf{w}$ . For each frame, this watermark is then perceptually shaped so that the watermark insertion remains imperceptible. Each element  $i$  of the watermark is scaled by the local activity  $\lambda(i)$  of the frame, given for instance by Laplacian filtering. The flatter the region is, the lower the local activity is. This is coherent with the fact that the human eye is more sensitive to noise addition in flat regions of an image. Finally, the watermark is scaled by a global embedding strength  $s$  and added to the frame  $\mathbf{f}$  to obtain the watermarked frame  $\check{\mathbf{f}}$ . Therefore, the overall embedding process can be expressed as:

$$\check{\mathbf{f}}(i) = \mathbf{f}(i) + s \cdot \lambda(i) \cdot \mathbf{w}(i) \quad (3.2)$$



**Figure 3.1:** JAWS embedding procedure.


On the detector side, the incoming frames are folded, summed and stored in an  $M \times M$  buffer  $\mathbf{b}$ . The detector looks then for all the occurrences of the reference

pattern  $\mathbf{p}_r$  in the buffer with a two dimensional cyclic convolution. Since such an operation is most efficiently computed in the frequency domain, this leads to Symmetrical Phase Only Matched Filtering (SPOMF) detection which is given by the following equation:

$$\text{SPOMF}(\mathbf{b}, \mathbf{p}_r) = \text{IFFT} \left[ \phi \left( \text{FFT}(\mathbf{b}) \right) \cdot \phi \left( \text{FFT}(\mathbf{p}_r)^* \right) \right]$$

$$\text{with } \phi(x) = \begin{cases} x/|x| & \text{if } x \neq 0 \\ 1 & \text{if } x = 0 \end{cases} \quad (3.3)$$

where  $\text{FFT}(\cdot)$  (resp.  $\text{IFFT}(\cdot)$ ) denotes the forward (resp. inverse) Fast Fourier Transform and  $x^*$  the complex conjugation operation. Figure 3.2 shows the result of such a detection. Two peaks can be isolated which correspond to the two occurrences of  $\mathbf{p}_r$  in  $\mathbf{w}_r$ . The peaks are oriented accordingly to the sign before their associated occurrence of  $\mathbf{p}_r$  in Equation (3.1). Because of possible positional jitter, all the relative positions between the peaks cannot be used and relative positions are forced to be multiple of a grid size  $G$ . Once the detector has extracted the peaks, the hidden payload can be easily retrieved according to the estimated shift. It should be noted that this scheme is inherently shift invariant since a shifting operation does not modify the relative position of the peaks. Significant improvements have been added to this scheme afterwards. For example, shift invariance has been further exploited to increase the payload [132] and simple modifications enabled to obtain scale invariance [184].



**Figure 3.2:** Example of SPOMF detection.



## 3.2 Integration of the Temporal Dimension

The major shortcoming of considering video content as a succession of independent still images is that the new temporal dimension is not satisfactorily taken into account. On its side, the coding community made a big step forward when it decided to incorporate the temporal dimension in their coding schemes - for instance through motion prediction - and it is quite sure that it is the advantage of the watermarking community to also investigate such a path. Perceptual shaping is another issue which highlights the fact that the temporal dimension is a crucial point in video and that it should be taken into account to design efficient algorithms. Many researchers have investigated how to reduce the visual impact of embedding a watermark within still images by considering the properties of the Human Visual System (HVS) such as frequency masking, luminance masking and contrast masking. Such studies can be easily exported to video with a straightforward frame-by-frame strategy. However, the obtained watermark is not optimal in terms of visibility since it does not consider the *temporal sensitivity* of the human eye. Motion is indeed a very specific feature of the video and new video-driven perceptual measures need to be designed to be exploited in digital watermarking [99]. The next subsections will consequently rapidly present different ways to handle this temporal dimension in digital video watermarking.

### 3.2.1 Video as a Monodimensional Signal

Quite surprisingly, one of the pioneer works in video watermarking considers video content as a one dimensional signal [79]. In other words, the algorithm discards any notion of dimensionality, should it be spatial or temporal, and looks at the video signal as a collection of signal samples. Such a signal is acquired in a line-scanning fashion as depicted in Figure 3.3 i.e. spatiotemporal dimensions are used for scanning and completely ignored afterwards. The remainder of this algorithm is further detailed below.

#### SS: Spread Spectrum

Let the sequence  $\mathbf{a}(j) \in \{-1, 1\}$  represents the watermark bits to be embedded. This sequence is spread by a chip-rate  $cr$  according to the following equation:

$$\mathbf{b}(i) = \mathbf{a}(j), \quad j.cr \leq i < (j+1).cr, \quad i \in \mathbb{N} \quad (3.4)$$

The spreading operation enables to add redundancy by embedding one bit of information into  $cr$  samples of the video signal. The obtained sequence  $b(i)$  is then scaled locally by an adjustable factor  $\lambda(i) \geq 0$  and modulated by a pseudo-random binary sequence  $\mathbf{p}(i) \in \{-1, 1\}$ . Finally, the spread spectrum watermark  $\mathbf{w}(i)$  is added to the line-scanned video signal  $\mathbf{v}(i)$ , which gives the watermarked

video signal  $\check{\mathbf{v}}(i)$ . The overall embedding process is consequently described by the following equation:

$$\check{\mathbf{v}}(i) = \mathbf{v}(i) + \mathbf{w}(i) = \mathbf{v}(i) + \lambda(i) \cdot \mathbf{b}(i) \cdot \mathbf{p}(i), \quad i \in \mathbb{N} \quad (3.5)$$


The adjustable factor  $\lambda(i)$  may be tuned according to local properties of the video signal, e.g. spatial and temporal masking of the Human Visual System (HVS), or kept constant depending on the targeted application.



**Figure 3.3:** Line scan of a video stream.

On the detector side, recovery is easily accomplished with a simple correlation. However, to reduce cross-talk between watermark and video signals, the watermarked video signal is high-pass filtered, yielding a filtered watermarked video signal  $\check{\mathbf{v}}(i)$ , so that major components of the video signal itself are isolated and removed. The second step is demodulation. The filtered watermarked video signal is multiplied by the pseudo-random noise pattern  $\mathbf{p}(i)$  used for embedding and summed over the window for each embedded bit. The correlation score  $\mathbf{s}(j)$  for the  $j$ -th bit is given by the following equation:

$$\begin{aligned} \mathbf{s}(j) &= \sum_{i=j.cr}^{(j+1).cr-1} \mathbf{p}(i) \cdot \check{\mathbf{v}}(i) \\ &= \sum_{i=j.cr}^{(j+1).cr-1} \mathbf{p}(i) \cdot \mathbf{v}(i) + \sum_{i=j.cr}^{(j+1).cr-1} \mathbf{p}(i) \cdot \lambda(i) \cdot \mathbf{b}(i) \cdot \mathbf{p}(i) \\ &= \Sigma_1 + \Sigma_2 \end{aligned} \quad (3.6)$$



**Figure 3.4:** Independent Component Analysis of the *piano* sequence [170].

The correlation consists of two terms  $\Sigma_1$  and  $\Sigma_2$ . The main purpose of filtering was to leave  $\Sigma_2$  untouched while reducing  $\Sigma_1$  down to 0. As a result, the correlation sum becomes:

$$\mathbf{s}(j) \approx \Sigma_2 \approx \sum_{i=j.cr}^{(j+1).cr-1} \mathbf{p}(i)^2 \cdot \lambda(i) \cdot \mathbf{b}(i) \approx \mathbf{a}(j).cr.\text{mean}(\lambda(i)) \quad (3.7)$$

The hidden bit is then directly given by the sign of  $\mathbf{s}(j)$ .

### 3.2.2 Video as a Temporal Signal

Even if the SS pioneer method offers a very flexible framework, which can be used as a baseline for more elaborate video watermarking schemes, it completely ignores spatiotemporal dimensions. Therefore, the resulting embedded watermark is likely not to be optimal in terms of invisibility and robustness for instance. Other approaches have consequently been investigated to better cope with the temporal dimension and one of them is to look at video content as a purely temporal signal. A typical example is to insert a temporal watermark sequence at some key-dependent specific pixel locations [144]. To ensure watermark invisibility, those embedding locations have to be carefully chosen. Indeed, if modifying a single pixel in a textured area is imperceptible in each individual video frame, it might become visible when the video is rendered. In practice, pixels that change fast along the time axis or pixels in border areas of motionless regions have been shown to be good candidates for embedding. Nevertheless, using only a few pixels for watermark embedding drastically reduces the embedding capacity. In fact, one may prefer to compute some temporal transform on the whole video to have a larger embedding space. In particular, temporal wavelet decomposition can be useful to obtain a compact *multiresolution temporal representation* of the video [182]. With such a decomposition, one can isolate a *static* (no motion) component and several *dynamic* (motion) ones. The multiresolution nature of the wavelet transform allows the watermark to exist across multiple temporal scales. For instance, if a watermark is embedded in the lowest temporal frequency (DC) wavelet frame, it exists in *all the frames* of the considered video scene. Another promising temporal transform is Independent Component Analysis (ICA). This

transform produces a set of frames which can be used as independent sources to generate the processed video sequence. A typical example is depicted in Figure 3.4. The input video is a shot of a hand playing notes on a piano keyboard and ICA outputs six independent components: one is the background of the scene and the other ones are tuned to the individual keys that have been pressed. The highly semantic role of the extracted components open avenues to produce watermarks which are related with the video scene [177].

### 3.2.3 Video as a 3D Signal

Video content can also be regarded as a three dimensional signal. This point of view has already been considered in the coding community and can be extended to video watermarking. The usage of 3-dimensionnal transforms is usually motivated by visibility and robustness considerations. For instance, 3D DFT can be exploited to obtain an alternative representation of a video sequence [39]. In this case, mid frequencies, should they be spatial or temporal, are considered for watermark embedding to achieve a trade-off between invisibility and robustness. This defines two cylindrical annuli in the 3D DFT domain which are modified so that the resulting watermark resists to MPEG compression. 3D wavelet transform [107] and 3D Gabor transform [203] have also been investigated to produce robust video watermarks. Nevertheless, considering video as a three dimensional signal may be inaccurate. The three considered dimensions are indeed not homogeneous - there are two *spatial* dimensions and one *temporal* one - and should not be treated the same way. This consideration and also the required computational cost may have reduced the research effort in this direction. However this approach remains pertinent in some very specific cases. In medical imaging for example, different slices of a scanner can be seen as different frames of a video. In this case, the three dimensions are homogeneous and a 3D-transform can be used.

## 3.3 Exploiting the Video Compression Format

Another approach to video watermarking basically considers that video content is encoded for convenience with a specific video compression standard such as MPEG. For instance, video files are stored most of the time in a lossy compressed version to spare storage space. Similarly, video is usually streamed across digital distribution networks in a compressed form to cut down bandwidth requirements. Therefore, watermarking methods have been designed to directly embed the watermark into the compressed video stream by exploiting some very specific characteristics of the compression standard. The next subsections will rapidly describe how a digital watermark can be embedded at different levels of the compression


procedure. In particular, since video compression standards usually involve a signal transform e.g. DCT for MPEG, a watermark can be inserted by modifying some coefficients in the transformed domain (Subsection 3.3.1). Video coding also exploits motion prediction. Motion information can consequently be exploited to carry some information (Subsection 3.3.2). Finally, both spatial and temporal information are quantized and represented with some symbolic codewords such as run-level codewords for DCT coefficients before entropy coding. As a result, digital video watermarking can also be performed by modifying these codewords (Subsection 3.3.3).

### 3.3.1 Modify Coefficients in the Transformed Domain

Watermarking in the compressed stream can be seen as a form of video editing in the compressed domain [134]. Such editing is not trivial in practice and new issues are raised. The previously seen SS algorithm has been adapted so that the watermark can be directly inserted in the non-zero DCT coefficients of an MPEG video stream [79]. The first concern was to ensure that the watermarking embedding process would not increase the output bit-rate. Nothing ensures indeed that a watermarked DCT-coefficient will be VLC-encoded with the same number of bits than when it was unwatermarked. A straightforward strategy consists then to watermark only the DCT coefficients which do not require more bits to be VLC encoded. The second issue was to prevent the distortion introduced by the watermarking process to propagate from one frame to another one. The MPEG standard relies indeed on motion prediction and any distortion is likely to be propagated to neighbor frames. Since the accumulation of such propagating signals may result in a poor quality video, a drift compensation signal can be added if necessary. In this case, motion compensation can be seen as a constraint.

#### DEW: Differential Energy Watermarks

The DEW method was initially designed for still images and has been extended to video by watermarking the I-frames of an MPEG stream [113]. It is based on selectively discarding high frequency DCT coefficients in the compressed data stream. The embedding process is depicted in Figure 3.5. The  $8 \times 8$  pixels blocks of the video frame are first pseudo randomly shuffled. This operation forms the secret key of the algorithm and it spatially randomizes the statistics of pixel blocks i.e. it breaks the correlation between neighboring blocks. The obtained shuffled frame is then split into  $n$   $8 \times 8$  blocks. In Figure 3.5,  $n$  is equal to 16. One bit is embedded into each one of those blocks by introducing an energy difference between the high frequency DCT-coefficients of the top half of the block (region A) and the bottom half (region B). This is the reason why this technique is called a differential energy watermark.



**Figure 3.5:** DEW embedding procedure.

To introduce an energy difference, the block DCT is computed for each  $n \times 8 \times 8$  block and the DCT-coefficients are prequantized with quality factor  $Q_{\text{jpeg}}$  using the standard JPEG quantization matrix. The obtained coefficients are then separated in two halves and the high frequency energy for each region is computed according to the following equation:

$$E(c, n, Q_{\text{jpeg}}) = \sum_{b=0}^{n/2-1} \sum_{i \in S(c)} ([\theta_{i,b}]_{Q_{\text{jpeg}}})^2 \quad \text{with } S(c) = \{i \in \{0, 63\} | (i > c)\} \quad (3.8)$$

where  $\theta_{i,b}$  is the DCT coefficient with index  $i$  in the zigzag order in the  $b$ -th DCT block,  $[\cdot]$  indicates the prequantization with quality factor  $Q_{\text{jpeg}}$  and  $c$  is a given cut-off index which was fixed to 27 in Figure 3.5. The value of the embedded bit is encoded as the sign of the energy difference  $D = E_A - E_B$  between the two regions A and B. All the energy after the cut-off index  $c$  in either region A or region B is eliminated by setting the corresponding DCT coefficients to zero to obtain the appropriate sign for the difference  $D$ . It should be noted that this can be efficiently implemented directly in the compressed domain by shifting the End Of Block (EOB) marker of the corresponding  $8 \times 8$  DCT blocks toward the DC-coefficient up to the cut-off index. Finally, the inverse block DCT is computed and the shuffling is inversed to obtain the watermarked frame. On the detector side, the energy difference is computed and the embedded bit is determined according to the sign of the difference  $D$ . This algorithm has been further improved to adapt the cut-off index  $c$  to the frequency content of the considered  $n \times 8 \times 8$  block and so that the energy difference  $D$  is greater than a given threshold  $\tau_{\text{target}}$  [112].

### 3.3.2 Modify Motion Information

Another key element in video coding is motion estimation/compensation to reduce temporal redundancy. Indeed, successive video frames are highly similar and video coding basically aims at predicting one frame from another one using

motion prediction to reduce the amount of data to be transmitted. For instance, in the MPEG standard, there is a clear distinction between I frames which are encoded as still images and P or B frames which are respectively encoded in reference with one I frame and two other frames, either I or P. This results in a sequence of motion vectors which are transmitted to the decoder to perform motion compensation on the other side. It could be interesting to consider those motion vectors as potential candidates to carry a secret watermark. In this perspective, one can impose a parity rule on the components of the motion vector [89]. That is to say that, for example, the horizontal component of a motion vector is quantized to an even value if the bit to be hidden is equal to 0 and to an odd value otherwise. In the same fashion, for visibility reasons, one can also choose to consider only high magnitude motion vectors for embedding and to modify either the horizontal component or the vertical component of the motion vector according to its angle [202]. Alternatively, recent advances in digital watermarking with quantization schemes can also be considered for modifying motion information. In this perspective, motion vectors can be quantized with respect to a square grid or a circular grid or an angular grid [13, 14]. Such approaches have been demonstrated to be slightly more robust. But anyway, one of the major concern when motion information is modified is fidelity: it is very difficult to predict the perceptual impact of modifying motion vectors. Nevertheless, this issue may be not critical in some applications. For instance, motion information can be modified to perform partial encryption, also referred to as *waterscrambling* [15]. In this context, the goal is to degrade the video quality, but still enabling video content to be perceived by an end-user, giving him/her an idea of the original content to trigger an impulsive buying action.

### 3.3.3 Modify VLC Codewords

In many video encoders, transform domain coefficients and motion vectors are usually quantized, either with a scalar or a vector quantization. Next, the resulting information is represented with some symbols which are sent to an entropy encoder to obtain the final bitstream. For instance, in the MPEG standard, the quantized DCT coefficients are scanned in a zigzag order and represented with  $(run, level)$  tuples. The run is equal to the number of zeros preceding a coefficient and the level is equal to the value of the quantized coefficient. Those tuples are then input to an entropy encoder. In practice, some lookup tables are defined in the MPEG standard to associate a Variable Length Coded (VLC) codeword to each possible tuple. As a result, some researchers have investigated how to directly modify the bitstream i.e. the VLC codewords to avoid full compression and decompression which is time consuming. In this perspective, a pioneer work has identified a set of VLCs which can be modified without introducing strong visual artifacts [113]. This algorithm is further detailed below. Even if some variations

around this approach have been proposed [129], the most promising research track is the one which exploit recent works to make conventional VLCs exhibit resynchronization properties upon encountering bit errors [137, 138]. Such VLC are called reversible VLCS (RVLC) and are two-way decodable. The idea is then to use the error recovery power of such RVLCs to design reversible watermarking schemes: binary modifications due to the watermarking process are considered as channel errors and recovered.

**Table 3.2:** Example of lc-VLCs in Table B.14 of the MPEG-2 standard.

Variable length code	VLC size	Run	Level	LSB of Level
0010 0110 s	8+1	0	5	1
0010 0001 s	8+1	0	6	0
0000 0001 1101 s	12+1	0	8	0
0000 0001 1000 s	12+1	0	9	1
0000 0000 1101 0 s	13+1	0	12	0
0000 0000 1100 1 s	13+1	0	13	1
0000 0000 0111 11 s	14+1	0	16	0
0000 0000 0111 10 s	14+1	0	17	1
0000 0000 0011 101 s	15+1	1	10	0
0000 0000 0011 100 s	15+1	1	11	1
0000 0000 0001 0011 s	16+1	1	15	1
0000 0000 0001 0010 s	16+1	1	16	1

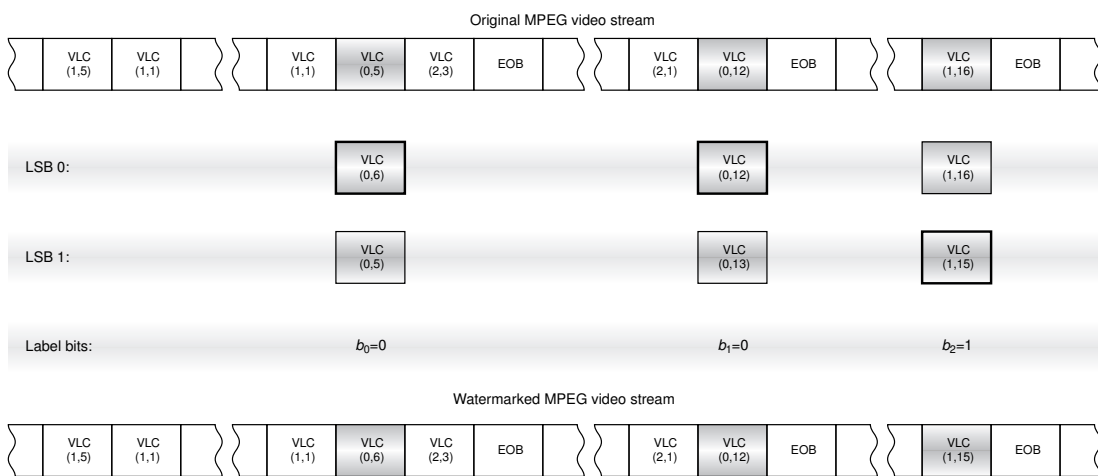
### Parity Bit Modification in the Bit Domain

Previous works have modified directly the bitstream of compressed video to embed a watermark. Such algorithms are very interesting because of the high achievable embedding rate and the low computational complexity. In the context of an MPEG video stream, a watermark consisting of  $l$  bits payload  $b_i$  ( $i = 0, 1, \dots, l - 1$ ) is embedded in the bitstream by selecting suitable VLCs and forcing the Least Significant Bit (LSB) of their *quantized level* to be equal to the payload bits [113]. To ensure that the change of VLC is perceptually invisible and that the size of the MPEG stream does not increase, only a few VLC called *label bit carrying VLC* (lc-VLC) are considered for embedding. Those VLCs have the interesting property that another VLC exists which has:

- the same run length,
- a level difference of 1,
- the same VLC length.



According to Tables B.14 and B.15 of the MPEG standards, there exists a lot of such lc-VLCS and a few examples are reported in Table 3.2. The symbol  $s$  represents the sign bit which indicates the sign of the DCT coefficient level. To insert the payload bits in an MPEG video bit-stream, all the VLCs in each macro block are tested, excepted the DC coefficients for visibility reasons. If an lc-VLC is found and the LSB of its level is unequal to the payload bit  $b_i$ , then this VLC is replaced by another one, whose level has the same LSB as the payload bit to be embedded. On the other hand, if the LSB of the original lc-VLC matches the payload bit  $b_i$ , the VLC is left untouched. This procedure is repeated until all payload bits are embedded. Figure 3.6 depicts the embedding process where 3 payload bits are inserted within a MPEG video stream. On the detector side,



**Figure 3.6:** VLC embedding process.

the payload bits are retrieved by testing all the VLCs in each macroblock. If an lc-VLC is found, the value of its LSB is retrieved and appended to the payload bitstream. The procedure is repeated until lc-VLCs are no longer found. Even if such algorithms are quite sensible against video editing, they are completely adapted for applications such as data hiding.



# 4

---

## Challenges

---


Digital video watermarking has been considered for along time as a simple extension of digital watermarking for still images. However, if watermarking still images and video content are similar problems, they are not identical. New issues, new challenges show up and have to be addressed. This chapter highlights a few one of them. In Section 4.1, it is reminded that the embedded watermark signal is assumed to resist a vast range of non-hostile video processing primitives. Next, the sensitivity to temporal desynchronization is pointed out in Section 4.2 and a few potential solutions are presented. Digital watermarking is often regarded as noise addition and efficient models have been constructed to predict the perceptual impact of noise addition on audio content and still images. However, as recalled in Section 4.3, very few works have investigated this issue in video. Another critical point when video watermarking is considered is that some applications may require real-time embedding and/or detection. Therefore, Section 4.4 briefly reviews a few techniques to achieve this goal. Finally, the security issue will be examined in depth in Section 4.5. Indeed, since each watermarked video frame can be considered as a single watermarked document, it opens avenues for collusion attacks.

### **4.1 Various Non Hostile Video Processing Primitives**

Robustness of digital watermarking has always been evaluated via the survival of the embedded watermark after attacks. Benchmarking tools have even been

developed to automate this process [22, 24, 147, 148, 173, 193]. In the context of video, the possibilities of attacking the video are multiplied. Many different non-hostile video processing primitives are indeed available. Non-hostile refers to the fact that even content providers are likely to process a bit their digital data to manage efficiently their resources.

**Photometric attacks.** This category gathers all the attacks which modify the pixel values in the frames. Those modifications can be due to a wide range of video processing primitives. Data transmission is likely to introduce some noise for example. Similarly, digital to analog and analog to digital conversions introduce some distortions in the video signal. Another common processing is to perform a gamma correction to increase the contrast. To reduce the storage needs, content owners often transcode, i.e. re-encode with a different compression ratio, their digital data. The induced loss of information is then susceptible to alter the performances of the watermarking algorithm. In the same fashion, customers are likely to convert their videos from a standard video format such as MPEG-1, MPEG-2 or MPEG-4 to a *popular* format e.g. DivX. Here again, the watermark signal is bound to undergo some kind of interferences. Spatial filtering inside each frame is often used to restore a low-quality video. Inter-frames filtering, i.e. filtering between adjacent frames of the video, has to be considered too. Finally, chrominance resampling (4:4:4, 4:2:2, 4:2:0) is a commonly used processing to reduce storage needs.



**Figure 4.1:** Example of distortion created by a handled camera (exaggerated).

**Spatial desynchronization.** Many watermarking algorithms rely on an implicit spatial synchronization between the embedder and the detector. A pixel at a given location in the frame is assumed to be associated with a given bit of the watermark. However, many non-hostile video operations introduce spatial desynchronization which may result in a drastic loss of performances of a watermarking scheme. The most common examples are

changes across display formats ( $4/3$ ,  $16/9$  and  $2.11/1$ ) and changes of spatial resolution (NTSC, PAL, SECAM and usual movies standards). Alternatively the pixel position is susceptible to jitter. In particular, positional jitter occurs for video over poor analog links e.g. broadcasting in a wireless environment. In the digital cinema context, distortions brought by the handheld camera can be considered as non-hostile since the purpose of the camera is not explicitly to remove the embedded watermark. It has been shown that the handheld camera attack can be separated into two geometrical distortions [40]: a bilinear transform, due to the misalignment between the camera and the cinema screen, and a curved transform, because of the lens deformations. This results in a curved-bilinear transform depicted in Figure 4.1 which can be modeled by twelve parameters.

**Temporal desynchronization.** Similarly temporal desynchronization may affect the watermark signal. For example, if the secret key for embedding is different for each frame, simple frame rate modification will make the detection algorithm fail. Since changing frame rate is a quite common processing, watermarks should be designed so that they survive such an operation. In the same fashion, frame dropping/insertion and frame decimation/duplication are also operations which are likely to affect the temporal synchronization of watermarking systems and they have consequently to be considered. Finally, frame transposition can disrupt the key schedule and thus result in desynchronization. Frame transposition has a similar effect to frame dropping with the difference being that this attack displaces the video frames in time instead of removing frames from the video.

**Video editing.** The very last kind of non-hostile attacks gathers all the operation that a video editor may perform. Cut-and-splice and cut-insert-splice are two very common processing primitives used during video editing. Cut-insert-splice is basically what happens when a commercial is inserted in the middle of a movie. Moreover, transition effects, like fade-and-dissolve or wipe-and-matte, can be used to smooth the transition between two scenes of the video. Such kind of editing can be seen as temporal editing in contrast to spatial editing. Spatial editing refers to the addition of a visual content in each frame of the video stream. This includes for example graphic overlay, e.g. logos or subtitles insertion, and video stream superimposition, like in the Picture-in-Picture technology. The detector sees such operations as a cropping of some part of the watermark. Such a severe attack is susceptible to induce a high degradation of the detection performances.

There are many various attacks to be considered as reminded in Table 4.1 and it may be useful to insert countermeasures [38] in the video stream to cope with the distortions introduced by such video operations. Moreover, the reader

**Table 4.1:** Examples of non-hostile video processing primitives.

Photometric	<ul style="list-style-type: none"> <li>- Noise addition, DA/AD conversion</li> <li>- Gamma correction</li> <li>- Transcoding and video format conversion</li> <li>- Intra and inter-frames filtering</li> <li>- Chrominance resampling (4:4:4, 4:2:2, 4:2:0)</li> </ul>
Spatial Desynchronization	<ul style="list-style-type: none"> <li>- Changes across display formats (4/3, 16/9, 2.11/1)</li> <li>- Changes of spatial resolution (NTSC, PAL, SECAM)</li> <li>- Positional jitter</li> <li>- Handled camera attack</li> </ul>
Temporal Desynchronization	<ul style="list-style-type: none"> <li>- Changes of frame rate</li> <li>- Frame dropping / insertion</li> <li>- Frame decimation / duplication</li> <li>- Frame transposition</li> </ul>
Video editing	<ul style="list-style-type: none"> <li>- Cut-and-splice and cut-insert-splice</li> <li>- Fade-and-dissolve and wipe-and-matte</li> <li>- Graphic overlay (subtitles, logo)</li> </ul>

should be aware that many other hostile attacks are likely to occur in the real world. Indeed, it is relatively easy today to process a whole movie thanks to the powerful available personal computers. It is virtually possible to do whatever transformation on a video stream. For example, for still images, StirMark introduces random local geometric distortions which succeed in trapping the synchronization of the detector. This software has been optimized for still images and, when used on each frame of the video stream, visible artifacts can be spotted when moving objects go through the fixed geometric distortion. However future versions of StirMark will surely address this visibility issue.

## 4.2 Temporal Synchronization

Despite the *purported* robustness of the watermark signal, temporal synchronization is still a critical issue in robust video watermarking detection, even in the absence of malicious attackers. Indeed, some applications places the watermarked content in conditions which are likely to damage the watermark signal or to confuse the detector. For instance, during streaming, network congestion may cause the watermark signal to be lost for an indeterminate amount of time i.e. many consecutive frames are dropped [198, 123]. If the watermark detector loses synchronization, it is necessary for the detector to resynchronize prior to resuming detection. In other words, the detector should know which watermark pattern to

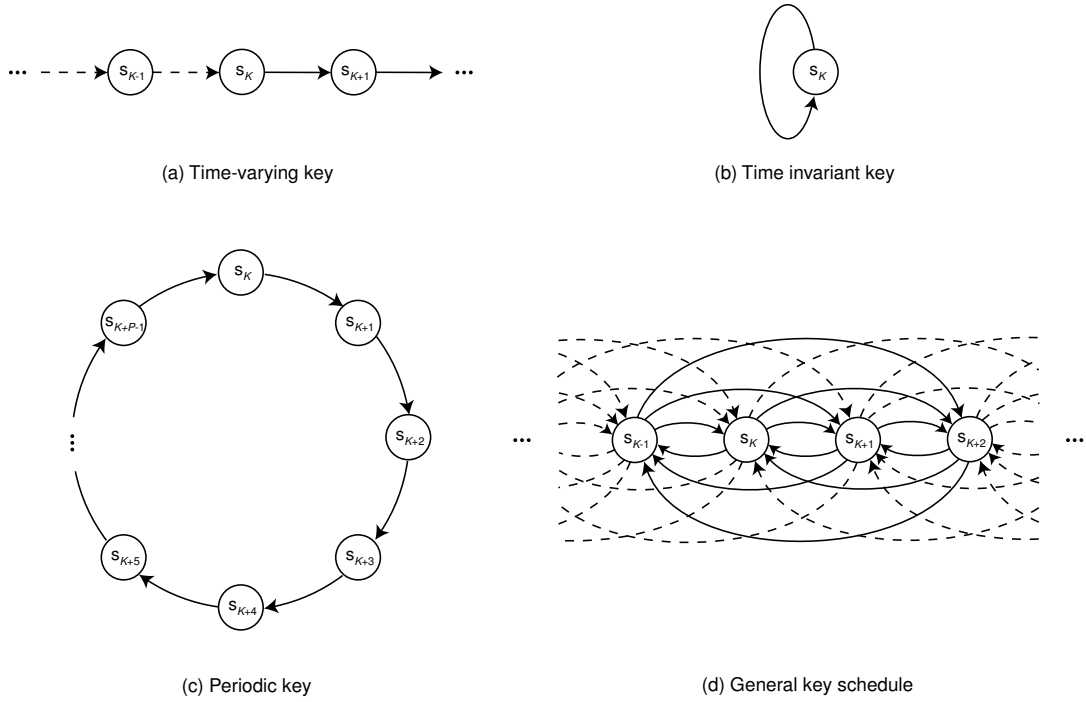
look for in each individual video frame. Another concern is initial synchronization, that is to say the ability of the detector to synchronize when the input signal first becomes available. The major challenge is that any portion of the watermarked video may be input to the detector, not necessarily the *beginning*. As a result, the detector should not rely on observing the beginning or any other specific part of the watermarked signal for resynchronization. The following subsections will present two alternative approaches to address this temporal synchronization issue. In Subsection 4.2.1, successive watermarks are related through a more or less complex key schedule whose knowledge is shared by both the embedder and the detector. Alternatively, in Subsection 4.2.2, temporal resynchronization is enabled by exploiting powerful image hashing functions.

### 4.2.1 Key Scheduling

When a watermark is embedded in a video frame, one of the parameters is the secret embedding key  $K$ . In video watermarking, successive keys used to generate the watermarks embedded in successive video frames can be considered as forming a key sequence  $\{K_t, K_{t+1}, \dots, K_{t+T}\}$ . This can be written:

$$\check{\mathbf{f}}_t = \mathbf{f}_t + \alpha \mathbf{w}_t = \mathbf{f}_t + \alpha \mathbf{w}(K_t) \quad (4.1)$$

where  $\mathbf{f}_t$  is the  $t$ -th original video frame and  $\check{\mathbf{f}}_t$  its watermarked version and  $\alpha$  some embedding strength. On the detector side, the goal is to know which key has been used to generate the embedded watermark to compute the proper correlation score. As a result, the idea is to impose a key schedule on the embedder side so that the detector can exploit this shared knowledge for resynchronization. Figure 4.2 depicts a few alternative key schedules, from the simplest to more elaborated ones. Always embedding the same watermark pattern, or also always using the same embedding key, is the best option against temporal desynchronization. Indeed the detector only has to check for the presence of the redundantly embedded watermark. However such a simple strategy introduces a weakness against estimation attacks (cf. Chapter 5). Alternatively, independent watermarks can be embedded in successive video frames. In this case, the detector is highly sensible to temporal desynchronization: a simple frame dropping breaks the chain of the key schedule. A possible countermeasure is to use a sliding correlator [79] but such a strategy does not scale. Instead, one can reinitialize the key schedule every  $P$  frames to obtain some periodic key sequences. A single frame dropping still disrupts the sequence of embedding key. However, the detector knows that if the video is watermarked a frame carrying the watermark generated with the initialization embedding key associated with  $s_K$  should show up in at most  $P$  frames. Unfortunately, temporal frequency analysis can reveal this periodic key schedule. All these examples are specific case of a generic state machine where all the states are connected i.e. from each state, the next state is determined according to the



**Figure 4.2:** Alternative key schedules for video watermarking. There exist a bijective function which associate to each state  $s_i$  an embedding key. It is assumed that the embedder starts in state  $s_K$  where  $K$  is the secret key.

secret key and possibly other parameters. For instance, features of the current video frame can be considered to determine the next state [122]. Furthermore, practical implementations have shown that introducing some redundancy in the key schedule can significantly enhance resynchronization performances. In other words, the same embedding key is used in several successive video frames before using the next one in the key schedule. This results in a trade-off between the randomness and the redundancy of the key schedule, which can be related with security against estimation attacks.

### 4.2.2 Frame Dependent Watermarking

Another approach to achieve temporal synchronization is to make the embedded watermarks dependent of some key features extracted from the video frames. This can be written:

$$\check{\mathbf{f}}_t = \mathbf{f}_t + \alpha \mathbf{w}_t = \mathbf{f}_t + \alpha \mathbf{w}(K, \mathbf{h}(\mathbf{f}_t)) \quad (4.2)$$

where  $\mathbf{h}(\mathbf{f}_t)$  are some robust features of the  $t$ -th video frame. This formula has to be carefully compared with Equation (4.1). On the detector side, assuming




that the considered features are robust enough, they can be extracted back from the input video frame and the watermark can be re-generated. As a result, the detector knows which watermark pattern is likely to be embedded and only has to check for its presence. Several alternative methods have been proposed to generate such frame dependent watermarks and most of them rely on efficient image hash functions [70, 84, 86, 106, 115, 135, 140, 180, 187]. The difference between cryptographic hash functions and image hash functions is related with their sensitivity with respect to the input image. If two images differ by a single bit, a cryptographic hash function will output two uncorrelated binary digests. On the contrary, an image hash function will output the same binary digest. In fact, the output digests are the same when the frames are *similar* whereas they are uncorrelated when the frames are completely different. Once such binary strings have been produced, the watermark is generated so that it degrades gracefully with an increased number of bit errors. It should be noted that such a strategy ensures that two video frames carry watermarks which are as correlated as the host video frames. Image hash function can also be exploited in some applications where the detector is allowed to have access to some side information e.g. the binary digests of some video frames with their associated timestamp [77, 178]. On the other side, the detector computes the binary hash of all the frames, compares them with the available side information and compensate for possible temporal distortions. Then, the usual watermark detection process can be launched.


### 4.3 Distortion Evaluation

Evaluating the impact of distorting a signal as perceived by a human being is a great challenge. The amount and perceptibility of distortions, such as those introduced by lossy compression or digital watermarking, are indeed tightly related to the actual signal content. For instance, when considering still images, it is commonly admitted that the human eye is more sensitive to changes in low frequencies than to changes in high frequencies. Furthermore, it is also recognized that bright areas can be changed by a larger amount without being noticed. This property is referred to as luminance masking. In the same fashion, contrast masking can also be considered. This is related with the reduction in visibility of a change in one frequency due to the energy present in that frequency. Those characteristics of the human eye can then be gathered to design efficient human vision models [194] which can be used to perform perceptual shaping to ensure watermark imperceptibility [32]. Similarly, masking is also relevant in audio and can be considered to obtain a model which provides a measure of the threshold of hearing in the presence of a given audio signal.

In the context of video, the Video Quality Experts Group (VQEG) [188] was formed in 1997 to devise objective methods for predicting video image quality



(a) Spatio-temporal contrast sensitivity [96].



(b) Information in natural images sequences vs. visual sensitivity as a function of temporal frequency [185].

**Figure 4.3:** Human Visual System and spatiotemporal signals.

after compression. In 1999, they stated that no objective measurement system at test was able to replace subjective testing and that no objective model outperforms the others in all cases. This explains while the Peak Signal to Noise Ratio (PSNR) is still the most often used metric today to evaluate the visibility of video watermarks. Nevertheless, earlier studies in human vision [96, 185] can be of interest to predict the impact of noise addition in a spatiotemporal video signal. In particular, the human spatiotemporal visual contrast sensitivity function reproduced in Figure 4.3 clearly exhibits a substantial dip at the low spatial, low temporal corner of the plot. Such a disparity of sensitivity can be useful in discovering robust watermark carriers which remain invisible to human beings. This analysis is corroborated by further studies which modeled the spatiotemporal filters of mammalian vision, based on the spatial and temporal statistics of natural images and on an optimization that maximizes the flow of information through noisy channels of limited dynamic range [185]. In the results reported in Figure 4.3, information is maximal at the lowest spatial and temporal frequencies even if the optimal filter reduces sensitivity drastically to these frequencies. This

discrepancy between information and sensitivity can motivate the modulation of low frequencies for watermark embedding. Indeed even if the sensitivity of the human visual system is reduced at low spatiotemporal frequencies, the high degree of information in these frequencies makes them difficult to distort without degrading the quality. In fact, in practice, most processes that are applied to moving pictures and result in *watchable* quality tend to reproduce these low frequency / high information components with high fidelity.

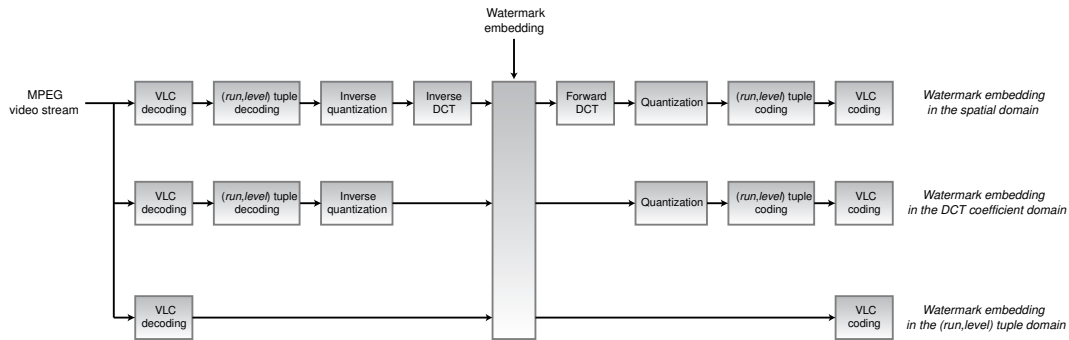
Contrarily to still images and audio, such observations have not led to well-known metrics or widespread perceptual shaping procedures. Hence, developing more sophisticated models which also include temporal masking remains an open research problem. However, from a subjective point of view, previous works [131, 197] have isolated two kinds of impairments which appear in video, when the embedding strength is increased:

1. *Temporal flicker*: Embedding uncorrelated watermarks in successive video frames usually results in annoying twinkle or flicker artifacts similar to the existing ones in video compression,
2. *Stationary pattern*: Embedding the same watermark pattern in all the video frames is visually disturbing since it gives the feeling that the scene has been filmed with a dirty camera when it pans across the movie set.

With this lack of objective perceptual metric for video, visual quality evaluation usually comes down to a subjective examination of the watermarked material, possibly by *golden eyes*.

## 4.4 Real-time Watermarking

Real-time can be an additional specification for video watermarking. As a matter of fact, it is not really a big concern with still images. When a user wants to embed a watermark or to check the presence of a watermark in an image, a few seconds is an acceptable delay. However, such a delay becomes unrealistic in the context of video. Individual video frames should indeed be processed at a fairly high rate, typically 25 frames per second, to obtain a temporally smooth video stream. Therefore, for specific applications, it may be required that at least the embedder or the detector, and even sometimes both of them, is able to handle such a rate. For instance, in the context of broadcast monitoring, the detector should be able to detect an embedded watermark in real-time. Alternatively, in a VoD environment, the video server should be able to insert the fingerprint watermark identifying the customer at the same rate as the one the video is streamed. As a result, different approaches will be surveyed in the next subsections as possible means to meet this real-time specification.



**Figure 4.4:** Embedding strategies with MPEG video streams.

#### 4.4.1 Low-Complexity Algorithms

A possible way to achieve real-time watermarking is obviously to reduce the complexity of the embedder and/or of the detector as much as possible. For convenience, video content is usually lossy compressed for storage or transmission. Such compression algorithms typically combine motion estimation/compensation, signal transformation - such as Discrete Cosine Transform (DCT), Discrete Fourier Transform (DFT) or Discrete Wavelet Transform (DWT) -, quantization, lossless coding and entropy coding. Subsequently, the embedding/detection process can be introduced at any location in the video codec. The different possible alternatives for a MPEG video stream have been depicted in Figure 4.4. If it is not possible to embed the watermark directly into the MPEG binary stream, one should try to avoid as many operations as possible and especially computationally expensive computations such as full decompression and recompression. Such a philosophy can lead to very efficient watermarking algorithms when the specificities of a given compression codec are properly exploited. For instance, with the MPEG standard, after variable length decoding, the video stream basically consists of a succession of  $(run, level)$  tuples. A very simple watermarking strategy could then be to directly modify these tuples [113]. In fact, the MPEG standard clearly defines some lookup tables between  $(run, level)$  tuples and Variable Length Coded (VLC) codewords. When examining closely these tables it is possible to isolate similar VLC codewords i.e. codewords with the same number of bits which encode tuples sharing the same run length but differing by a quantized level difference of one. As a result, such codewords can be used alternatively to hide one bit, without increasing the length of the binary stream and without introducing strong visual distortions.

### 4.4.2 Reduce the Rate of Expensive Computations

When a system designer tries to reduce the complexity of a video watermarking system, the resulting algorithm is usually heavily linked with the considered video codec. In the previous subsection for instance, the presented algorithm is strongly related with the MPEG standard and the DCT. As a result, such watermarks are likely not to survive very simple signal processing primitives such as transcoding with a video codec which uses another transform. This explains why elaborated algorithms are still investigated today. For instance, in Section 3.2.2, the JAWS system designed by Philips for video broadcast monitoring has been presented [93]. It operates in the spatial domain and even requires very complex computations for detection such as Fast Fourier Transform (FFT). In this case, how can they claim to achieve real-time detection which is a critical requirement for broadcast monitoring? Well, several tricks have been exploited to cut down the computational complexity of the detector which looks at first sight very expensive. The first ruse consists in computing simple operations at video rate and complex ones at much lower rate. In this perspective, the detection procedure only consists of additions (video frames folding) and a few FFT. This strategy has been demonstrated to enable real-time detection over analog links. When MPEG video streams are considered, the complexity of the detector seems to increase since the stream has now to be decoded. However, one can still reduce the computational cost down to the point where real-time detection is possible [94]. In this context, the first key observation is that IDCT transformation and folding can commute. Indeed, because of the linearity of the IDCT operator, the result of first applying IDCT to a number of DCT blocks and then adding them up is the same as first block addition followed by a single IDCT applied to the summed blocks. Hence, the number of IDCT computations is decreased. Next, it has been experimentally verified that a large portion of JAWS watermark energy is concentrated in residual frames i.e. in displaced frame differences without motion compensation. As a result, memory complexity associated with motion compensation in the MPEG decoder can be deleted. In summary, real-time detection for JAWS watermarks has been made possible by reducing the rate of complex computations such as IDCT and FFT.


### 4.4.3 Preprocessing to Facilitate Watermark Embedding

A major concern with cheap watermark embedders is that they usually favor simple algorithms which do not ensure that performances will remain the same for different cover works. For instance, an embedder might embed a high-fidelity, robust watermark in one video while completely failing to embed it in another video. For illustration, let us consider a zero-bit watermarking system which relies on a detector computing the linear correlation between the cover work  $\mathbf{c}$

and a pseudo-random pattern  $\mathbf{w}_o$ :

$$\text{lc}(\mathbf{c}, \mathbf{w}_o) = \frac{1}{N} \mathbf{c} \cdot \mathbf{w}_o = \frac{1}{N} \sum_{i=1}^N \mathbf{c}(i) \mathbf{w}_o(i) \quad (4.3)$$

If this detection score is greater than a threshold  $\tau_{\text{detect}}$ , then the detector reports that a watermark has been embedded. From a geometric perspective, this comes down to defining a plane which divides the media space in two regions. Depending on the region where the considered media content lies, it will be considered as watermarked or not.



(a) Blind watermarking: original contents are simply moved in direction  $\mathbf{w}_o$  with a fixed strength. The system ensures a fixed distortion but the resulting contents can be watermarked or not.

(b) Informed watermarking: original contents are moved in direction  $\mathbf{w}_o$  so that the resulting contents are at fixed distance inside the detection region. The system ensures 100% embedding effectiveness at the cost of variable distortion.

(c) Preprocessing before embedding: original contents are moved in direction  $\mathbf{w}_o$  to a prepping region close to but outside the detection region. Then simple blind embedding ensures same robustness performances as informed embedding.

**Figure 4.5:** Geometric interpretation of different embedding strategies when a linear correlation detector is considered. The vertical line denotes the detection boundary in the media space. Empty (resp. plain) circles denote unwatermarked (resp. watermarked) contents [34].

The simplest embedding strategy for such a system is blind embedding as depicted in Figure 4.5(a). It basically consists in adding the watermark pattern  $\mathbf{w}_o$  to the original host content with a fixed embedding strength so that the embedding process always introduces the same distortion. This is computationally easy and the watermark can be added for example directly to the video stream (in baseband) without requiring frame buffering. Nevertheless, such a simple system will fail to embed the watermark for some cover works, which make it unacceptable for many applications where the watermark *must* be embedded. Alternatively, informed embedding can ensure 100% embedding effectiveness. As depicted in Figure 4.5(b), such systems examine the cover work before embedding and adjust the embedding strength to make sure that the watermark is effectively embedded. In other words, constant robustness is guaranteed at the

cost of variable fidelity. Unfortunately, examining video content before embedding is increasing the computational complexity. Therefore, a two step process can be considered which basically combines informed preprocessing and blind embedding [34]. As depicted in Figure 4.5(c), the first step is more costly and is usually performed by content providers. The goal is to modify the original content so that subsequent blind embedding is successful. For instance, in the studied linear correlation detection system, host interference can be canceled i.e. the preprocessing step removes any correlation with the reference pattern  $\mathbf{w}_o$ . The second operation is in comparison far more simple. Thus, VoD servers and DVD recorders for example can do it in real-time. In summary, the philosophy is to split the computational load between content providers and devices which have to process data in real-time.

## 4.5 Security Issue and Collusion Attacks

Digital watermarking has always been regarded as a security related technology. Nevertheless, it is not really clear what the term *security* refers to. At the very beginning, this was kind of connected with the fact that watermarking embedding and detection processes are made dependent of a secret key. A direct analogy has then been drawn with cryptography and Kerckhoffs' principles to ensure security have been considered [97]. They state for example that even if the system under study is publicly known, it should not be broken down as long as the secret key is not disclosed. However, whereas *breaking down the system* means obtaining the plain text in cryptography, it might mean several different things in digital watermarking. For instance, unauthorized users should not be able to remove, detect, estimate, write or modify embedded watermarks [92]. But if extensive benchmarking is now performed to assert whether or not a watermarking system can cope with standard signal processing primitives, almost no work has been done to evaluate security. Thus, the remainder of this section will try to give some elements to define somewhat clearly what security is. First of all, the relationship between security and the need for trust in a hostile environment is exhibited in Subsection 4.5.1. Then, security is opposed to robustness in Subsection 4.5.2 to draw a line, even fuzzy, between both concepts. Next, Subsection 4.5.3 reminds that absolute security is not required in real life. In fact, even limited security can be valuable when the risk is properly managed. Finally, collusion attacks are introduced in Subsection 4.5.4 as a possible way to evaluate security.

### 4.5.1 Trust in a Hostile Environment

In many watermarking applications, there is a need to trust the information which is conveyed by the watermarking channel. When fingerprinting watermarks are

exploited to trace back malicious customers who have broken their license agreement, content owners basically want to rely on the information extracted from the embedded watermark. Thus, in such a scenario, it should not be possible to produce multimedia contents carrying fake valid watermarks to prevent innocent consumers from being framed. Additionally, since fingerprinting watermarks can be used to identify the source of leakage in a content distribution system, they are likely to be attacked and they should consequently survive as many removal attacks as possible. On the other hand, it is not critical if the attacker succeeds in detecting/reading the watermark. However this last point is no longer valid when digital watermarking is used for steganography. In this case, the watermarking channel is basically exploited to establish a covert communication channel between two parties whose existence remains unknown for other people. Therefore, the presence of a hidden watermark should not be even detected if the secret key is not known. Alternatively, in an authentication perspective, unauthorized watermark removal is not really important: digital content will be seen as non valid and discarded. In a completely different strategy, digital watermarks can be used to insert useful information into the cover data e.g. annotation watermarks, error recovery watermarks, metadata binding watermarks... In such cases, altering the watermark is likely to remove the associated additional information or service. In other words, consumers have no longer interest to eliminate embedded watermarks and hostile behaviors disappear. The environment is collaborative instead of hostile and security requirements are now superfluous.

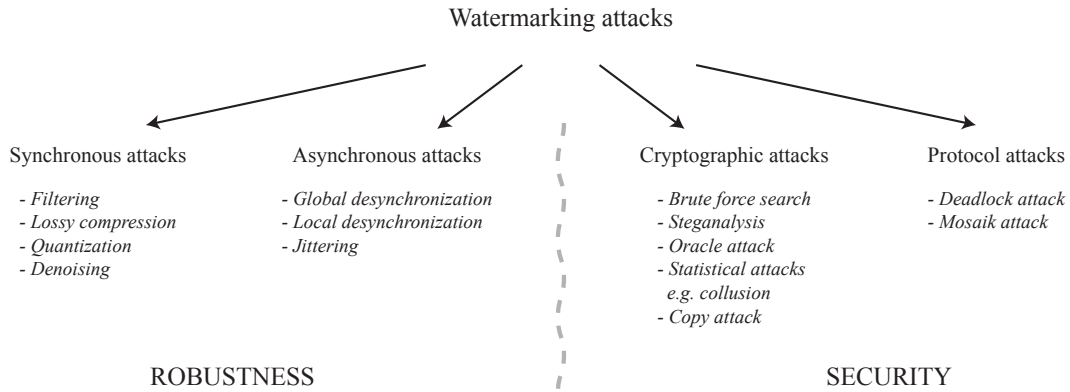
In summary, the environment in which the watermarking technology is going to be released has to be carefully studied. Depending on the targeted applications, customers will not judge digital watermarking the same way and will consequently adopt different attitudes. Generally speaking, the more the embedded watermarks *disturb* the customers, the more hostile will be their behaviors and the higher will the security specifications need to be raised. Hence, security issues basically arise when two conditions are met. On one side, content providers value the information conveyed by the watermarking channel and expect the embedded watermark to survive to provide some service e.g. copy control to prevent copying multimedia content without paying royalties, traitor tracing to identify the source of leakage in multimedia content distribution systems... On the other hand, customers see the watermarking signal as a disturbing technology and deploy highly hostile strategies to defeat the protection system. In other words, the notion of security is inherently tied with the need for trust in a hostile environment. This is a key aspect to consider when a watermarking system is under study. In particular, IP protection related applications should identify which operations are critical or not in their framework. Nevertheless, it should be reminded that a large range of applications using digital watermarking do not have security specifications at all e.g. when embedded watermarks are used to convey some useful information.



### 4.5.2 Security versus Robustness

This subsection aims at drawing a line, even fuzzy, between two very important concepts in digital watermarking: security and robustness. These notions have indeed been mixed up for a long time, even in the watermarking community itself. The first noticeable difference is that security implicitly assumes a hostile environment, as reminded in the previous subsection, which is not the case for robustness. Robustness addresses the impact of regular signal processing primitives on watermarked multimedia content. Typically, a usual customer who is compressing multimedia content for storage/transmission convenience is not regarded as a security threat even if lossy compression is likely to increase the probability of missing an embedded watermark or the probability of retrieving a message with bit errors. The main point is that the customer does not intend to remove the watermark. In summary, one can say that robustness is only concerned by regular customers while security cares more specifically about hackers whose well-thought-out goal is to defeat the system. A second distinction is that security attacks usually aims at gaining some knowledge on the protection system. There is a clear gap between a customer who is blindly JPEG compressing an image and a hostile attacker who is collecting a whole database of watermarked documents to check whether some information leaks about the watermarking process or not. This knowledge can then be exploited to detect, remove or edit the embedded watermarks. A direct consequence is that robustness attacks are usually more generic/universal than security ones. Finally, a last dissimilitude between security and robustness is that the later one is only concerned by watermark removal whereas the first one also cares about unauthorized watermark detection, estimation, modification and writing. In summary, robustness deals with common consumers which perform regular signal processing operations which are likely to degrade the performances of watermark detection even if it is not their original goal. On the other hand, security handles hostile hackers who attack the protection system on purpose to gain first some knowledge about it and then to design dedicated attacks which are not limited to watermark removal.

Keeping these observations in mind, usual attacks which are relevant in digital watermarking have been separated as depicted in Figure 4.6 depending on whether they are a matter of security or robustness. This separation is basically an extension of previous classifications proposed in the watermarking literature [151, 192]. On the robustness side, attacks against digital watermarks have been split into synchronous and asynchronous attacks. Synchronous attacks refer to common signal processing primitives such as filtering, lossy compression, quantization, denoising which are likely to directly affect the watermark signal and thus interfere with the watermark detector. On the other hand, asynchronous attacks include all the operations which perturb the signal sample positions. As a result, the synchronization convention shared by the embedder and the decoder



**Figure 4.6:** Watermarking attacks breakdown depending on whether they address robustness or security issues.

becomes obsolete. Hence, such attacks do not explicitly remove the watermark signal but still, they are considered as watermark removal attacks since the detector is no longer able to retrieve the watermark [92, 173]. A very well known example is the Random Bending Attack (RBA) [151] which is now a reference attack to evaluate robustness. It basically simulates the effects of D-A/A-D conversion: geometrical distortions are introduced to take into account lens parameters and imperfect sensors alignment, some noise is added to mimic the response of non ideal sensors and a mild JPEG compression is performed for storage convenience.

On the security side, watermarking attacks are divided into protocol and cryptographic attacks. The first set of attacks exploits some general knowledge on the watermarking framework. For instance, in a copyright protection perspective, if a multimedia item is found to carry more than a single watermark, many schemes do not provide an intrinsic way of detecting which of the watermarks was embedded first [36]. It is a deadlock and nobody can claim the ownership of the document. Alternatively, automated search robots can be used to browse the Internet and check if web sites host illegally copyrighted material. A simple way to circumvent such web crawlers is to split the images into many small pieces and to embed them in a suitable sequence in a web page [1]. The juxtaposed images appear stuck during rendering, that is to say as the original image, but the watermark detector is confused. On the other hand, cryptographic attacks aim at gaining some knowledge about the watermark signal itself i.e. the secret key or the utilized pseudo-random sequence. Brute force search basically tries all the possible keys until the correct one is found. Steganalysis objective is to isolate some characteristics of watermarking methods to enable non authorized watermark detection [23]. In another fashion, the Oracle attack uses publicly available detectors to iteratively modify watermarked contents until the detector

fails to retrieve the embedded watermark [126]. On their side, statistical attacks, also referred to as collusion attacks, collect several watermarked documents and combine them to obtain non watermarked documents. This attacking strategy will be further examined in the next part of this thesis. Finally, to the best knowledge of the author, the only example of unauthorized watermark writing is the copy attack [110, 83]: the watermark is estimated from a watermarked document and successfully inserted back into a non-protected one.

### 4.5.3 Security in the Real World

Nowadays, it is commonly admitted that a perfectly secure system does not exist. If a motivated hacker has no limit of time, computing resources and money, he/she will succeed in defeating the protection system, for instance with a brute force key search approach. This is also true for digital watermarking. Does it mean that security is useless? Not at all! Customers value even limited forms of content protection. Let us for instance examine the case of copy protection for Digital Versatile Disk (DVD) distribution. When content owners release a new movie, they know that most of the sales are going to be done within the first months. As a result, they basically want the copy protection mechanism to last a few months and they do not really care if a pirate hacks the protection after one year and release it on the Internet. If the protection technology lasts long enough, customers who are eager to consume new released products will not wait until a pirated copy appears on Peer-to-Peer (P2P) networks. A second important point is that just because a technology can be circumvented does not necessarily implies that customers will effectively do it [3, 4]. For example, if the case study of DVD is continued, the proposed copy/playback protection basically divides devices into compliant DVD players which cannot read illegal DVDs and noncompliant DVD players which cannot read legal DVDs [12]. The expense of owning two DVD players can then be exploited to help *keep honest people honest*. In summary, in real life, all that matters is that the cost of breaking the system (complexity, time, money...) should be higher than the cost of doing things legally.

Coming back to down-to-earth considerations, money has also to be taken into account. Who will pay to introduce a secure protection system? This is a big issue and also the point where the different concerned parties usually disagree. There are typically three actors: content owners (cinema studios, music majors...), consumer electronics manufacturers and consumers associations. Content owners want to protect their high valuable multimedia items once they are released to the public but they are most of the time not ready to bear all the cost of the protection system. From the manufacturer point of view, more security means more hardware or software, more expensive devices and consequently less sales and lower profits. And of course, consumers are not really enthusiastic about the idea of paying for some security mechanism which is going to restrict


the possible usages of multimedia data. Such conflicts of interest can come to a dead end. A typical example is the introduction of digital watermarks inside DVDs which has been almost abandoned. Thus, the efficiency of the protection technology has to be balanced with the economic interests at stake. Once again insecure technologies may be worth being deployed if the risk is properly managed [104]. Although credit card networks are based on insecure magnetic stripe technology, risk management tools have been able to maintain fraud rates below 0.1% of transaction volume.

Nevertheless, even with an efficient risk management strategy, it is useful to investigate how to obtain secure watermarks. The main issue here is that watermarked contents cannot be updated once released as the antivirus softwares are when a new threat appears. Watermark embedders and detectors can of course be improved but the already released items will not benefit from these enhancements. In fact, as soon as the watermark is removed from a multimedia item, this one is likely to be broadcasted on P2P networks. Today, a single leak on the Internet and everything is done. As a result it is still relevant to anticipate hostile behaviors to iteratively propose more secure algorithms i.e. a longer time for removing the watermark once the protected item is released to the public. In another fashion, research is also conducted to design systems, such as conditional access for multimedia players [101], where a defeated entity does not compromise the security of the whole system.

#### 4.5.4 Collusion Attacks

Collusion is a well-known attacking strategy which basically refers to a set of malicious customers who gather their individual knowledge about the protection system, whatever it is, to obtain unprotected multimedia content. It has first been mentioned in cryptography when protocols have been established to part a secret between different individuals. Typical examples include secret sharing, also referred to as threshold cryptography, and conference keying [133]. The idea of secret sharing is to start with a secret and then to divide it into pieces called *shares* which are distributed amongst users such that the pooled shares of specific users allow reconstruction of the original secret. These schemes can be exploited to enable shared control for critical actions. Vault deposit accounts are a good illustration of such a procedure. Both the customer key and the bank manager key are required to grant access to the account. If any part of the secret (key) is missing, the door of the vault remains closed. This idea can be extended to more than two people. Access to a top secret laboratory can for instance be controlled by access badges: admittance necessitates a security guard badge and a researcher badge. Since there are many researchers and guards in the lab, this results in two groups of badges and one badge from each group is required to enter the lab. From a more general perspective, secret sharing

split knowledge between individuals so that at least  $u$  users have to merge their shares to retrieve the secret knowledge. In such a framework, colluders are  $c$  users who try to build fake valid shares or even to reconstruct the secret despite the fact that  $c < u$ . On the other side, conference keying is slightly different. Whereas secret sharing can be seen as a key pre-distribution technique wherein the recovered secret is static and usually the same for all groups, conference keying protocols allow to have *session keys* which are different for different groups and which dynamically adapt to the individuals in the group. These protocols are particularly interesting for applications which need secure group communications such as telephone/video conferences, interactive meetings and Video on Demand (VoD) [67]. Most concerns come from the need to manage members joining or leaving groups, which has an impact on session keys. In such scenarios, the goal of the colluders is to create some new keys to join the sessions without paying the fee. In summary, collusion has already been studied in cryptography. The riposte which has been introduced to circumvent such behaviors is basically a dissuasive weapon. Distributed keys are build in such a way that, if some colluders combine several of them to produce an illegal key, this one contains some evidence of the pirate identities which can be used to take legal measures [27, 28]. Once there is a threat of being caught, there are far more less candidates for collusion.



**Figure 4.7:** Collusion in watermarking: Colluders collect several watermarked documents and combine them to produce digital content without underlying watermarks.

In digital watermarking, collusion attacks were first mentioned in the context of fingerprinting [200]. In such applications, content owners want to distribute high valued items to a large audience. However, they are concerned about their copyright and want to be able to trace illegal content back to the source of leakage. To this end, instead of selling the same item to all the customers, they assign slightly different copies to each customer. As a result, each customer owns a *unique* copy carrying its own imperceptible and randomly located tracers.

Thus, if a customer decides to make his/her copy freely available on the Internet, the content owner is able to find the identity of the traitor using these secret markers. In this case, colluders will typically compare several marked documents to identify where the secret markers are located and then remove them. Now, in terms of digital watermarking, collusion consists in collecting several watermarked documents and in applying a process which succeeds in producing unwatermarked content as depicted in Figure 4.7. Traditionally, two alternative approaches can be enforced: combining watermarked documents can either aim at estimating directly the original unwatermarked content or in estimating some properties of the watermark signal which can be exploited to remove the embedded watermark in a second step. Solutions have already been proposed in the literature. Secure codes can for instance be used to prevent the watermark to be removed when multiple customers average their protected items [16].

Nevertheless, when video content is considered, the situation is significantly more challenging. Each frame of the video can indeed be seen as an individual watermarked document. This approach is all the more pertinent since many video watermarking schemes enforce a frame-by-frame embedding strategy [48]. An attacker can consequently collect multiple video frames and combine them to produce unwatermarked video frames. In this perspective, early studies have exhibited two main collusion strategies [82, 174]. When uncorrelated host video frames are studied, the goal is to isolate some hidden structure of the watermark signal. On the other hand, when similar video frames are collected, the objective is to fuse these frames so that the embedded watermark is no longer detectable. Both approaches will be further developed in the next part of the thesis.

## Part II

# Security Issue and Collusion Attacks





---

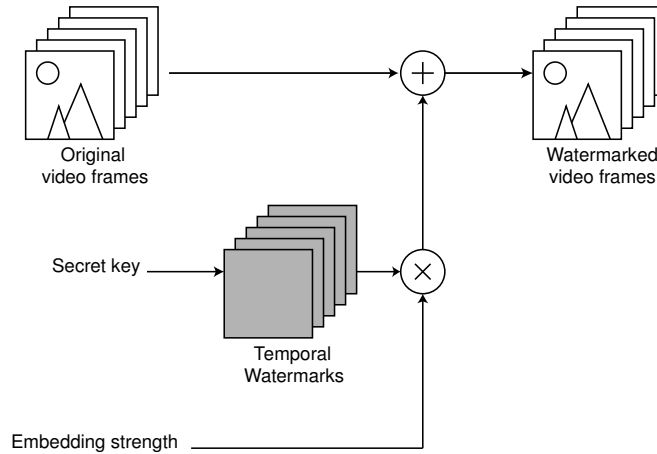
## Eavesdropping the Watermarking Channel

---

Video watermarking is usually considered as a simple extension of digital watermarking for still images. Since a video sequence is made of several video frames, a simple way to watermark the video is indeed to reuse previously designed watermarking algorithms for still image. Such approaches are possible as long as the watermarked content is not released in a hostile environment with attackers willing to remove the embedded signal. In this case, security issues have to be addressed as it has been previously highlighted in Section 4.5. In this chapter, collusion attacks will be surveyed to exhibit security pitfalls of frame-by-frame approaches to video watermarking. To begin with, two reference video watermarking algorithms are introduced in Section 5.1 as well as two simple collusion attacks. Next, both watermark modulation and watermarking strength modulation are investigated to provide higher performances against collusion attacks. Nevertheless, it will also be shown that a smart attacker is still able to confuse the system with elaborated estimate-and-remodulate attacking strategies. Although all this study is illustrated using additive Spread-Spectrum (SS) watermarks, Section 5.4 rapidly show that such weaknesses against collusion attacks are also an issue with Scalar Costa Schemes (SCS).

### 5.1 Baseline Framework

The next subsections describe in details the baseline framework which will be considered in the remainder of this chapter. In Subsection 5.1.1, two very basic



**Figure 5.1:** SS system: A different watermark is embedded in each video frame.

frame-by-frame embedding strategies are introduced. Next, two simple collusion attacks are described in Subsection 5.1.2 as possible means to defeat the previously presented watermarking schemes. This baseline will then be used as a starting point to propose possible enhancements for the watermarking system as well as more elaborated collusion attacks in an iterative fashion.

### 5.1.1 Frame-by-Frame Watermarking

As reminded in Chapter 3, some video watermarking algorithms exploit the specificities of a compression standard. Others embed a watermark in a three dimensional transform. However, watermarking digital video is mostly considered today as watermarking a sequence of still images [48]. Once this approach is enforced, two major embedding strategies are used: either a *different watermark* is inserted in each video frame or the *same watermark* is embedded in all the video frames. For sake of simplicity, both strategies will be illustrated with an additive watermark based on the Spread Spectrum (SS) theory [153] in the next subsections.

#### SS System

In the pioneering spread spectrum based video watermarking technique [79], video was considered as a one-dimensional signal. From a frame-by-frame perspective, this can be seen as a system which *always embeds a different watermark* as depicted in Figure 5.1. In such a *SS system*, the embedder inserts a pseudo-random watermark in each video frame:

$$\check{\mathbf{f}}_t = \mathbf{f}_t + \alpha \mathbf{w}_t(K), \quad \mathbf{w}_t(K) \sim \mathcal{N}(0, 1) \quad (5.1)$$

where  $\mathbf{f}_t$  is the luminance of the  $t$ -th video frame,  $\check{\mathbf{f}}_t$  the luminance of the  $t$ -th watermarked frame,  $\alpha$  the embedding strength and  $K$  a secret key. The inserted watermark  $\mathbf{w}_t(K)$  has a normal distribution with zero mean and unit variance and is different at every instant  $t$ . Using  $K + t$  as a seed for the pseudo-random generator is a simple way to obtain this property. Perceptual shaping can be introduced to improve the invisibility of the watermark even if a global embedding strength has been used in practice. From a subjective point of view, always changing the embedded watermark introduces an annoying flicker artifact [131].

The detector computes then the following correlation score:

$$\rho(\{\check{\mathbf{f}}_t\}) = \frac{1}{T} \sum_{t=1}^T \check{\mathbf{f}}_t \cdot \mathbf{w}_t = \alpha + \frac{1}{T} \sum_{t=1}^T \mathbf{f}_t \cdot \mathbf{w}_t \approx \alpha \quad (5.2)$$

where  $T$  is the number of considered video frames and  $\cdot$  denotes the linear correlation operation. This score should be equal to  $\alpha$  if a watermark is present in the video, while it should be almost equal to zero if no watermark has been inserted. Moreover, host interference can be cancelled in a preprocessing step [34] to enhance the detection statistics. As a result, the computed score is compared to a threshold  $\tau_{\text{detect}}$  to assert the presence or absence of the watermark. The value  $\alpha/2$  has been chosen in practice to obtain equal false positive and false negative probabilities<sup>1</sup>.

### SS-1 System

The SS system is highly sensitive to temporal desynchronization. A simple frame drop or insertion succeeds in confusing the detector. The alternative *SS-1 system* depicted in Figure 5.2 has consequently been introduced. It basically *always embeds the same watermark* [93]. In other words, the embedder redundantly inserts the same pseudo-random watermark in each video frame:

$$\check{\mathbf{f}}_t = \mathbf{f}_t + \alpha \mathbf{w}(K), \quad \mathbf{w}(K) \sim \mathcal{N}(0, 1) \quad (5.3)$$

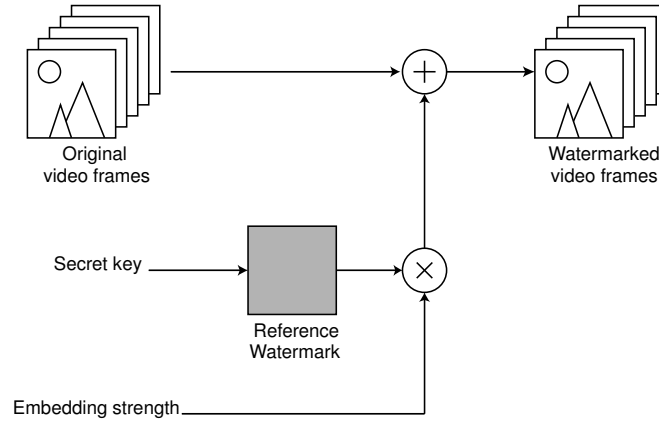
where  $\mathbf{w}(K)$  is a key-dependent reference watermark. From a subjective perspective, this embedding strategy produces an annoying persistent pattern [131] when the camera moves.

On the detector side, the correlation score defined in Equation (5.2) is computed. Now that the same watermark is embedded in each video frame, the linearity of the operator  $\cdot$  can be exploited to reduce the number of computations [93] required for detection:

$$\rho(\{\check{\mathbf{f}}_t\}) = \frac{1}{T} \sum_{t=1}^T \check{\mathbf{f}}_t \cdot \mathbf{w} = \left( \frac{1}{T} \sum_{t=1}^T \check{\mathbf{f}}_t \right) \cdot \mathbf{w} \quad (5.4)$$

---

<sup>1</sup>Adding some noise to the watermarked video introduces an interfering term in (5.2), which has zero mean and a variance proportional to  $1/\sqrt{T}$ . In other words, modifying  $T$  enables to adjust the false positive and false negative probabilities.



**Figure 5.2:** SS-1 system: The same reference watermark is redundantly embedded in each video frame.

This means that averaging several correlations between different video frames and the same watermark is equivalent to computing a single correlation between the average of the video frames and this watermark. Here again, the correlation score should be equal to  $\alpha$  if a watermark is present in the video, while it should be almost equal to zero if no watermark has been inserted. As a result, the computed score is compared to a threshold  $\tau_{\text{detect}}$ , which is set equal to  $\alpha/2$  in practice to obtain equal false positive and false negative probabilities.

### 5.1.2 Weaknesses Against Simple Collusion Attacks

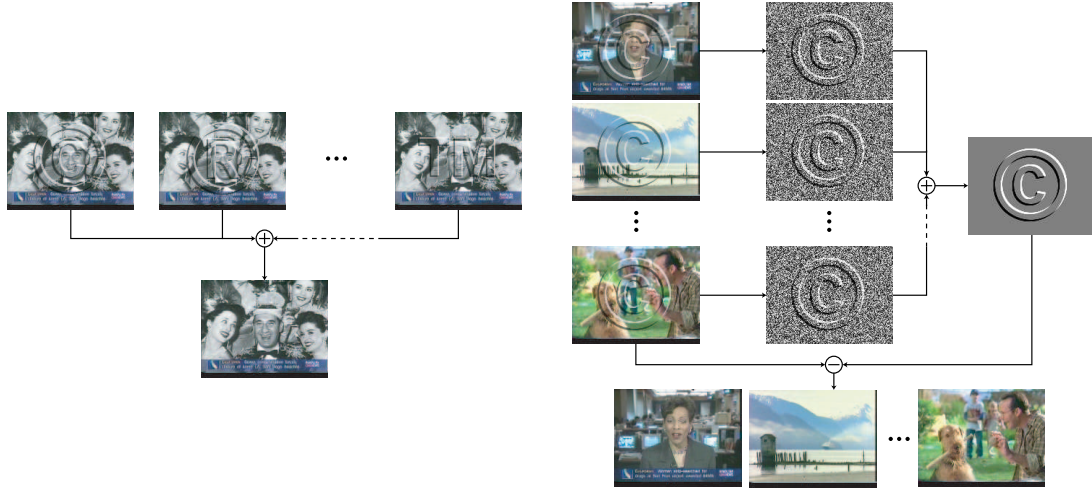
In a watermarking perspective, security can be seen as the resistance of the hidden watermark against hostile intelligence. In particular, collusion attacks have to be considered. The goal of collusion attacks is to produce unwatermarked content by combining several watermarked contents. Since each video frame can be seen as a single watermarked document, collusion is crucial in digital video watermarking. Two basic collusion attacks depicted in Figure 5.3 will be introduced hereafter to show how simple it is to defeat frame-by-frame video watermarking systems such as SS and SS-1.

#### Temporal Frame Averaging (TFA)

Since neighboring video frames are highly similar, temporal low-pass filtering can be performed without introducing much visual distortion:

$$\hat{\mathbf{f}}_t = L_w(\mathcal{F}_t), \quad \mathcal{F}_t = \{\mathbf{f}_u, -w/2 \leq t - u < w/2\} \quad (5.5)$$

where  $w$  is the size of the temporal window,  $L_w$  is a temporal low-pass filter and  $\hat{\mathbf{f}}_t$  is the resulting  $t$ -th attacked video frame. In experiments, a simple 3-frames



(a) Temporal Frame Averaging (TFA): Similar video frames carrying uncorrelated watermarks are averaged to produce unwatermarked content.

(b) Watermark Estimation Remodulation (WER): Several watermark estimations obtained from different video frames are combined to refine the estimation and allow watermark removal.

**Figure 5.3:** Visual illustration of basic collusion attacks.

temporal averaging filter has been used. Assuming that a watermarked video  $\{\check{\mathbf{f}}_t\}$  is temporally averaged, the following correlation score is obtained on the detector side:

$$\rho(\{\check{\mathbf{f}}_t\}) \approx \frac{\alpha}{wT} \sum_{t=1}^T \left( \sum_{u \in [-\frac{w}{2}, \frac{w}{2}[} \mathbf{w}_{t+u} \cdot \mathbf{w}_t \right) \quad (5.6)$$

If the same watermark has been redundantly embedded (SS-1 system), all the correlation terms  $\mathbf{w}_{t+u} \cdot \mathbf{w}_t$  are equal to 1 and the correlation score is equal to  $\alpha$ . In other words the TFA attack fails. Alternatively, if uncorrelated watermarks have been inserted in successive video frames (SS system), the term corresponding to the index  $u = 0$  in the second summation is the only one not to be null and the correlation score is reduced to  $\alpha/w$ . As a result, for  $w$  greater than 2, the correlation score drops below the detection threshold  $\tau_{\text{detect}}$  and the attack is a success. Averaging many video frames is likely to result in poor quality video in dynamic scenes. This attack is consequently more relevant in static scenes even if it can be adapted to cope with dynamic ones thanks to frame registration as it will be discussed in Section 6.1.

### Watermark Estimation Remodulation (WER)

Computing the difference  $\Delta_o(\check{\mathbf{f}}) = \check{\mathbf{f}} - \mathbf{f}$  is the optimal approach to estimate the watermark embedded in a given video frame. However, the attacker does not have access to the original digital content and has to blindly estimate the hidden watermark in practice. Digital watermarks are usually located in high

frequencies. A rough estimation of the watermark can consequently be obtained with denoising techniques, or more simply by computing the difference between the watermarked frame and its low-pass filtered version [191]:

$$\Delta(\check{\mathbf{f}}) = \check{\mathbf{f}} - L(\check{\mathbf{f}}) \quad (5.7)$$

where  $L(\cdot)$  is a low-pass filter e.g. a simple  $5 \times 5$  spatial averaging filter. Then, estimations obtained from different video frames are averaged [82]:

$$\tilde{\mathbf{w}}_T = \frac{1}{T} \sum_{t=1}^T \tilde{\mathbf{w}}_t = \frac{1}{T} \sum_{t=1}^T \Delta(\check{\mathbf{f}}_t) \quad (5.8)$$

where  $T$  is the number of considered video frames for collusion. In practice, the estimator defined in Equation (5.7) produces badly estimated samples around discontinuities (edges or textured areas). An additional thresholding operation is consequently performed to discard samples whose magnitude is greater than  $\tau_{\text{valid}}$ . The threshold value has been set to 8 for experiments and the number of valid estimations for each watermark sample has been counted to allow pertinent normalization in Equation (5.8). The resulting watermark  $\tilde{\mathbf{w}}_T$  is then subtracted from each watermarked video frame with a remodulation strength  $\beta$ . This strength is chosen to introduce a distortion similar to the one due to the watermarking process in terms of Peak Signal to Noise Ratio (PSNR). The attacked video frames are thus given by:

$$\hat{\mathbf{f}}_t = \check{\mathbf{f}}_t - \alpha \tilde{\mathbf{w}}_{T_n} = \check{\mathbf{f}}_t - \alpha \frac{\tilde{\mathbf{w}}_T}{\sqrt{\tilde{\mathbf{w}}_T \cdot \tilde{\mathbf{w}}_T}} \quad (5.9)$$

Assuming that the attacker has access to the estimator  $\Delta_o(\cdot)$ , when a watermarked video is submitted to the WER attack, the detector obtains the following correlation score:

$$\rho(\{\hat{\mathbf{f}}_t\}) \approx \alpha \left[ 1 - \frac{1}{T^2 \sqrt{\tilde{\mathbf{w}}_T \cdot \tilde{\mathbf{w}}_T}} \sum_{t=1}^T \sum_{u=1}^T \mathbf{w}_u \cdot \mathbf{w}_t \right] \quad (5.10)$$

If the watermarks embedded in different video frames are uncorrelated (SS system), the correlation term  $\mathbf{w}_u \cdot \mathbf{w}_t$  is equal to  $\delta_u^t$  where  $\delta$  is the Kronecker delta and the correlation score after attack is equal to  $\alpha(1 - 1/\sqrt{T})$  which is almost equal to  $\alpha$  for large  $T$ . As a result, the attack does not succeed in removing an embedded watermark if a strategy which *always embeds a different watermark* is enforced. On the other hand, if the same watermark has been redundantly embedded in all the video frames (SS-1 system), each correlation term is equal to 1 and the correlation score drops down to zero. This result has to be contrasted since the attacker has not access to  $\Delta_o(\cdot)$ . However, combining several individual estimates as in Equation (5.8) refines the final one and the attack proves to be a success in practice [82]. In fact, the more the video frames are different, the more each individual watermark estimate refines the final one i.e. the attack is more relevant in dynamic scenes.

## 5.2 Switching Between Orthogonal Watermarks

Subsection 5.1.2 highlighted two important facts. First, uncorrelated watermarks can be washed out with temporal frame averaging. Second, a redundant watermark can be estimated and later removed via remodulation. Watermark modulation is explored in the remainder of this section: for each video frame, the watermark is picked out from a finite pool of reference watermark patterns. The superiority of this strategy in terms of security is demonstrated both theoretically and experimentally. Its limitations against an expert attacker are also outlined.

### 5.2.1 SS-N System

Periodic watermark schedules have been investigated for temporal synchronization [121]. However, from a security point of view, repeating the same sequence of watermarks allows an attacker to group frames carrying the same watermark before performing a WER attack. Thus, for each video frame, the watermark should rather be randomly chosen from a finite set of  $N$  watermarks  $\{\mathbf{w}_i\}$  as depicted in Figure 5.4. Both previous systems are specific cases of this novel architecture:  $N = 1$  for SS-1 system and  $N = \infty$  for SS system. Watermarks are orthonormalized to prevent cross-talk on the detector side. The embedding process can then be rewritten:

$$\check{\mathbf{f}}_t = \mathbf{f}_t + \alpha \mathbf{w}_{\Phi(t)}, \quad P(\Phi(t) = i) = p_i \quad (5.11)$$

where the  $p_i$ 's are the emission probabilities of the system. From a subjective point of view, changing the watermark pattern still introduces a flicker artifact.

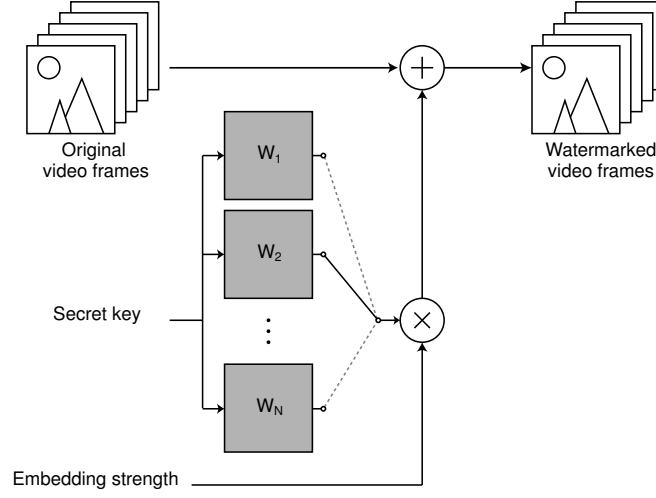
On the detector side, a new correlation score<sup>2</sup> is computed:

$$\rho(\{\check{\mathbf{f}}_t\}) = \frac{1}{T} \sum_{t=1}^T \sum_{i=1}^N |\check{\mathbf{f}}_t \cdot \mathbf{w}_i| \quad (5.12)$$

For each video frame,  $N$  linear correlations are computed and their absolute values are summed before being temporally averaged. This detection process does not require synchronization. However, the complexity<sup>3</sup> of the detector is increased by a factor  $N$  and the linearity of the operator  $\cdot$  cannot be exploited as in Equation (5.4) because of the absolute values. Immediately after embedding,

<sup>2</sup>Changing the detector has of course an impact on the detection statistics. In particular, the variance is increased by a factor  $\sqrt{N}$  in comparison with SS and SS-1 systems i.e. more frames need to be accumulated to have the same false positive and false negative probabilities.

<sup>3</sup>Complexity can be reduced by using non full frame watermark patterns. In other words, each frame is partitioned in  $N$  non-overlapping areas and each watermark pattern is spread over one of these areas. As a result, each  $\mathbf{f}_t \cdot \mathbf{w}_i$  has  $N$  times fewer terms. However, this also alters detection statistics i.e. robustness performance.



**Figure 5.4:** SS- $N$  system: the embedder inserts a watermark randomly chosen from a collection of  $N$  reference watermarks.

the detector obtains:

$$\begin{aligned} \rho(\{\tilde{\mathbf{f}}_t\}) &= \frac{1}{T} \sum_{t=1}^T \sum_{i=1}^N \left| \mathbf{f}_t \cdot \mathbf{w}_i + \alpha \mathbf{w}_{\Phi(t)} \cdot \mathbf{w}_i \right| \\ &\approx \frac{\alpha}{T} \sum_{t=1}^T \sum_{i=1}^N \delta_{\Phi(t)}^i \approx \alpha \end{aligned} \quad (5.13)$$

Host interference is cancelled in a preprocessing step [34] to improve detection statistics. The correlation score is then equal to  $\alpha$  if a watermark is present in the video and to zero otherwise. This score is consequently compared to a threshold  $\tau_{\text{detect}}$ , which is set equal to  $\alpha/2$  in practice, in order to assert the presence or absence of a hidden watermark.

### 5.2.2 Enhanced Security

If a watermarked video is temporally averaged with a large window size  $w$  i.e. a strong attack without any concern for video quality, the attacked video frames are then given by:

$$\hat{\mathbf{f}}_t = \frac{1}{w} \sum_{u \in [-\frac{w}{2}, \frac{w}{2}[} \mathbf{f}_{t+u} + \alpha \sum_{i=1}^N p_i \mathbf{w}_i \quad (5.14)$$

Thus, the detector obtains the following correlation score:

$$\rho(\{\hat{\mathbf{f}}_t\}) = \frac{1}{T} \sum_{t=1}^T \sum_{i=1}^N |\hat{\mathbf{f}}_t \cdot \mathbf{w}_i| \approx \frac{\alpha}{T} \sum_{t=1}^T \sum_{i=1}^N p_i \approx \alpha \quad (5.15)$$



TFA spreads the energy of a watermark embedded in a video frame over its neighboring frames. In the SS system, when the detector checks for the presence of the watermark that should be embedded in each video frame, it misses most of the watermark signal. On the other hand, the SS- $N$  detector checks the presence of *all the watermarks* of the set  $\{\mathbf{w}_i\}$  in each video frame and thus retrieves all the parts of each watermark. As a result, the TFA attack fails.

Assuming that the attacker has access to the estimator  $\Delta_o(\cdot)$ , if a watermarked video is submitted to the WER attack, the final watermark estimate is equal to  $\tilde{\mathbf{w}}_T = \sum_{i=1}^N p_i \mathbf{w}_i$ . After remodulation, the following video frames are produced:

$$\hat{\mathbf{f}}_t = \mathbf{f}_t + \alpha \left[ \left(1 - \frac{p_t}{\nu}\right) \mathbf{w}_{\Phi(t)} - \sum_{i \neq \Phi(t)} \frac{p_i}{\nu} \mathbf{w}_i \right] \quad (5.16)$$

where  $\nu = \sqrt{\tilde{\mathbf{w}}_T \cdot \tilde{\mathbf{w}}_T}$ . Subsequently, the detector obtains the following correlation score:

$$\begin{aligned} \rho(\{\hat{\mathbf{f}}_t\}) &\approx \frac{\alpha}{T} \sum_{t=1}^T \sum_{i=1}^N \left| \left(1 - \frac{p_{\Phi(t)}}{\nu}\right) \delta_{\Phi(t)}^i - \sum_{j \neq \Phi(t)} \frac{p_j}{\nu} \delta_j^i \right| \\ &\approx \alpha \sum_{i=1}^N p_i \left[ \left(1 - \frac{p_i}{\nu}\right) + \sum_{j \neq i} \frac{p_j}{\nu} \right] \end{aligned} \quad (5.17)$$

If all the  $p_i$  are equal to  $1/N$ , the norm  $\nu$  is equal to  $1/\sqrt{N}$  and Equation (5.17) becomes:

$$\rho(\{\hat{\mathbf{f}}_t\}) = \alpha \left[ 1 + (N-2) \frac{\sqrt{N}}{N} \right] \quad (5.18)$$

In other words, for  $N$  greater or equal to 2, the correlation score is greater or equal to  $\tau_{\text{detect}}$  and the attack fails. Here, using several watermarks has interfered with the watermark estimation process. Thus, the attacker can only remove a small fraction  $\sqrt{N}/N$  of the embedded watermark in each video frame. On the other hand, a small part of all the other watermarks from the set  $\{\mathbf{w}_i\}$  is also removed. Then, summing the *absolute values* of the linear correlations succeeds in compensating the loss of correlation with the originally embedded watermark. Absolute values play a key role in fact. If they are removed from Equation (5.12), the algorithm is still immune to TFA but the WER attack causes then the correlation score to drop to zero. Equation (5.18) also reminds that the WER attack is a success for  $N = 1$  (SS-1 system).

### 5.2.3 Experimental Results

Five videos ( $704 \times 576$ , 25 frames per second, 375 frames) are used for experiments. Their content is summarized in Table 5.1. They are watermarked with the three

presented watermarking schemes, with a global embedding strength equal to 3. The PSNR is consequently around 38 dB which ensures the watermark invisibility. Four different watermarks have been used for the SS- $N$  system. The watermarked videos are then submitted to TFA on one hand and to the WER attack on the other. Finally, the correlation score is computed for all the videos.


**Table 5.1:** Description of the videos used for experiments.

<i>Video shot</i>	<i>Short description</i>
Ping-Pong	Moving players, camera zoom/static/pan
Ski	Fast moving skier tracked by the camera
Susie	Girl on phone close-up, lips/eye/head motion
Train	Many moving objects (train, ball, calendar), camera pan
Tree	Static landscape background, camera static/pan

Each watermarking scheme is represented in Figure 5.5 by a specific symbol: crosses for SS system, triangles for SS-1 system and circles for SS- $N$  system. The figure has also been divided into four quadrants whose borders are defined by the detection threshold  $\tau_{\text{detect}} = 1.5$ . The crosses are located in the upper-left quadrant, which confirms that the SS system resists the WER attack while it is weak against TFA. In fact they are in the neighborhood of the line defined by  $\mathbf{y} = w\mathbf{x}$  ( $w = 3$  in the experiments) as can be predicted from theoretical results in Subsection 5.1.2. On the other hand, the triangles are in the lower-right quadrant, which supports conjectures asserting that the SS-1 system is robust against TFA while the WER attack succeeds in stirring out the embedded watermark, even if this latter attack is more or less efficient depending on the video content of the shot. Finally, the circles are in the upper-right quadrant, meaning that the SS- $N$  system effectively resists both TFA and WER attacks. The WER attack even increases the correlation score as asserted in Equation (5.18).

## 5.2.4 Potential Weaknesses

Attackers are likely to modify and adjust their approach according to this novel watermarking strategy. The security of the SS- $N$  system basically relies on the assumption that attackers are unable to build sets of frames carrying the same watermark. Otherwise, a simple WER attack performed on each subset succeeds in estimating the pool of secret watermarks. It will be shown that a successful brute force attack can be theoretically designed. However, its computational complexity may prevent its use in practice. Alternatively, individual watermark estimates  $\{\tilde{\mathbf{w}}_t\}$  obtained from different video frames can be seen as vectors in a very high



**Figure 5.5:** Resilience of the three presented video watermarking systems (SS, SS-1 and SS-N) against TFA and WER collusion attacks.

dimensional space. Since these vectors should approximate the embedded watermarks  $\{\mathbf{w}_i\}$ , the problem comes down to vector quantization. Therefore, a method will be described hereafter, which defines  $N$  clusters  $\mathcal{C}_i$  whose centroids  $\mathbf{c}_i$  are good estimates of the secret watermarks. Moreover, its impact on the SS- $N$  system will be investigated.

### Brute Force Approach

According to Kerckhoffs' principle [97], the watermarking system is publicly known and the attacker is aware that  $N$  alternative watermarks have been randomly embedded in the video. One approach consists then in distributing the watermarked video frames between  $N$  sets  $\mathcal{S}_i$  and estimating the watermarks  $\tilde{\mathbf{w}}_i$  from each one of those sets. Assuming that the attacker has access to the ideal watermark estimator  $\Delta_o(\cdot)$ , the following watermarks are obtained:

$$\tilde{\mathbf{w}}_i = \sum_{j=1}^N n_{i,j} \mathbf{w}_j \quad (5.19)$$

where  $n_{i,j}$  is the number of frames carrying the watermark  $\mathbf{w}_j$  in the set  $\mathcal{S}_i$ . Moreover, if each set  $\mathcal{S}_i$  contains  $P$  frames, the  $n_{i,j}$ 's verify:

$$\sum_{i=1}^N n_{i,j} = P \quad \text{and} \quad \sum_{j=1}^N n_{i,j} = P \quad (5.20)$$

Once those watermark estimations have been obtained, the attacker should have a criterion to determine if one or more watermarks have been correctly estimated.

When the  $N$  sets  $\mathcal{S}_i$  are built, the  $n_{i,j}$  are unknown. The attacker can only compute the different correlations between the estimated watermarks  $\{\tilde{\mathbf{w}}_i\}$  defined as follows:

$$c_{i_1, i_2} = \tilde{\mathbf{w}}_{i_1} \cdot \tilde{\mathbf{w}}_{i_2} = \sum_{j=1}^N n_{i_1, j} \cdot n_{i_2, j} \quad (5.21)$$

For a given estimated watermark  $\tilde{\mathbf{w}}_{i_0}$ , the sum of the correlations with the set of estimated watermarks  $\{\tilde{\mathbf{w}}_i\}$  is equal to:

$$\sum_{i=1}^N c_{i_0, i} = \sum_{i=1}^N \sum_{j=1}^N n_{i_0, j} \cdot n_{i, j} = P^2 \quad (5.22)$$

Now, let assume that there exists an index  $i_0$  for which  $c_{i_0, i_0} > mP^2$  with  $m$  in  $[0.5, 1]$ . It can be shown that  $n_{i_0, j^*} = \max_j(n_{i_0, j})$  is greater than  $mP$ . Since  $m$  is greater than 0.5, the correlation score between  $\tilde{\mathbf{w}}_{i_0}$  and the video frames carrying  $\mathbf{w}_{j^*}$  is higher than with the other ones. As a result, the attacker can distinguish the frames carrying  $\mathbf{w}_{j^*}$ , obtain a finer estimate for  $\mathbf{w}_{j^*}$  and iterate the attack with a reduced set of video frames to estimate the remaining watermarks i.e. with a reduced complexity.

In summary, this demonstrates that the attacker can estimate and remove the embedded watermarks. However, the complexity of this brute force approach is very high. Since the probability that  $c_{i, i}$  is greater than  $mP^2$  is difficult to estimate, the probability that  $n_{i, j}$  is greater than  $mP$  will be considered below to obtain a lower bound for the complexity. It is obvious that the probability that  $n_{i, j}$  is equal to  $k$  is given by:

$$\begin{aligned} P(n_{i, j} = k) &= \binom{P}{k} P(\mathbf{w} = \mathbf{w}_i)^k P(\mathbf{w} \neq \mathbf{w}_i)^{P-k} \\ &= \binom{P}{k} \left(\frac{1}{N}\right)^k \left(1 - \frac{1}{N}\right)^{P-k} \end{aligned} \quad (5.23)$$

These probabilities can then be summed to obtain the probability  $p_L$  that  $n_{i, j}$  is strictly greater than  $L$ .

$$p_L = \sum_{n=\lfloor L+1 \rfloor}^P \binom{P}{n} \left(\frac{1}{N}\right)^n \left(1 - \frac{1}{N}\right)^{P-n} \quad (5.24)$$

When  $N$  grows large, this probability  $p_L$  is almost reduced to a single term:

$$\begin{aligned} p_L &\approx \binom{P}{\lfloor L+1 \rfloor} \left(\frac{1}{N}\right)^{\lfloor L+1 \rfloor} \left(1 - \frac{1}{N}\right)^{P-\lfloor L+1 \rfloor} \\ &\approx \binom{P}{\lfloor L+1 \rfloor} N^{-\lfloor L+1 \rfloor} \end{aligned} \quad (5.25)$$

As a result, the attacker should in average distribute the video frames between  $N$  sets  $1/p_L$  times before obtaining a distribution that can be exploited. Moreover, for each one of these distributions,  $N(N + 1)/2$  correlations between estimated watermark patterns are computed. In other words, the number of correlation  $n_{corr}^N$  is equal to:

$$\begin{aligned} n_{corr}^N &= \frac{1}{2} \binom{P}{\lfloor L+1 \rfloor}^{-1} N^{\lfloor L+1 \rfloor} N(N+1) \\ &\approx \frac{1}{2} \binom{P}{\lfloor L+1 \rfloor}^{-1} N^{\lfloor L+3 \rfloor} \end{aligned} \quad (5.26)$$

When the attacker has performed all those operations, a single watermark has been estimated and the process has to be continued to estimate the  $N - 1$  remaining watermarks. Consequently the total number of computed correlations is equal to:

$$n_{corr} = \sum_{i=2}^N n_{corr}^i \approx \frac{1}{2 \lfloor L+4 \rfloor} \binom{P}{\lfloor L+1 \rfloor}^{-1} N^{\lfloor L+4 \rfloor} \quad (5.27)$$

In practice,  $L$  is equal to  $mP$  with  $m$  in  $[0.5, 1]$  which means that the attacker is required to compute at least  $O(N^{\lfloor mP+4 \rfloor})$  correlations between estimated watermarks to terminate the proposed brute force attack. Thus, on one side, checking for the presence of  $N$  alternative watermarks in each video frame increases the complexity of the embedding algorithm by a factor  $N$ . On the attacker side, the complexity of a brute force attack is proportional to  $N^{\lfloor mP+4 \rfloor}$ . For instance, for  $N = 64$ ,  $P = 50$  and  $m = 0.5$ , Equation (5.27) means that the brute force attack requires at least  $3.10^{36}$  correlation computations. Such a high computational cost is likely to discourage most attackers.

### Watermark Estimations Clustering and Remodulation (WECR)

The  $k$ -means algorithm is a simple way to perform vector quantization. In a first step, the individual watermark estimates  $\{\tilde{\mathbf{w}}_t\}$  are distributed amongst different clusters  $\{\mathcal{C}_i\}$ , so that each vector is assigned to the cluster associated with its nearest centroid  $\mathbf{c}_i$  according to the distance below:

$$d(\tilde{\mathbf{w}}_t, \mathbf{c}_i)^2 = \frac{1}{P} \left[ \sum_{x \in \mathcal{V}} \left( \tilde{\mathbf{w}}_t(x) - \mathbf{c}_i(x) \right)^2 + \sum_{x \notin \mathcal{V}} \mathbf{c}_i^2(x) \right] \quad (5.28)$$

where  $P$  is the frame dimension and  $\mathcal{V}$  the set of valid samples i.e. whose magnitude is lower than  $\tau_{\text{valid}}$ . The first term in Equation (5.28) measures how close the observation  $\tilde{\mathbf{w}}_t$  is from the centroid  $\mathbf{c}_i$  considering only the valid samples. The other term is a penalty term which favors observations having more valid samples. In a second step, the centroids are updated using only valid samples and the algorithm iterates until convergence.

To avoid random initialization, a splitting strategy [124] has been introduced. The basic idea is to start with a single cluster and to increment iteratively the number of clusters. Once the  $k$ -means algorithm has run until convergence, the log-likelihood  $L_i$  of each cluster is computed:

$$L_i = -\frac{|\mathcal{C}_i|}{2} \left[ 1 + \log \left( \frac{2\pi}{|\mathcal{C}_i|} \sum_{\tilde{\mathbf{w}}_t \in \mathcal{C}_i} d(\tilde{\mathbf{w}}_t, \mathbf{c}_i)^2 \right) \right] \quad (5.29)$$

where  $|\mathcal{C}_i|$  is the number of vectors contained in the cluster  $\mathcal{C}_i$ . The worst cluster, the one with the lowest log-likelihood, is then identified and its associated centroid  $\mathbf{c}_{\text{worst}}$  is split in  $\mathbf{c}_{\text{worst}} \pm \epsilon \mathbf{d}$  where  $\epsilon$  is a very small value and  $\mathbf{d}$  is a direction to be set. This direction can be fixed, random or even better the direction of principal variation in the cluster. After each split, the  $k$ -means algorithm is run until convergence. This splitting strategy is stopped when the last split has not significantly reduced the average of the distances between each watermark estimate  $\tilde{\mathbf{w}}_t$  and its nearest centroid.

At this point,  $M$  centroids have been obtained which are assumed to estimate the embedded watermark patterns. Thus, they can be remodulated to alter the watermark signal:

$$\dot{\mathbf{f}}_t = \check{\mathbf{f}}_t - \alpha \frac{\mathbf{c}_{\tilde{\phi}(t)}}{\sqrt{\mathbf{c}_{\tilde{\phi}(t)} \cdot \mathbf{c}_{\tilde{\phi}(t)}}} \quad (5.30)$$

where  $\tilde{\phi}(t) = \arg \max_i \check{\mathbf{f}}_t \cdot \mathbf{c}_i$ . If the attacker knows how many watermarks have been used during embedding, an additional merging step [166] can be introduced to have exactly the same number  $N$  of centroids. The basic idea consists in successively merging the two most similar centroids, according to a given metric such as the correlation coefficient for example:

$$\mathbf{c}_{i \cup j} = \frac{|\mathcal{C}_i| \mathbf{c}_i + |\mathcal{C}_j| \mathbf{c}_j}{|\mathcal{C}_i| + |\mathcal{C}_j|} \quad (5.31)$$

### Impact of WE CR

The videos presented in Table 5.1 have been watermarked with the SS- $N$  system using 4 alternative watermarks and an embedding strength  $\alpha$  equal to 3. Next, the watermarked videos have been submitted to the WE CR attack with and without an additional merging step. The detection score has been computed before and after the attack and the results have been gathered in Table 5.2. The value in brackets indicates the detection score when a merging step is introduced. It is clear that the efficiency of the attack depends on the content of the video. The more dynamic the video content, the more different the individual watermark estimates and the more effective the watermark estimation refinement process. Furthermore, if the video contains long static shots, it can interfere with the

splitting strategy and results in *bad* centroids i.e. which gathers video frames not carrying the same watermark pattern  $\mathbf{w}_i$ . Adding a merging step may then alter the efficiency of the attack (*ping-pong* video). In real life, an attacker would not use successive frames from a video, but would rather extract some key frames of the watermarked video. As an example, a TV news video with commercial breaks has been watermarked with the SS- $N$  system and 325 key frames have been extracted to perform the WECR attack. In this case, almost 90% of the watermark signal has been properly estimated, which succeeds in lowering the correlation score from 2.91 to 0.52 (0.45) i.e. a score below the detection threshold.

**Table 5.2:** Impact of the WECR attack on the detection score of the SS- $N$  system.

<i>Video shot</i>	<i>Before WECR attack</i>	<i>After WECR attack</i>
Ping-Pong	2.92	1.73 (3.24)
Ski	2.82	0.46 (0.45)
Susie	3.00	0.30 (0.27)
Train	2.89	0.70 (0.54)
Tree	2.37	1.63 (1.02)

## 5.3 Embedding Strength Modulation


The SS- $N$  system exploits watermark modulation to obtain superior performance against collusion attacks. However, an expert attacker can still remove the embedded watermark with an attack based on vector quantization. A new geometrical interpretation is consequently introduced in this section to obtain a novel perspective and thus a better understanding of the weaknesses of the previous watermarking schemes. From these observations, embedding strength modulation is explored to achieve security. Limitations of such an approach against hostile intelligence are also evaluated.

### 5.3.1 Novel Perspective

The presented three video watermarking systems all embed a normally distributed watermark  $\mathbf{w}_t$  with zero mean and unit variance in each frame  $\mathbf{f}_t$  with a fixed embedding strength  $\alpha$ :

$$\check{\mathbf{f}}_t = \mathbf{f}_t + \alpha \mathbf{w}_t(K), \quad \mathbf{w}_t(K) \sim \mathcal{N}(0, 1) \quad (5.32)$$

The embedded watermark can be seen as a low-power pseudo-random image of  $P$  pixels which is scaled and added to a video frame. Alternatively, it can be considered as a disturbing random vector drawn from a  $P$ -dimensional space which is added to a host vector. In this case, the norm of the first vector has to be far lower than the norm of the latter to fulfill the invisibility constraint. Since watermarks are zero mean, they are in fact drawn from a  $(P - 1)$  dimensional subspace. Furthermore, they are bound to lie on the unit sphere associated with the distance  $d(\mathbf{x}, \mathbf{y}) = \sqrt{(\mathbf{x} - \mathbf{y}) \cdot (\mathbf{x} - \mathbf{y})}$  as they have unit variance. Now, even if the presented watermarking schemes share a common framework, they enforce a different embedding strategy. This has a direct impact on how the different watermarks are distributed over the unit sphere as illustrated in Figure 5.6.



**Figure 5.6:** Distribution of the embedded watermarks over the unit sphere depending on the enforced watermarking strategy in a 3-dimensional watermarking subspace.

This geometric approach sheds a new light on the link between embedding strategies and security issues. When embedded watermarks are uniformly distributed over the unit sphere (SS system), averaging successive watermarks results then in a very small vector in the middle of the unit sphere i.e. there is very little residual watermark energy. Alternatively, when watermarks are gathered in a single narrow area (SS-1 system), or even several areas (SS- $N$  system), the watermarks can be distributed amongst well-identified clusters. As a conclusion, successive watermarks define a trajectory over the unit sphere and this watermark trajectory should have some properties to resist collusion attacks. First it should be continuous so that averaging successive watermarks results in a watermark near the surface of the unit sphere. Second, the trajectory should not have accumulation points to prevent weaknesses against WECD attacks.



### 5.3.2 SS- $\alpha$ System

The SS- $N$  system relies on watermark modulation to achieve security. However, an alternative strategy exploiting the embedding strength can also be explored. The basic idea consists in using a time dependent embedding strength  $\beta(t)$ :

$$\check{\mathbf{f}}_t = \mathbf{f}_t + \alpha\beta(t)\mathbf{w}(K), \quad \mathbf{w}(K) \sim \mathcal{N}(0, 1) \quad (5.33)$$

Here the embedding strength is modulated for security reasons and not to improve watermark invisibility as usual. With this end in view, the modulation function  $\beta(t)$  has to respect the three following constraints:

1. It should vary smoothly in time to be immune to TFA attacks,
2. It should be zero mean to resist a potential WER attacks,
3. It should have a large number of values after discrete sampling to avoid WEER attacks.

Keeping these specifications in mind, a set  $\{\mathbf{w}_i\}$  of  $N$  orthonormal watermark patterns is built. The embedding procedure of the SS- $\alpha$  system is then defined as follows:

$$\check{\mathbf{f}}_t = \mathbf{f}_t + \alpha \sum_{i=1}^N \beta_i(t)\mathbf{w}_i = \mathbf{f}_t + \alpha\mathbf{w}_t \quad (5.34)$$

The modulation functions  $\beta_i(t)$  have to be chosen in accordance with the previously cited specifications to achieve security. The SS- $N$  system can indeed be seen as a specific case of this new system where the modulation functions are equal to  $\beta_i(t) = \delta_{\Phi(t)}^i$ . However, such modulation functions only give  $N$  possible combinations of watermarks and this system can be defeated by a WEER attack. An additional constraint is introduced so that embedded watermarks  $\mathbf{w}_t$  all lie on the unit sphere. In other words, the modulation functions should verify:

$$\forall t \quad \sum_{i=1}^N \beta_i^2(t) = \mathbf{w}_t \cdot \mathbf{w}_t = 1 \quad (5.35)$$

As a result, the embedding process introduces a Mean Square Error (MSE) equal to  $\alpha^2$  and an embedding strength  $\alpha$  equal to 3 induces a distortion of about 38 dB. The detector computes the energy<sup>4</sup> contained in the subspace spanned by

---

<sup>4</sup>As for the SS- $N$  system, changing the detector has an impact on the detection statistics. Here again, the variance is increased and more frames need to be accumulated to obtain similar false positive and false negative probabilities than for SS or SS-1 systems

the watermark patterns  $\mathbf{w}_i$ :

$$\begin{aligned}\rho(\{\check{\mathbf{f}}_t\}) &= \sqrt{\frac{1}{T} \sum_{t=1}^T \sum_{i=1}^N (\check{\mathbf{f}}_t \cdot \mathbf{w}_i)^2} \\ &\approx \alpha \sqrt{\frac{1}{T} \sum_{t=1}^T \sum_{i=1}^N \beta_i^2(t)} = \alpha\end{aligned}\quad (5.36)$$

Host interference is cancelled in a preprocessing step [34] to enhance detection statistics. The detection score should be equal to  $\alpha$  if a watermark is present in the video, while it should be almost equal to zero if no watermark has been inserted. It is consequently compared to a threshold  $\tau_{\text{detect}}$ , which is set equal to  $\alpha/2$  in practice, to assert the presence or absence of a hidden watermark.

### Sinusoidal Modulation

A sinusoidal embedding strength [52] can be used to have a practical implementation of this strategy:

$$\beta_i(t) = \sqrt{\frac{2}{N}} \sin(\Omega t + \phi_i) \quad (5.37)$$

where  $\Omega$  is a shared radial frequency and  $\phi_i$  are phases to be set appropriately. From a communication perspective, this system can be considered as transmitting the same low-power temporal signal  $\sin(\Omega t)$  along several non-interfering channels  $\mathbf{w}_i$  with some phase differences  $\phi_i$ . The square norm of the embedded watermarks  $\mathbf{w}_t$  is then given by:

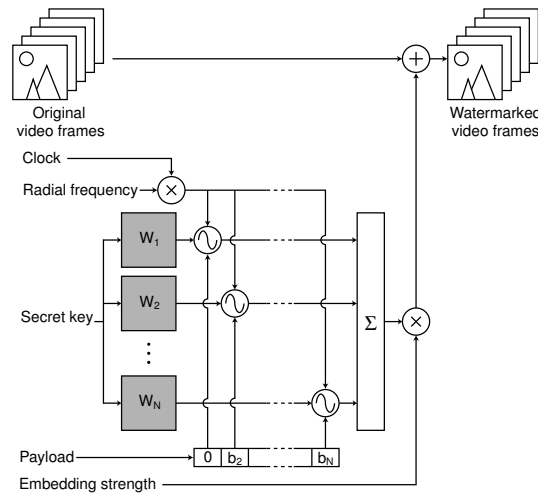
$$\begin{aligned}\mathbf{w}_t \cdot \mathbf{w}_t = 1 &- \frac{\cos(2\Omega t)}{N} \sum_{i=1}^N \cos(2\phi_i) \\ &+ \frac{\sin(2\Omega t)}{N} \sum_{i=1}^N \sin(2\phi_i)\end{aligned}\quad (5.38)$$

The phase differences  $\phi_i$  should be chosen so that both sums are equal to zero to fulfill Equation (5.35). The  $N$ -th roots of unity in  $\mathbb{C}$  can be taken into account and  $2\phi_i = i2\pi/N$  modulo  $2\pi$ . An ambiguity regarding the value of  $\phi_i$  still remains, leaving room for embedding a moderate payload:

$$\phi_1 = 0, \quad \phi_i = \left(\frac{i}{N} + b_i\right)\pi \pmod{2\pi} \quad (5.39)$$

where  $b_i \in \{0, 1\}$  is a bit of payload. Since the detector will only be able to estimate phase differences, the phase  $\phi_1$  is set 0 to allow payload retrieval. The

whole embedding process is depicted in Figure 5.7. On its side, the detector correlates each incoming video frame  $\mathbf{f}_t$  with all the watermark patterns  $\mathbf{w}_i$  to obtain an estimate  $\tilde{\beta}_i(t) = \mathbf{f}_t \cdot \mathbf{w}_i$  of the temporal signal transmitted along each communication channel. Next, the detection score given in Equation (5.36) is computed to assert whether an underlying watermark is present in the video or not. If a watermark is detected, the payload bits are extracted by estimating the phase differences  $\phi_i$ . This can be easily done by computing the unbiased cross-correlation between the reference signal  $\tilde{\beta}_1(t)$  and the other  $\tilde{\beta}_i(t)$  whose phase difference encodes a payload bit  $b_i$  according to Equation (5.39).



**Figure 5.7:** SS- $\alpha$  system with sinusoidal modulation: the embedder inserts a linear combination of  $N$  reference watermark patterns, whose mixing coefficients are temporally sinusoidal.

### Security Constraints

Even if this novel system has been designed to resist collusion attacks, some parameters need to be carefully chosen. First, the radial frequency  $\Omega$  should remain secret or pseudo-secret to prevent an attacker from separating the watermarked video frames into distinct sets of frames carrying almost the same watermark signal<sup>5</sup>. Otherwise, a WER attack can then be successfully applied to each set. Now, if an attacker performs a WER attack on the whole video using the optimal

<sup>5</sup>In fact,  $\Omega$  can be estimated with a very simple temporal spectral estimation. This is a major security flaw for any watermarking system based on periodic watermark schedule. However this system is only presented for illustrative purpose. In the general case, this attack does not defeat the SS- $\alpha$  system.

watermark estimator  $\Delta_o(\cdot)$ , the following watermark estimate is obtained:

$$\tilde{\mathbf{w}}_T = \sum_{i=1}^N \left( \frac{\alpha}{T} \sum_{t=1}^T \beta_i(t) \right) \mathbf{w}_i = \sum_{i=1}^N \lambda_i(T) \mathbf{w}_i \quad (5.40)$$

The more video frames are considered, the closer the coefficients  $\lambda_i(T)$  are to zero. Since the attacker does not have access to the optimal watermark estimator in practice, each watermark estimation is noisy and accumulating several watermark estimations decreases the power of the watermark signal. That is to say that combining several individual watermark estimates hampers the final watermark estimation, which is in complete contradiction with the paradigm behind the original attack. The same property can be demonstrated with non-adjacent video frames. The radial frequency  $\Omega$  should also be set so that a given mixture of sinusoidal coefficients  $\beta_i(t)$  is never used twice. It should consequently be selected from  $\mathbb{R} - \pi\mathbb{Q}$  so that any WE CR attack is then doomed to fail.

Alternatively, an attacker can perform a TFA attack and obtain the following attacked video frames:

$$\hat{\mathbf{f}}_t = \frac{1}{w} \sum_{u \in [-\frac{w}{2}, \frac{w}{2}[} \mathbf{f}_{t+u} + \alpha \gamma_w \mathbf{w}_t, \quad \gamma_w = \frac{\text{sinc}(\frac{w\Omega}{2})}{\text{sinc}(\frac{\Omega}{2})} \quad (5.41)$$

Looking at Equation (5.34), TFA has basically scaled the embedded watermark signal by a signed attenuation factor  $\gamma_w$ . The larger the temporal window size  $w$ , the lower the attenuation factor. Similarly, the higher the radial frequency  $\Omega$ , the closer the attenuation factor to zero. As a result, the radial frequency  $\Omega$  should be chosen in such a way that the attenuation factor remains higher than a threshold value  $\gamma_{\text{lim}}$  as long as the temporal window size is lower than a given value  $w_{\text{max}}$ . If a larger window size is used, the content provider considers that the video has lost its commercial value due to the loss of visual quality. In other words, the parameters  $\gamma_{\text{lim}}$  and  $w_{\text{max}}$  give a higher bound for the radial frequency  $\Omega$  so that TFA only results in a small attenuation of the hidden signal.

### 5.3.3 Watermarking Subspace Estimation Draining (WSED)

Embedded watermarks  $\mathbf{w}_t$  are always a linear combination of a small number of reference watermark patterns  $\mathbf{w}_i$  as written in Equation (5.34). In other words, embedded watermarks are restricted to a low dimensional watermarking subspace which can be estimated<sup>6</sup> using space dimension reduction techniques [50]. Having

<sup>6</sup>It should be noted that this estimation of the watermarking subspace can be exploited to enhance the previously described WE CR attack. The watermark estimates  $\tilde{\mathbf{w}}_t$  are projected onto the estimated subspace  $\mathcal{E}$  prior to vector quantization. Once the coordinates of the clusters have been identified in the watermarking subspace, the centroids  $\mathbf{c}_i$  can then be easily retrieved.

a collection of  $T$  individual watermark estimates of size  $P$  and knowing that the embedded watermarks are contained in a  $N$ -dimensional subspace ( $N \ll P$ ), the attacker wants to find  $N$  vectors  $\mathbf{e}_i$  which span the same subspace as the one generated by the secret patterns  $\mathbf{w}_i$ :

$$\mathcal{W} = \text{span}(\mathbf{w}_i) = \text{span}(\mathbf{e}_i) = \mathcal{E} \quad (5.42)$$

With this end in view, Principal Component Analysis (PCA) can be performed since it is an optimal dimension reduction technique. Let  $\tilde{\mathbf{W}}$  be a  $P \times T$  matrix whose columns are the individual watermark estimates  $\tilde{\mathbf{w}}_t$ . The goal is to find a  $P \times N$  matrix  $\mathbf{E}$  and a  $N \times T$  matrix  $\mathbf{V}$  which minimize the norm  $\|\tilde{\mathbf{W}} - \mathbf{E}\mathbf{V}\|$ . Each column of the matrix  $\mathbf{V}$  can be viewed as the coordinates of the associated watermark estimate in the matrix  $\tilde{\mathbf{W}}$  in the principal subspace spanned by the vectors defined by the columns of matrix  $\mathbf{E}$ .

### Attack Description

As standard methods for PCA require too much memory for high dimensional data, an approach based on the Expectation-Maximization (EM) algorithm is exploited [163]. The PCA procedure is then reduced to an iterative algorithm using two steps:

$$\text{E-step: } \quad \mathbf{V} = (\mathbf{E}^T \mathbf{E})^{-1} \mathbf{E}^T \tilde{\mathbf{W}} \quad (5.43)$$

$$\text{M-step: } \quad \mathbf{E} = \tilde{\mathbf{W}} \mathbf{V}^T (\mathbf{V} \mathbf{V}^T)^{-1} \quad (5.44)$$

where  $\cdot^T$  denotes the transposition operator. A major asset of this approach is that it can be performed *online* using only a single watermark estimate at a time, which significantly reduces storage requirements. Moreover, the EM framework supports missing data, i.e. non pertinent estimated samples. During the E-step, for each incomplete watermark estimate  $\tilde{\mathbf{w}}_t$ , the coordinates  $\mathbf{v}_t$  in the current estimated subspace are computed using only valid samples and missing information is completed so that the distance to the current principal subspace is minimized. The completed watermark estimate  $\tilde{\mathbf{w}}_t^*$  is then used for the M-step. After the PCA iterations, a  $N$ -dimensional subspace  $\mathcal{E}$  has been estimated which is assumed to be close to the watermarking subspace  $\mathcal{W}$ . Thus it is drained from any energy:

$$\begin{aligned} \dot{\mathbf{f}}_t &= \check{\mathbf{f}}_t - \sum_{i=1}^N (\check{\mathbf{f}}_t \cdot \mathbf{e}_i) \mathbf{e}_i \\ &\approx \mathbf{f}_t + \sum_{i=1}^N \alpha_i(t) \left( \mathbf{w}_i - \sum_{j=1}^N (\mathbf{w}_i \cdot \mathbf{e}_j) \mathbf{e}_j \right) \end{aligned} \quad (5.45)$$

where  $\{\mathbf{e}_i\}$  is an orthonormalized basis of the subspace  $\mathcal{E}$  e.g. the eigenvectors of matrix  $\mathbf{E}$ . If the watermarking subspace  $\mathcal{W}$  has been finely estimated, the terms  $\mathbf{w}_i - \sum_{j=1}^N (\mathbf{w}_i \cdot \mathbf{e}_j) \mathbf{e}_j$  are null and the embedded watermark is removed.

### Impact of WSED

A TV news video with commercial breaks has been watermarked with the sinusoidal implementation of the SS- $\alpha$  system. An 8 bit payload has been hidden using  $N = 9$  watermark patterns  $\mathbf{w}_i$  and the embedding strength  $\alpha$  has been set equal to 3. Previous experiments have shown that collusion attacks are more efficient when the several individual watermark estimates originate from video frames with uncorrelated contents. As a result, key frames of the watermarked video have been extracted and used to estimate the watermarking subspace  $\mathcal{W}$ . Eventually, all the frames of the watermarked video were drained of any energy contained in the estimated subspace  $\mathcal{E}$ . This WSED attack has reduced the detection score given in Equation (5.36) from 2.96 to 0.53. In other words, there is no longer enough watermark energy and the attack is a success. This result however has to be contrasted. First, for a given dimension  $N$ , the more watermarked video frames  $\tilde{\mathbf{w}}_t$  are considered, the finer is the estimated watermarking subspace and the more efficient is the attack. Second, with a given number  $T$  of watermarked video frames, the greater is the dimension of the watermarking subspace  $\mathcal{W}$ , the harder it is to estimate [50].

## 5.4 Scalar Costa Schemes (SCS)

Although the presented study has been illustrated with additive spread-spectrum watermarks, it can also be extended to other reference watermarking schemes. Lately, a lot of attention has been paid to watermarking techniques whose goal is to approach the result predicted by Costa [30], that is to say that the capacity of the Gaussian data hiding channel is not affected by the host signal. The most popular techniques are the ones which are usually referred to as *quantization based* methods, since they approximate the random codebook used in Costa's proof by a structured one consisting of uniform quantizers. In the remainder of this section, it will rapidly be shown that such algorithms also leak some information that can be exploited to confuse the system. The so-called Scalar Costa Scheme (SCS) [65], also known as Quantization Index Modulation (QIM) or Dither Modulation (DM) [25], is one of the most popular quantization-based methods. The starting point is to quantize each input sample according to the message to be hidden:

$$\mathbf{y}(i) = Q_{\Lambda_{\mathbf{m}(i)}}(\mathbf{x}(i)) \quad (5.46)$$

where  $\mathbf{x}(i)$  (resp.  $\mathbf{y}(i)$ ) is the  $i$ -th sample of the original (resp. watermarked) signal and  $\Lambda_m$  a lattice shifted according to the  $i$ -th  $M$ -ary symbol  $\mathbf{m}(i)$  to be transmitted:

$$\Lambda_{\mathbf{m}(i)} \triangleq \Delta\mathbb{Z} - \mathbf{m}(i)\frac{\Delta}{M} \quad (5.47)$$


with  $\Delta$  being the quantization step. In other words, the whole embedding process can be written:

$$\mathbf{y}(i) = \Delta \left\lfloor \frac{\mathbf{x}(i) - \mathbf{m}(i) \frac{\Delta}{M}}{\Delta} + 0.5 \right\rfloor + \mathbf{m}(i) \frac{\Delta}{M} \quad (5.48)$$

Furthermore, to address imperceptibility requirement, a distortion compensation parameter  $\alpha$  which belongs to the interval  $[0, 1]$  is usually introduced and Equation (5.46) is replaced by the following one:

$$\mathbf{y}(i) = \mathbf{x}(i) + \alpha \left( \mathcal{Q}_{\Lambda_{\mathbf{m}(i)}}(\mathbf{x}(i)) - \mathbf{x}(i) \right) \quad (5.49)$$

In other words, the quantization error  $\mathcal{Q}_{\Lambda_{\mathbf{m}(i)}}(\mathbf{x}(i)) - \mathbf{x}(i)$  is embedded with a given embedding strength  $\alpha$ . It should be noted that Equation (5.46) corresponds to the specific case where  $\alpha = 1$ . On the receiver side, the detector simply quantizes each watermark sample with lattice having a quantization step equal to  $\Delta/M$ . The hidden  $M$ -ary symbol  $\mathbf{m}(i)$  is then determined according to the label of the quantization bin.



**Figure 5.8:** Probability density function of the watermark signal with a Gaussian host and a random binary message  $\mathbf{m}$ .

However, from a security perspective, if an attacker consider the probability density function of the watermarked signal, the lattice defined by Equation (5.47) can be somewhat easily estimated. The attacker can then exploit this knowledge to defeat the system, for instance by quantifying the watermarked signal with a lattice shifted by  $\Delta/2M$ . To avoid such a pitfall, a *dithering* term  $\mathbf{d}$  is usually introduced in Equation (5.47) as follows:

$$\Lambda_{\mathbf{m}(i)} \triangleq \Delta\mathbb{Z} - \mathbf{m}(i) \frac{\Delta}{M} - \mathbf{d}(i) \quad (5.50)$$

In other words, for each sample, the original lattice  $\Lambda_{\mathbf{m}(i)}$  is shifted by a random amount  $\mathbf{d}(i)$ . When this dithering term is made key-dependent, security is achieved since it randomizes the codebook. As a result, if an attacker examines the probability density function of the watermarked signal, he/she will only be able to recognize some Gaussian distribution as depicted in Figure 5.8. Nevertheless, if the attacker adopts a collusion strategy, the situation remains the same. For given sample position  $i$ , the lattice defined in Equation (5.50) remains the same for all the collected documents. Therefore, if the probability density function of the different watermarked samples at this specific location is considered, the attacker will be able to identify the characteristics of the lattice and to exploit this knowledge to confuse the system.

Finally, a common extension of SCS is to combine both spread spectrum and quantization-based watermarks. This is the so-called Spread Transform Scalar Costa Scheme (ST-SCS) or Spread Transform Dither Modulation (ST-DM) [25, 150]. In this case the correlations between the host signal and some pseudo-random reference patterns  $\mathbf{w}_i$  are quantized rather than the host sample values themselves. Nevertheless, the embedded watermark is then bound to a low-dimensional subspace  $\mathcal{W} = \text{span}(\mathbf{w}_i)$ . As a result, it can be estimated using the approach described in Subsection 5.3.3 and the attacker can estimate this knowledge to confuse the detector.

## 5.5 Discussion

Robustness is usually considered as a key-property for watermarking systems. However, it is only a first requirement when the watermarking technology is to be deployed in a hostile environment. In this case, malicious users will design some advanced dedicated attacks to defeat the system. The security issue has consequently to be addressed and malicious behaviors have to be anticipated. In this chapter, it has been demonstrated that using a redundant watermarking structure can lead to critical security pitfalls. Indeed, such strategies have been shown to leak information about the watermarking system and the attacker can exploit this knowledge to defeat the system. In other words, the attacker eavesdrops the watermarking channel to isolate suspicious redundant patterns and then uses this information to confuse the detector. Similar studies have been conducted independently which aim at quantifying the amount of information leakage using information theory [20, 21]. To this end, the ignorance about the system is evaluated through conditional entropy, which Shannon named *equivocation*:

$$H(K|\mathbf{o}_1, \dots, \mathbf{o}_N) = H(K) - I(K; \mathbf{o}_1, \dots, \mathbf{o}_N) \quad (5.51)$$

where  $\{\mathbf{o}_i\}$  is a set of watermarked documents and  $K$  the secret to be estimated, should it be a secret key or a secret pseudo-random sequence. In other words, the



information leakage is measured by the mutual information between the observations and the secret. Such an approach has the advantage to give some kind of objective metric to compare the security level of different watermarking systems.

Although many security pitfalls have been exhibited in this chapter, almost no countermeasure has been proposed. The geometrical perspective introduced in Subsection 5.3.1 can give some intuitive insight to define a secure embedding strategy for video watermarking. The trajectory defined by successive watermarks should be continuous, without any accumulation point and should cover the whole media space. However, such a theoretical statement does not give any clue on how such trajectories can be built in practice. All the presented watermarking systems can be labeled as *blind* as they do not in any way consider the data to be watermarked. Considering the host data may have a significant impact on performances and possible tracks for future work are given below:

1. *Anchor-based watermarks*: Security is somewhat related to statistical invisibility [176]. In such an approach, two watermarks should be as similar as the associated host video frames. An implementation of this idea consists in embedding small watermark patches at some anchor locations of the video frames [175]. These anchor points should be pseudo-secret, and also host signal dependent.
2. *Image signature*: Another approach to obtain such coherent watermarks exploits key-dependent image signatures [70, 41]. The goal is to obtain binary strings related with the host content i.e. image signatures should be as correlated as the associated images. They can then be used to generate a watermark pattern which degrades gracefully with an increased number of bit errors.
3. *Informed coding*: Recently, dirty paper codes [25, 65, 136] have been explored to make the embedded watermark dependent of the host signal. Basically, for a given payload, a constellation of possible watermarks is defined on the unit sphere and the nearest watermark from the host signal is embedded. As a result, the induced watermark trajectory varies as smoothly as the host content and links several points of the constellation. Furthermore, recent studies [21] have reported that trellis dirty-paper watermarks [136] are more secure than other watermarks.



---

## Jamming the Watermarking Channel

---

The previous chapter has stressed the fact that using a redundant watermarking structure is likely to induce some information leaks. Considering multiple watermarked contents, a hostile attacker is able to gain some knowledge about the embedded watermark signal and exploit it to confuse the detector. Nevertheless, completely independent watermarks are not the solution. If an attacker can collect similar contents carrying uncorrelated watermarks, averaging them will sum the watermark samples to zero. Multimedia digital data is highly redundant: successive video frames are highly similar in a movie clip, most songs contain some repetitive patterns, etc. This property can consequently be exploited to successively replace each part of the signal with a similar one taken from another location in the same signal or with a combination of similar parts. Such an approach is all the more pertinent when video content is considered since such signals exhibit both temporal and spatial self-similarities. In Section 6.1, temporal redundancy across successive video frames is exploited to confuse the watermark detector. The basic idea consists in approximating the background of each video frame, using information contained in the neighbor ones. The enforced strategy basically comes down to Temporal Frame Averaging after Registration (TFAR). Next, Section 6.2 highlights that each single video frame contains also some spatial self-similarities. This is typically the property exploited in fractal coding. Several methods to design efficient Block Replacement Attacks (BRA) are then presented.

## 6.1 Temporal Frame Averaging after Registration

Temporal Frame Averaging (TFA) has been depicted in Figure 5.3 as a possible collusion strategy to confuse watermark detectors. Nevertheless, a major shortcoming of this approach is that it is limited by the content of the considered video. When the scene consists of dynamic content, e.g. fast moving object and/or camera motion, video frames cannot be directly temporally averaged without strongly degrading the video quality. Although neighboring frames are highly correlated, they still require to be registered to permit efficient averaging [49, 51]. Each video frame is indeed a projection of a single 3D movie set and different video frames from a shot can be seen as different 2D projections of the same scene. Thus, frame registration can be exploited to bring all these projections onto the same reference frame so that all the projections of a given 3D point overlap. As a result, temporal averaging can be done with a large temporal window without introducing much visual distortion. A detailed description of Temporal Frame Averaging after Registration (TFAR) is given in Subsection 6.1.1 and the whole process is depicted in Figure 6.1. Furthermore, the performances of this attack are reported in Subsection 6.1.2.

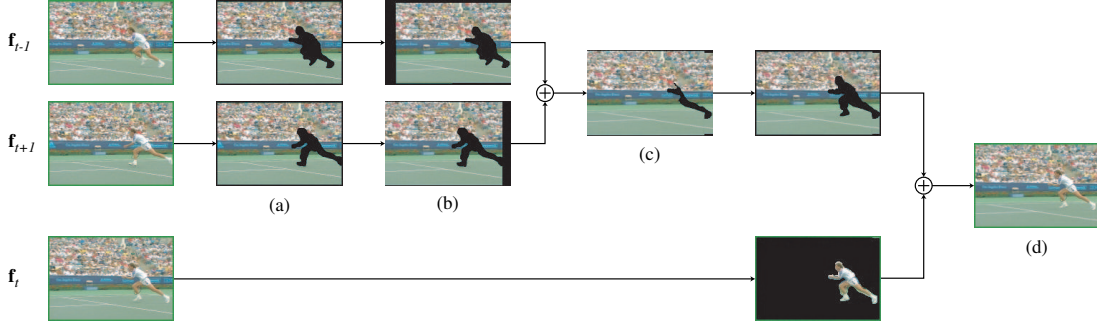
### 6.1.1 Attack Description

The goal is to estimate a given video frame  $\mathbf{f}_t$  from its neighboring ones  $\mathbf{f}_{t+\delta}$  thanks to frame registration. However these frames may contain objects which cannot be used to reconstruct the target video frame. As a result, a binary mask  $\mathbf{m}_t$  has to be built for each frame to distinguish useful areas in the frame (e.g. the background) from useless ones (e.g. moving objects). This mask is somewhat similar to the Video Object Plane (VOP) in the MPEG-4 video coding standard [105]. Once this mask has been defined, the background  $\mathbf{b}_t$  and the moving objects  $\mathbf{o}_t$  can be retrieved using the following equations:

$$\mathbf{o}_t = \mathbf{f}_t \otimes \mathbf{m}_t \quad \text{and} \quad \mathbf{b}_t = \mathbf{f}_t \otimes \bar{\mathbf{m}}_t \quad (6.1)$$

where  $\otimes$  is the pixel-wise multiplication operator and  $\bar{\cdot}$  is the binary negation operator. No specific work has been done to design an object-based segmentation in this thesis and an existing algorithm based on semi-automatic initial segmentation of the first video frame, followed by an automatic tracking of the selected objects [171] has been reused.

Once several observations  $\mathbf{b}_{t'}$  of the movie set have been obtained from neighboring frames, they can be exploited to estimate the background  $\tilde{\mathbf{b}}_t$  of the current frame. To this end, it is necessary to find a registration function which *pertinently* associates to each pixel position  $(x_t, y_t)$  in the current frame  $\mathbf{f}_t$  a position  $(x_{t'}, y_{t'})$  in a neighboring frame  $\mathbf{f}_{t'}$  i.e. which minimizes for example the mean square error between the target background  $\mathbf{b}_t$  and the registered one  $\mathbf{b}_{t'}^{(t)}$ . In other words, the



**Figure 6.1:** Temporal Frame Averaging after Registration (TFAR): Once the video objects have been removed (a), neighbor frames are registered (b) and combined to estimate the background of the current frame (c). Next, the missing video objects are inserted back (d). In this illustration, the temporal window half-size  $w$  is equal to 1.

goal is to define a model which describes the apparent displacement generated by the camera motion. Physically, camera motion is a combination of traveling displacements (horizontal, vertical, forward and backward translations), rotations (pan, roll and tilt) and zooming effects (forward and backward). As the background of the scene is often far from the camera, pan and tilt rotations can be assimilated, for small rotations, to translations in terms of 2D apparent motion. Thus, the zoom, roll and traveling displacements can be represented, under some assumptions, by a first order polynomial motion model [143] as follows:

$$\begin{cases} x_{t'} = t_x + z(x_t - x_o) - z\theta(y_t - y_o) \\ y_{t'} = t_y + z(y_t - y_o) + z\theta(x_t - x_o) \end{cases} \quad (6.2)$$

where  $z$  is the zoom factor,  $\theta$  the 2D rotation angle,  $(t_x, t_y)$  the 2D translational vector and  $(x_o, y_o)$  the coordinates of the camera optical center. Obviously, this model is quite simple and may not be accurate when the camera displacement or the scene structure is very complicated. More complex motion representations can be introduced such as the affine model [143], the projection model [183] or the trifocal motion model [179]. Nevertheless, the model described in Equation (6.2) has been used in this thesis for simplicity reasons.

The computed registered backgrounds  $\mathbf{b}_{t+\delta}^{(t)}$ , obtained from the video frames in the temporal window, are then combined to obtain an estimation  $\tilde{\mathbf{b}}_t$  of the background in the current frame. For each pixel position  $\mathbf{p}$  in the frame, the value of the background is estimated:

$$\tilde{\mathbf{b}}_t(\mathbf{p}) = \begin{cases} \frac{\sum_{\delta \in [-w, w]^*} \mathbf{b}_{t+\delta}^{(t)}(\mathbf{p})}{\sum_{\delta \in [-w, w]^*} \mathbf{m}_{t+\delta}^{(t)}(\mathbf{p})} & \text{if the denominator is not equal to 0} \\ 0 & \text{otherwise} \end{cases} \quad (6.3)$$

where  $\bar{\mathbf{m}}_t^{(t)}$  is the registered binary mask and  $w$  the temporal window half-size. In other words, the registered backgrounds are temporally averaged using the proper normalization factor. A binary mask  $\mathbf{r}_t$  is also built to indicate, for each pixel position, whether a background value has been effectively estimated ( $\mathbf{r}_t(\mathbf{p}) = 1$ ) or not ( $\mathbf{r}_t(\mathbf{p}) = 0$ ). The whole reconstruction process can then be written as follows:

$$\hat{\mathbf{f}}_t = \underbrace{\tilde{\mathbf{b}}_t \otimes \bar{\mathbf{m}}_t}_{\text{Background}} + \underbrace{\mathbf{f}_t \otimes \mathbf{m}_t}_{\text{Objects}} + \underbrace{\mathbf{f}_t \otimes (\bar{\mathbf{m}}_t \& \bar{\mathbf{r}}_t)}_{\text{Missing pixels}} \quad (6.4)$$

where  $\&$  is the binary AND operator. The first term is associated with the current estimated background: pixel values have to be discarded if the related positions do not belong to the current background binary mask  $\bar{\mathbf{m}}_t$ . The second term indicates that moving video objects  $\mathbf{o}_t$  from the original frame are inserted back. The last term in Equation (6.4) points out that, at this point, some background pixels may have not been estimated. In this case, the pixel values from the original video frame  $\mathbf{f}_t$  are retrieved. It should be noted that this attack does not affect the moving video objects  $\mathbf{o}_t$ . As a result, if such objects occupy most of the video scene, the attack is not likely to trap the detector. However, the background is usually the main part in many video shots and the attack is still pertinent.

From a coding perspective, the presented TFAR attack can be seen as encoding the background with an advanced forward-backward predictive coder e.g. B-frames in MPEG. Alternatively, it can also be considered as temporal averaging along the motion axis. Whatever, since most watermarking algorithms do not consider the evolution of the structure of the scene during embedding, such a processing is likely to confuse the watermark detector as it will be verified in the next subsection. Skeptical people might argue that such attacks are too computationally intensive to be realistic. However, video mosaics or sprite panoramas are expected to be exploited for efficient background compression in the upcoming video standard MPEG-4 and such video coding algorithms will have a similar impact on embedded watermarks [105].

### 6.1.2 TFAR Evaluation

The two reference SS and SS-1 systems have been considered to evaluate the impact of TFAR. Although they have already been presented in Subsection 5.1.1, their description is briefly reminded below. They almost share the same embedding procedure. Indeed, they both add a normally distributed with zero mean and unit variance watermark signal  $\mathbf{w}_t(K)$  to each video frame  $\mathbf{f}_t$  as follows:


$$\tilde{\mathbf{f}}_t = \mathbf{f}_t + \alpha \mathbf{w}_t, \quad \mathbf{w}_t \sim \mathcal{N}(0, 1) \quad (6.5)$$

where  $\alpha$  is some embedding strength and  $K$  a secret key. The only difference is that the SS-1 system always embeds the same reference watermark pattern

( $\mathbf{w}_t = \mathbf{w}_o$  for all  $t$ ) whereas the SS system always embeds a different watermark ( $\mathbf{w}_t \neq \mathbf{w}_{t'}$  when  $t \neq t'$ ). A common variation is to make the embedding strength  $\alpha$  dependent of the local content of the video frame for perceptual reasons [189]. However, in practice a global embedding strength equal to 3 has been used so that the resulting distortion is around 38 dB in terms of PSNR. On the receiver side, the detector simply checks if each input video frame  $\hat{\mathbf{f}}_t$  contains the expected watermark using a simple linear correlation:

$$\rho(\hat{\mathbf{f}}_t, \mathbf{w}_t) = \hat{\mathbf{f}}_t \cdot \mathbf{w}_t = \mathbf{f}_t \cdot \mathbf{w}_t + \epsilon \alpha \mathbf{w}_t \cdot \mathbf{w}_t \approx \epsilon \alpha \quad (6.6)$$

where  $\epsilon$  is equal to 0 (resp. 1) when a watermark has not (resp. has) been embedded. Host interference can be cancelled during embedding to enhance detection statistics [34]. Hence, the correlation score should be almost equal to  $\alpha$  (resp. 0) when the watermark is present (resp. absent). It is enough to compare this score to a threshold  $\tau_{\text{detect}}$  to assert the presence or absence of the watermark. In practice, this threshold is set equal to  $\alpha/2$  for equal false positive and false negative probabilities.



**Figure 6.2:** Impact of TFAR on SS and SS-1 systems.

Two different video sequences have been considered for the experiments. The first one has been artificially created from a panoramic view of the *Cap d'Antibes*. It consists of 141 frames of size  $352 \times 240$  and the synthetic displacement has been set to 2 pixels per frame horizontal translation from right to left. In this case, the displacements are exactly known and frame registration is perfect. The second surveyed video is the well-known *Stefan* sequence which consists of 91 frames of size  $352 \times 240$ . In this case, registration parameters have to be estimated and frame registration is not perfect. In practice, previous work from Professor Henri Nicolas, IRISA, France has been exploited to obtain these parameters [143, 142]. Both videos have been watermarked using SS and SS-1 embedding strategies and

each watermarked video has then been submitted to TFAR. The correlation score has been computed before and after TFAR and the results are reported in Figure 6.2. It is immediate to see that TFAR succeeds in removing the embedded watermark. Whereas the detection scores oscillates around  $\alpha$  before TFAR, it almost drops down to 0 after the attack. The only noticeable singularity is around frame 5 with the video *Stefan* when the SS-1 strategy is enforced. At this very moment, there is almost no camera motion and the video is static. Therefore, TFAR has no impact and the correlation score bumps above the detection threshold  $\tau_{\text{detect}}$ . This is due to the fact that the SS-1 strategy is optimal with respect to TFAR when the video sequence is static.

## 6.2 Block Replacement Attack

If temporal redundancy can be easily exhibited in successive video frames as noticed in the previous section, less obvious self-similarities are also present in each single video frame. Such spatial self-similarities have previously been exploited to obtain efficient compression tools [68]. As a result, using this kind of redundancy, an attacker can design a Block Replacement Attack (BRA) which replaces each input signal block with another one which is similar to the input block. Alternative methods to obtain such valid replacement blocks will be reviewed in the upcoming subsections. There exists indeed a trade-off between fidelity and attack efficiency. On one hand, the replacement block should be similar enough to the input one so that the attack does not introduce perceptible artifacts. On the other hand, if the replacement is too close from the input one, it is also likely to still carry the watermark signal. As a result, several methods will be surveyed to try to optimize this trade-off. Finally, the performances of such BRA are examined through intensive testing.

### 6.2.1 Block Restoration

Error concealment techniques have initially been designed to recover blocks which have been lost or corrupted during digital transmission. As depicted in Figure 6.3, when a missing block is detected, the neighborhood of this block is considered to obtain a prediction of the missing information. Such approaches can be exploited to design an efficient block replacement attack. Sequentially, each block of the signal is considered as missing and the error concealment procedure computes a replacement block [199]. However, this strategy suffers from two major shortcomings. First, there is no direct attacking strength i.e. there is no possibility to adapt the impact of the attack according to the watermarking strength. Second, each block is considered as *missing* which is not really the case. In other words, some information is ignored and it is likely to result in a relatively poor qual-



ity attacked signal. For both those reasons, such approaches will not be further considered in the remainder of the thesis.



**Figure 6.3:** Error concealment techniques: when a block is detected as corrupted or missing, it is discarded and the algorithm tries to predict it using blocks in the vicinity.

### 6.2.2 Block Swapping

Most watermarking algorithms have exhibited weaknesses against desynchronization attacks and especially non global ones. The random bending attack [151] has been considered for a long time now as a reference for benchmarking watermarking systems. However, countermeasures have appeared which basically exploit the fact that this processing does not drastically modify the *geography* of the embedded watermark. Each watermark sample is slightly displaced but it remains in the neighborhood of its original location. As a result, local block-matching based detectors [80, 100, 156] have been shown to be able to recover watermarks altered by such attacks. Alternatively, the block swapping attack [158] aims at shuffling the watermark samples while keeping the host data synchronized. The basic idea is to replace each block of the signal with a similar one, which *does not carry the same watermark signal*. In other words, the *geography* of the embedded watermark is strongly altered so that resynchronization is no longer possible, and thus the detector is confused.

The pseudo-code of the block swapping attack is detailed in Table 6.1. For each block  $\mathbf{b}_T$  of the input signal, a search window is defined and a codebook  $\mathcal{Q}$  built. Next, photometric compensation is necessary, at least with still images,

**Table 6.1:** Block swapping attack.

For each block $\mathbf{b}_T$ of the signal	
1	Build the block codebook $\mathcal{Q}$
2	Perform photometric compensation
3	Sort the blocks $\mathbf{b}_{Q_i}$ according to the MSE
4	Set $\mathbf{b}_R$ as the most similar block
5	Replace $\mathbf{b}_T$ by $\mathbf{b}_R$

to obtain a good pool of candidate blocks for replacement. Otherwise, the codebook  $\mathcal{Q}$  is unlikely to contain a block which is similar enough to  $\mathbf{b}_T$  and the replacement process will introduce a strong distortion. As a result, each block  $\mathbf{b}_{Q_i}$  is transformed in  $s\mathbf{b}_{Q_i} + o\mathbf{1}$ , where  $\mathbf{1}$  is a block containing only ones, so that the Mean Square Error (MSE) with the target block  $\mathbf{b}_T$  is minimized. This is a simple least squares problem and the scale  $s$  and offset  $o$  can be determined as follows:

$$s = \frac{(\mathbf{b}_T - m_T\mathbf{1}) \cdot (\mathbf{b}_{Q_i} - m_{Q_i}\mathbf{1})}{|\mathbf{b}_{Q_i} - m_{Q_i}\mathbf{1}|^2} \quad (6.7)$$

$$o = m_T - s \cdot m_{Q_i} \quad (6.8)$$


where  $m_T$  (resp.  $m_{Q_i}$ ) is the mean value of block  $\mathbf{b}_T$  (resp.  $\mathbf{b}_{Q_i}$ ),  $\cdot$  is the linear correlation defined as:

$$\mathbf{b} \cdot \mathbf{b}' = \frac{1}{S_T} \sum_{i=1}^{S_T} \mathbf{b}(i)\mathbf{b}'(i) \quad (6.9)$$

and  $|\mathbf{b}|$  is the norm defined as  $\sqrt{\mathbf{b} \cdot \mathbf{b}}$ . At this point, the transformed blocks  $s\mathbf{b}_{Q_i} + o\mathbf{1}$  are sorted in ascending order according to their similarity with the target block  $\mathbf{b}_T$ . The most similar block is then retained and used for replacement. In this version, the block replacement attack is equivalent to image compression with a fractal coder [68]. A visual interpretation of this attack is depicted in Figure 6.4. In the same fashion, an alternative approach consists in building iteratively sets of similar blocks and randomly shuffling their positions [152, 102] until all the blocks have been replaced.

Performing photometric compensation and computing the MSE can become computationally prohibitive as the number of blocks in the codebook  $\mathcal{Q}$  increases. Furthermore, there is no real need to perform explicitly photometric compensation for each block  $\mathbf{b}_{Q_i}$ . In fact, photometric compensation needs to be done only for a single block of the codebook, the one which will be used for replacement. There exists a relationship between  $\text{MSE}(s\mathbf{b}_{Q_i} + o\mathbf{1}, \mathbf{b}_T)$  and the correlation coefficient:

$$\mathbf{b}_{Q_i} \odot \mathbf{b}_T = \frac{\mathbf{b}_{Q_i} - m_{Q_i}\mathbf{1}}{|\mathbf{b}_{Q_i} - m_{Q_i}\mathbf{1}|} \cdot \frac{\mathbf{b}_T - m_T\mathbf{1}}{|\mathbf{b}_T - m_T\mathbf{1}|} \quad (6.10)$$



**Figure 6.4:** Block swapping attack: each block is replaced by the one in the search window which is the most similar modulo a geometrical and photometric transformation.

After a few derivations, the following equation can be obtained:


$$\text{MSE}(s\mathbf{b}_{Q_i} + o\mathbf{1}, \mathbf{b}_T) = |\mathbf{b}_T - m_T\mathbf{1}|^2 \left(1 - (\mathbf{b}_{Q_i} \odot \mathbf{b}_T)^2\right) \quad (6.11)$$

It means that sorting the blocks in ascending  $\text{MSE}(s\mathbf{b}_{Q_i} + o\mathbf{1}, \mathbf{b}_T)$  is equivalent to sorting the blocks in descending  $(\mathbf{b}_{Q_i} \odot \mathbf{b}_T)^2$ . This property can be exploited to sort the blocks of the codebook without explicitly building the modified blocks  $s\mathbf{b}_{Q_i} + o\mathbf{1}$ .

Exchanging highly similar blocks is likely to be imperceptible. However, it is also likely not to affect the watermark signal. A threshold  $\tau_{\text{low}}$  can consequently be introduced to force a minimum distortion between the replacement block  $\mathbf{b}_R$  and the target block  $\mathbf{b}_T$  which is to be replaced. In other words, the step [3](#) is modified so that the replacement block is no longer the most similar block in the codebook  $\mathcal{Q}$  modulo a geometrical and photometric transformation, but rather the most similar block *whose distortion is if possible above*  $\tau_{\text{low}}$ . This additional parameter can be regarded as an attacking strength and introduces a trade-off between the efficiency of the attack and its impact in terms of fidelity.

### 6.2.3 Blocks Combination

In the previous subsection, a threshold  $\tau_{\text{low}}$  has been introduced to ensure that the replacement block  $\mathbf{b}_R$  is not similar to the target block  $\mathbf{b}_T$  up to the point that it also contains the watermark signal. On the other hand, there is no guaranty that this replacement block will be similar enough to be imperceptible once it



**Figure 6.5:** Influence of the number of blocks  $N$  used for combination once the thresholds  $\tau_{\text{low}}$  and  $\tau_{\text{high}}$  have been set. Light (resp. dark) gray blocks indicate *too good* (resp. *too bad*) blocks.

has been substituted with the original block. In fact, experimental results have shown that blocks are likely to be badly estimated with a single block, even if photometric compensation is performed. Following previous advances in fractal coding [73, 145], the idea is then to combine several blocks  $\mathbf{b}_{Q_i}$  in the codebook  $\mathcal{Q}$  to obtain a better replacement block:

$$\mathbf{b}_R = \sum_{i=1}^N \lambda_i \mathbf{b}_{Q_i} \quad (6.12)$$

where the  $\lambda_i$ 's are mixing coefficients. To obtain the best possible replacement block, those mixing coefficients are chosen so that  $\text{MSE}(\mathbf{b}_R, \mathbf{b}_T)$  is minimized<sup>1</sup>. This is a traditional least squares problem which can be easily solved using common linear algebra tools. From this novel perspective, the block replacement attack is more related with intra-signal collusion attacks [53] i.e. combining several watermarked contents to obtain unwatermarked content.

### Fixed number of blocks:

With this new approach in mind, a novel block replacement attack can be designed as depicted by the pseudo-code given in Table 6.2. For each block  $\mathbf{b}_T$  of the input signal, a search window is defined and a codebook  $\mathcal{Q}$  built. Then the blocks  $\mathbf{b}_{Q_i}$  are sorted in ascending order according to their similarity with the target block  $\mathbf{b}_T$  using Equation (6.11). At this point, a *fixed* number of blocks e.g. the first  $N$  blocks in the codebook  $\mathcal{Q}$  are considered to compute an optimal replacement block  $\mathbf{b}_R$  in a least squares sense [103]. Finally, this candidate block is put into the place of the original one  $\mathbf{b}_T$ . Here again, the step [3] can be modified to prevent the candidate replacement block  $\mathbf{b}_R$  from being too similar

<sup>1</sup>It should be noted that the block  $\mathbf{1}$  can be artificially added to the codebook  $\mathcal{Q}$  to permit automatic mean value adjustment.

to the target block  $\mathbf{b}_T$ . To this end, the first  $N$  blocks *whose distortion is above a threshold*  $\tau_{\text{low}}$  can be considered, rather than the  $N$  first ones, to compute the optimal replacement block. The expectation is that using poorer blocks from the codebook  $\mathcal{Q}$  will output a poorer candidate block for replacement.

**Table 6.2:** Fixed number of blocks combination attack.

For each block $\mathbf{b}_T$ of the signal	
1	Build the block codebook $\mathcal{Q}$
2	Sort the blocks $\mathbf{b}_{Q_i}$
3	Build the optimal replacement block $\mathbf{b}_R$ using the first $N$ blocks in $\mathcal{Q}$
4	Replace $\mathbf{b}_T$ by $\mathbf{b}_R$

However, the attacker would rather like to be able to ensure that the final distortion  $\text{MSE}(\mathbf{b}_R, \mathbf{b}_T)$  is between the two thresholds  $\tau_{\text{low}}$  and  $\tau_{\text{high}}$ . Indeed, the replacement block should not be *too good* ( $\text{MSE}(\mathbf{b}_R, \mathbf{b}_T) < \tau_{\text{low}}$ ). Otherwise, it is also likely to carry the watermark signal. Furthermore, it should not be *too bad* either ( $\text{MSE}(\mathbf{b}_R, \mathbf{b}_T) > \tau_{\text{high}}$ ) so that the block replacement attack does not introduce perceptible artifacts. Unfortunately, it is difficult to predict this distortion from the similarity of the blocks  $\mathbf{b}_{Q_i}$  used for combination with the target block  $\mathbf{b}_T$ . It can only be checked a posteriori. Figure 6.5 shows the localization of the *too good* and *too bad* blocks once the two thresholds  $\tau_{\text{low}}$  and  $\tau_{\text{high}}$  have been fixed and that the number  $N$  of blocks used for combination is varying. The first observation is that the number of *too bad* blocks decreases as  $N$  increases, while the number of *too good* blocks increases. Secondly, the number of blocks needed to make the distortion  $\text{MSE}(\mathbf{b}_R, \mathbf{b}_T)$  drop below  $\tau_{\text{high}}$  seems to be related with the content of blocks: flat blocks require fewer blocks to obtain a valid replacement block  $\mathbf{b}_R$  after combination in comparison with textured blocks. This calls for a new approach which automatically adjusts the number of blocks used for combination with respect to the content of the block.

### Adaptive number of blocks:

The previous subsection has highlighted the fact that using a *fixed* number of blocks is somewhat limiting. Each block does not indeed need the same number of blocks to be finely enough approximated e.g. flat vs. textured blocks. An improved algorithm whose pseudo-code is given in Table 6.3 is consequently introduced so that the number and the set of blocks chosen for combination are adaptively modified to obtain a candidate replacement block  $\mathbf{b}_R$  whose distortion  $\text{MSE}(\mathbf{b}_R, \mathbf{b}_T)$  is between  $\tau_{\text{low}}$  and  $\tau_{\text{high}}$ . The basic idea is to modify the step 3 in the previous algorithm by checking the distortion  $\Delta = \text{MSE}(\mathbf{b}_R, \mathbf{b}_T)$  of the computed candidate block for replacement. Depending on the value of this

distortion, different rules are enforced.

- If  $\Delta$  is between  $\tau_{\text{low}}$  and  $\tau_{\text{high}}$ , a valid candidate block has been found for replacement. A flag is consequently set to 1 to terminate the adaptive algorithm.
- If  $\Delta$  is greater than  $\tau_{\text{high}}$ , it means that the obtained candidate block for replacement does not approximate the target block  $\mathbf{b}_T$  well enough. As a result, the attack would introduce perceptible artifacts if they were substituted.  $N$  is consequently incremented so that more blocks are considered during combination and thus a better candidate block is obtained.
- If  $\Delta$  is lower than  $\tau_{\text{low}}$ , the candidate replacement block  $\mathbf{b}_R$  is too similar to the target block  $\mathbf{b}_T$ . It is likely to also carry the watermark signal. The offset  $\Phi$  is consequently incremented so that poorer blocks from  $\mathcal{Q}$  are considered during block combination. Furthermore, the number of combined blocks  $N$  is reset to 1.

It should be noted that this algorithm inherently assumes that a candidate block whose distortion falls within the bounds  $\tau_{\text{low}}$  and  $\tau_{\text{high}}$  will be found. However, nothing ensures that it will be the case in practice. In particular, for small codebooks or close threshold values, such a block might not exist. The algorithm consequently needs to be slightly modified to handle such exceptions. For example, if this case occurs, the candidate block whose distortion minimizes  $\max(\sqrt{\tau_{\text{low}}} - \sqrt{\Delta}, \sqrt{\Delta} - \sqrt{\tau_{\text{high}}})$  can be retained for replacement.

**Table 6.3:** Adaptive number of blocks combination attack.

---

For each block $\mathbf{b}_T$ of the signal	
1	Build the block codebook $\mathcal{Q}$
2	Sort the blocks $\mathbf{b}_{Q_i}$ Set $\Phi = 0$ , $N = 1$ , flag = 0
3	While (flag = 0) AND ( $\Phi + N \leq  \mathcal{Q} $ ) <ul style="list-style-type: none"> <li>(a) Build the optimal replacement block <math>\mathbf{b}_R</math> using <math>N</math> successive blocks from <math>\mathcal{Q}</math> starting with block <math>\mathbf{b}_{Q_{\Phi+1}}</math></li> <li>(b) Compute <math>\Delta = \text{MSE}(\mathbf{b}_R, \mathbf{b}_T)</math></li> <li>(c) If <math>\tau_{\text{low}} \leq \Delta \leq \tau_{\text{high}}</math>, set flag = 1</li> <li>(d) Else if <math>\Delta &gt; \tau_{\text{high}}</math>, increment <math>N</math></li> <li>(e) Else increment <math>\Phi</math> and reset <math>N</math> to 1</li> </ul>
4	Replace $\mathbf{b}_T$ by $\mathbf{b}_R$

---

### 6.2.4 Block Projection

The previous attack gives some good results as will be reported in Subsection 6.2.5. However, it is in some sense suboptimal. In step [3], when the computed candidate block for replacement is found to be too similar to the target block  $\mathbf{b}_T$ , the offset  $\Phi$  is incremented to consider poorer blocks during combination. Nevertheless, this does not ensure that a poorer block will be obtained *after combination*. In fact, this is only a way of getting alternative candidate blocks for replacement until one is found to be in the target interval  $[\tau_{\text{low}}, \tau_{\text{high}}]$ . In this case, all the possible blocks combinations should be computed instead of a restricted subset. But this is not possible in practice because of the prohibitive computational cost. As a result, a substitute approach is investigated below.

From a geometrical point of view, finding the mixing coefficients  $\lambda_i$  which minimize the distortion  $\text{MSE}(\sum_{i=1}^N \lambda_i \mathbf{b}_{Q_{\Phi+i}}, \mathbf{b}_T)$  is equivalent to computing the coordinates of the target block  $\mathbf{b}_T$  in the subspace spanned by the  $N$  blocks  $\mathbf{b}_{Q_{\Phi+i}}$ . In other words, the block replacement attack comes down to finding a subspace  $\mathcal{S}$  for each block  $\mathbf{b}_T$  so that  $\text{MSE}(\mathbf{b}_T^{\mathcal{S}}, \mathbf{b}_T)$  is between  $\tau_{\text{low}}$  and  $\tau_{\text{high}}$ , where  $\mathbf{b}_T^{\mathcal{S}}$  is the projection of the block  $\mathbf{b}_T$  onto the subspace  $\mathcal{S}$ . In the approaches described in Subsection 6.2.3, most of the computational cost is due to the fact that the basis vectors of the subspace  $\mathcal{S}$  - in this case the blocks  $\mathbf{b}_{Q_i}$  of the codebook  $\mathcal{Q}$  - are not orthogonal. Thus, a least squares problem has to be solved to obtain the coordinates  $\lambda_i$ 's of the target block in the generated subspace  $\mathcal{S} = \text{span}\{\mathbf{b}_{Q_i}\}$ . The problem would have been much easier if the basis vectors were orthogonal: successive projections on each vector gives then the coordinates. This has consequently motivated further research to investigate how to obtain such an orthogonal basis. In particular, approaches exploiting Gram-Schmidt Orthonormalization (GSO) and Principal Component Analysis (PCA) have been surveyed.

#### Gram-Schmidt Orthonormalization

The Gram-Schmidt orthonormalization procedure takes a non-orthogonal set of linearly independent vectors and constructs an orthogonal basis [29]. Thus, the goal is to incorporate it into a framework which iteratively builds an orthogonal basis in a *best possible match* fashion. First, the algorithm finds the block  $\mathbf{b}_{Q_i}$  in  $\mathcal{Q}$  which minimizes:

$$\text{MSE}(\mathbf{b}_T, \lambda_i \mathbf{b}_{Q_i}) \quad \text{with} \quad \lambda_i = \frac{\mathbf{b}_T \cdot \mathbf{b}_{Q_i}}{|\mathbf{b}_{Q_i}|^2} \quad (6.13)$$

Once this optimal block has been found, it is inserted into the basis  $\{\mathbf{s}_i\}$  which spans the subspace  $\mathcal{S} = \text{span}\{\mathbf{s}_i\}$ . Next, both the target block  $\mathbf{b}_T$  and the codebook  $\mathcal{Q}$  are projected onto the subspace orthogonal to  $\mathcal{S}$  as follows:

$$\mathbf{b}^{\mathcal{S}^\perp} = \mathbf{b} - \sum_{\mathbf{s}_i \in \mathcal{S}} \frac{\mathbf{b} \cdot \mathbf{s}_i}{|\mathbf{s}_i|^2} \mathbf{s}_i \quad (6.14)$$

where  $\mathbf{b}$  is some original input block and  $\mathbf{b}^{\mathcal{S}^\perp}$  its projection on  $\mathcal{S}^\perp$ . Then, the algorithm search for the best block as in Equation (6.13) and it is inserted into the basis which spans the subspace  $\mathcal{S}$ . The algorithm iterates until the distortion  $\text{MSE}(\mathbf{b}_T, \mathbf{b}_T^{\mathcal{S}})$  between the target block  $\mathbf{b}_T$  and its projection on the constructed subspace  $\mathcal{S}$  falls within the interval  $[\tau_{\text{low}}, \tau_{\text{high}}]$ . Nevertheless, this approach has two major shortcomings. First, the whole procedure requires many projection and correlation computations, which is likely to rapidly become intractable as the size of the codebook grows. Second, the basis is iteratively built in a *best possible match* way. However, nothing ensures that combining two blocks, which have been successively found to be the best possible match, will output a better candidate block than another combination of two blocks in the codebook.

### Principal Component Analysis

Principal Component Analysis [88] basically takes a set of vectors and outputs its centroid and a set of eigenvectors which can be seen as the directions of variations of the vectors in the set. Furthermore, each eigenvector is associated with an eigenvalue which indicates how much the set of vectors varies in this direction. The higher the eigenvalue is, the more variations there are in the associated direction. Such a procedure can be exploited to design an efficient block replacement attack as depicted in Table 6.4. First, a PCA is performed considering the different blocks  $\mathbf{b}_{Q_i}$  in the codebook  $\mathcal{Q}$ . This gives a centroid  $\mathbf{c}$  defined as follows:

$$\mathbf{c} = \frac{1}{|\mathcal{Q}|} \sum_{\mathbf{b}_{Q_i} \in \mathcal{Q}} \mathbf{b}_{Q_i} \quad (6.15)$$

and a set of eigenblocks  $\mathbf{e}_i$  associated with their eigenvalues  $\varepsilon_i$ . Those eigenblocks are then sorted by descending eigenvalues i.e. the direction  $\mathbf{e}_1$  contains more information than any other one in the basis. Then, a candidate block for replacement  $\mathbf{b}_R$  is computed using the  $N$  first eigenblocks so that the distortion with the target block  $\mathbf{b}_T$  is minimized. In other words, the block  $\mathbf{b}_T - \mathbf{c}$  is projected onto the subspace spanned by the  $N$  first eigenblocks. As a result, the replacement block can be written:

$$\mathbf{b}_R = \mathbf{c} + \sum_{i=1}^N \frac{(\mathbf{b}_T - \mathbf{c}) \cdot \mathbf{e}_i}{|\mathbf{e}_i|^2} \mathbf{e}_i \quad (6.16)$$

Of course, the distortion  $\Delta = \text{MSE}(\mathbf{b}_T, \mathbf{b}_R)$  gracefully decreases as the number  $N$  of combined eigenblocks increases. Thus, an adaptive framework is introduced to identify which value  $N$  should have so that the distortion  $\Delta$  falls within the range  $[\tau_{\text{low}}, \tau_{\text{high}}]$ . It may happen that no value of  $N$  gives a candidate block within this interval. In this case, the value  $N$  which gives the candidate block whose distortion minimizes  $\max(\sqrt{\tau_{\text{low}}} - \sqrt{\Delta}, \sqrt{\Delta} - \sqrt{\tau_{\text{high}}})$  is retained. The major



interest of this method is that it considers the *whole* codebook  $\mathcal{Q}$  to compute the orthogonal basis used for projection. Furthermore, experiments have shown that it was slightly quicker than the attack presented in Subsection 6.2.3. It should be noted that the underlying assumption is that most of the watermark energy will be concentrated in the last eigenblocks since the watermark can be seen as details. As a result, if a valid candidate block can be built without using the last eigenblocks, the watermark signal will not be reintroduced.

**Table 6.4:** Block projection on a PCA-defined subspace attack.

---

For each block $\mathbf{b}_T$ of the signal	
1	Build the block codebook $\mathbf{Q}$
2	Perform photometric compensation
3	Performs the PCA of the blocks in $\mathcal{Q}$ to obtain a set of orthogonal eigenblocks $\mathbf{e}_i$ associated with their eigenvalues $\varepsilon_i$ Set $N = 1$ , flag = 0
4	While (flag = 0) AND ( $N \leq S_T$ ) <ul style="list-style-type: none"> <li>(a) Build the optimal replacement block <math>\mathbf{b}_R</math> using the eigenblocks <math>\mathbf{e}_i</math> associated with the first <math>N</math> eigenvalues</li> <li>(b) Compute <math>\Delta = \text{MSE}(\mathbf{b}_R, \mathbf{b}_T)</math></li> <li>(c) If <math>\tau_{\text{low}} \leq \Delta \leq \tau_{\text{high}}</math>, set flag = 1</li> <li>(d) Increment <math>N</math></li> </ul>
5	Replace $\mathbf{b}_T$ by $\mathbf{b}_R$

---

### 6.2.5 BRA Evaluation

The description of the different block replacement attacks has been kept general on purpose. No hypothesis has been made on the data to be processed to offer a generic framework. This attack can consequently be applied to different types of multimedia content. Previous work from Microsoft Research has focused on audio data [102, 103, 152]. The previously presented TFAR attack in video [49, 51] can also be regarded as some sort of block replacement attack which exploits the temporal redundancy in successive video frames. In this section, image documents will be considered as an extension of earlier work [158]. The next subsections introduce the enforced watermarking scheme during the experiments as well as two basic signal processing operations which will be used as references. Finally, the efficiency of the different proposed algorithms is surveyed in the last subsection.

#### Watermarking Scheme

A basic additive spread spectrum watermark has been considered during the experiments. A secret key  $K$  is used as a seed to generate a pseudo-random

watermark pattern  $\mathbf{w}(K)$ , whose samples have zero mean and unit variance. This watermark is then scaled by an embedding strength  $\alpha$  and added in the spatial domain to the luminance component  $\mathbf{i}_o$  of the original image as follows:

$$\mathbf{i}_w = \mathbf{i}_o + \alpha \mathbf{w}(K) \quad \mathbf{w}(K) \sim \mathcal{N}(0, 1) \quad (6.17)$$

where  $\mathbf{i}_w$  is the resulting watermarked luminance component. Perceptual shaping can be introduced to improve the invisibility of the watermark by making for example the embedding strength  $\alpha$  dependent of the local content of the host image. In practice a global embedding strength equal to 3 has been used which results in a distortion of 9 in terms of MSE, or 38 dB in terms of Peak Signal to Noise Ratio (PSNR).

On the other side, when an image is presented to the detector for verification, the pseudo-random watermark  $\mathbf{w}(K)$  is re-generated using the shared secret key  $K$ . Then, the detector computes a simple linear correlation as follows:

$$\rho(\mathbf{i}, K) = \mathbf{i} \cdot \mathbf{w}(K) = (\mathbf{i}_o + \epsilon \alpha \mathbf{w}(K)) \cdot \mathbf{w}(K) \approx \epsilon \alpha \quad (6.18)$$

where  $\epsilon$  is equal to 1 or 0 depending if the tested luminance component  $\mathbf{i}$  is watermarked or not. If host interference ( $\mathbf{i}_o \cdot \mathbf{w}(K)$ ) is neglected, the correlation score should be equal to  $\alpha$  when the watermark  $\mathbf{w}(K)$  is present in the tested image, while it should be almost equal to zero if  $\mathbf{w}(K)$  has not been embedded. In practice, host interference can be cancelled in a preprocessing step [33] during embedding to enhance the detection statistics. Finally, the correlation score is compared to a threshold  $\tau_{\text{detect}}$  to assert whether or not the watermark  $\mathbf{w}(K)$  has been embedded. This threshold can for example be set to  $\alpha/2$  to have equal false positive and false negative probabilities.

### Reference Attacks

For comparison, the impact of two reference attacks will also be reported. Since watermarking is done in the luminance component of the images, attacks will also be performed only on the luminance component. First, linear filtering and in particular Gaussian filtering has been considered. The filters are computed as follows:

$$G_\sigma[x, y] = \frac{g_\sigma[x, y]}{\sum_{x, y} g_\sigma[x, y]} \quad \text{with} \quad g_\sigma[x, y] = e^{-\frac{x^2 + y^2}{2\sigma^2}} \quad (6.19)$$

where  $\sigma$  is the width of the Gaussian filter. The range of  $x, y$  is limited so that all large values of  $G_\sigma[x, y]$  are included. The filtered image is then obtained by convolving the image with the computed filter. The larger the filter width is, the more distorted is the filtered image. The second reference attack is lossy compression and especially JPEG compression [90]. This standard specifies the quantization values for DCT coefficients by multiplying a quantization matrix

(Table 6.5) by a global quantization level  $Q$ , which is related to a user specified *quality factor*  $QF$  in the range of 0 to 100:

$$Q = \begin{cases} 50/QF & \text{if } QF < 50 \\ 2 - 0.02 QF & \text{if } QF \geq 50 \end{cases} \quad (6.20)$$

For example, if  $QF = 25$ , the global quantization level is equal to 2 and the DC term is quantized with a quantization level of  $q = 32$ . In JPEG, loss of information only occurs during quantization of DCT coefficients. As a result, it is sufficient to perform this quantization to simulate the effects of JPEG compression. The following operation is performed to obtain the quantized value  $\bar{x}$  of a DCT coefficient  $x$

$$\bar{x} = q \left\lfloor \frac{x}{q} + 0.5 \right\rfloor \quad (6.21)$$

where  $q$  is the quantization value computed as described above. The lower the JPEG quality factor is, the more distorted is the compressed image.


**Table 6.5:** Luminance quantization matrix used in JPEG.

16	11	10	16	24	40	51	61
12	12	14	19	26	58	60	55
14	13	16	24	40	57	69	56
14	17	22	29	51	87	80	62
18	22	37	56	68	109	103	77
24	35	55	64	81	104	113	92
48	64	78	87	103	121	120	101
72	92	95	98	112	100	103	99

## Performances

A database of 500 images of size  $512 \times 512$  has been considered for experiments. It contains snapshots, synthetic images, drawings and cartoons. All the images are first watermarked using the algorithm described in Subsection 6.2.5. Since the detection is based on the computation of a correlation score, distortion vs. correlation curves can be plotted to study the impact of a given attack. To this end, each watermarked images has been submitted to 4 alternative attacks (Gaussian blurring, DCT quantization, adaptive number of blocks combination, and block projection on PCA-defined subspace) with predefined attacking parameter settings. For the reference attacks, the width  $\sigma$  of the filter and the quality factor  $QF$  can be varied. On the other hand, for both block replacement attacks, the thresholds  $\tau_{\text{low}}$  and  $\tau_{\text{high}}$  have to be set. However they can be set equal so that the resulting parameter  $\tau_{\text{target}} = \tau_{\text{low}} = \tau_{\text{high}}$  basically sets a target distortion in

terms of MSE that the attack should introduce. Furthermore,  $8 \times 8$  blocks have been used with a 4-pixels overlapping. Using overlapping blocks is indeed really important to avoid annoying blocking artifacts with high values for  $\tau_{\text{target}}$ . At this point, for each image in the database, a distortion vs. correlation curve can be drawn for each one of the 4 surveyed attacks. The different curves associated with a given attack are then *averaged* to obtain a single curve per attack which depicts the statistical behavior of the image database for a particular attack. The obtained 4 curves are reported in Figure 6.6.



**Figure 6.6:** Correlation score vs. distortion curves for the different surveyed attacks.

The goal of the attacker is to decrease the correlation score computed by the detector while maintaining the image quality. As a result, if a curve is below another one in a distortion vs. correlation plot, it means that the first attack has a stronger impact on the watermark than the second one. Looking at Figure 6.6, it is obvious that both proposed block replacement attacks outperform Gaussian blurring and JPEG compression. In particular, the correlation score drops below the detection threshold  $\tau_{\text{detect}} = 1.5$  around 40 dB with block replacement attacks while it is necessary to introduce a distortion around 36 dB to obtain the same result with the reference attacks. Furthermore, assuming that the parameters of the attacks are set so that the introduced distortion is similar to the one due to the embedding process (38 dB), block replacement attacks trap the detector while watermarks submitted to reference attacks can still be detected. In other words, from an attacker perspective, the introduced block replacement strategy allows to improve the trade-off distortion vs. correlation in comparison with other standard reference attacks. Both block replacement attacks exhibit roughly

the same performance. However, block projection on a PCA-defined subspace requires fewer computations than adaptive number of blocks combination.

## 6.3 Discussion

Both presented attacks (TFAR and BRA) follow the same strategy: consider the signal spatio-temporal redundancy to replace each signal block, should it be a full frame or a small  $8 \times 8$  pixels block, with another or a combination of other ones. In other words, the attacker simply exploits the fact that watermarking algorithms usually do not consider the signal self-similarities during embedding. As a result, it is possible to build some sets of similar blocks which on the other hand are not assumed to carry similar watermark samples. This is a weak link of most watermarking schemes today and a witty attacker is likely to exploit it to defeat the protection system. Of course, this brings up an interesting question: which countermeasures can be introduced by technology providers to disable, or at least decrease the impact, of such attacks? Intuitively, if similar signal blocks carry similar watermarks, the presented block replacement strategy is likely to be ineffective. That is to say that the introduced watermark has to be coherent with the self-similarities of the host signal. This can be seen as an intermediary specification between the security requirements for steganography - the embedded watermark should be statistically invisible [165] so that an attacker cannot even detect the presence of the hidden watermark - and the absence of any one for non-secure applications such as data hiding. Unfortunately this intuitive statement does not point to straightforward constructive ideas on how to obtain such coherent watermarks in practice. Therefore, the next part of the thesis will present two complementary video watermarking strategies to address temporal redundancy on one side and spatial redundancy on the other side.



**Part III**

**Signal Coherent Watermarking**





---

## Motion Compensated Watermarking

---

The results presented in the previous part basically recommend to watermark correlated video frames with the same watermark on one hand, and uncorrelated video frames with uncorrelated watermarks on the other one. These rules have subsequently been extended to give the following well-known fundamental embedding principle. *Watermarks embedded in distinct frames should be as correlated as the host video frames*, as written below:

$$\forall(t, t') \quad \rho(\mathbf{w}_t, \mathbf{w}_{t'}) \approx \rho(\mathbf{f}_t, \mathbf{f}_{t'}), \quad (7.1)$$

where  $\rho(\cdot)$  is a given correlation score, e.g. the correlation coefficient. Alternative approaches have been proposed to meet this specification e.g. the embedded watermark can be made frame-dependent [82], a frame-dependent binary string can be exploited to generate a watermark pattern which degrades gracefully with an increased number of bit errors [70, 41], the watermark can be embedded in some frame-dependent positions [174].

None of these solutions is however perfect and they are even likely not to be stable enough to be deployed in a video framework [119]. Furthermore, this *same correlation* specification may not be enough to ensure security. The correlation score between a video frame and a shifted version of it may be very low. Nonetheless, the embedded watermarks should not be completely uncorrelated. In fact, the watermark embedded in the shifted frame should also be a shifted version of the watermark embedded in the reference frame. Otherwise, Temporal Frame


Averaging after Registration (TFAR) introduced in Section 6.1 is likely to remove the embedded watermark. Skeptical people might argue that TFAR is too intensive in terms of computations to be realistic. However, video mosaics or *sprite panoramas* are expected to be exploited for efficient background compression in the upcoming video standard MPEG-4 [105]. Such video coding algorithms will have a similar impact on a potentially embedded watermark. As a result, this issue has to be addressed. In this perspective, a watermarking strategy based on motion compensation will be presented in details in this chapter.

## 7.1 Watermarking Exploiting Video Mosaicing (SS-Reg)

For a given scene, backgrounds of video frames can be considered as several 2D projections of the same 3D movie set. The weakness of SS and SS-1 embedding strategies against TFAR is due to the fact that camera motion is not considered at all. These watermarking systems are completely *blind*. As a result, a given 3D point which is projected in different locations in different video frames is associated with uncorrelated watermark samples. Thus, averaging registered video frames succeeds in confusing the watermark detector. The goal is consequently to inform the embedder about camera motion and to find an embedding strategy which forces each 3D point to carry the same watermark sample whenever it is visible in the video scene. In other words, the basic idea is to simulate a utopian world where the movie set would already be watermarked. In this perspective, video mosaicing will be considered in the remaining of this section to design a new watermarking system.

### 7.1.1 Watermark Embedding

Video mosaicing consists in aligning all the frames of a video sequence to a fixed coordinate system [85]. The resulting mosaic image provides a snapshot view of the video sequence i.e. an estimation of the background of the scene if the moving objects have been removed. A straightforward and naive approach would consist in embedding a digital watermark in the mosaic representation of the considered video scene. Next, the resulting watermark mosaic would be used as the background of the video frames. However, such a process requires double interpolation (frame  $\rightarrow$  mosaic  $\rightarrow$  frame) which is likely to reduce the performances of the detector. Therefore, an alternative but somewhat equivalent approach is depicted in Figure 7.1. First of all, warping parameters are computed for each video frame with respect to the considered motion model. For instance, if the motion model defined in Equation (6.2) is exploited, the warping parameters  $\theta$ ,  $z$ ,  $(x_o, y_o)$  and  $(t_x, t_y)$  are computed for each video frame. Hence, each frame  $\mathbf{f}_t$  is



**Figure 7.1:** Embedding procedure for camera motion coherent watermarking (SS-Reg): The part of the watermark pattern which is associated with the current video frame is retrieved and registered back. Next, it is embedded in the background portion of the video frame.

associated with a set of warping parameters i.e. the frame background  $\mathbf{b}_t$  is associated with a portion  $\mathbf{b}_m^{(t)}$  of the video mosaic. Next, a key-dependent watermark  $\mathbf{w}_m$  is generated which has the same dimensions as the mosaic representation of the video shot. Now, using the same warping parameters as the ones used for building the mosaic, a portion  $\mathbf{w}_m^{(t)}$  of this watermark can be associated to each video frame  $\mathbf{f}_t$ . Finally, the resulting watermark portion only has to be registered back to obtain the watermark signal  $\mathbf{w}_t$  to be embedded in the video frame. The overall embedding process can consequently be written as follows:

$$\check{\mathbf{f}}_t = \mathbf{f}_t + \alpha \bar{\mathbf{m}}_t \otimes \mathbf{w}_t, \quad \mathbf{w} \sim \mathcal{N}(0, 1). \quad (7.2)$$

where  $\check{\mathbf{f}}_t$  is the  $t$ -th watermarked video frame,  $\alpha$  the embedding strength,  $\otimes$  the pixel-wise multiplication operator and  $\bar{\cdot}$  the binary negation operator. Additionally, perceptual shaping can be introduced to make the embedded watermark less noticeable. This novel embedding strategy will be referred to as the *SS-Reg* strategy. Similarly to TFAR, a binary mask  $\mathbf{m}_t$  has been defined to isolate the background  $\mathbf{b}_t$  from moving objects  $\mathbf{o}_t$  using Equation (6.1). It should be noted that according to Equation (7.2) moving video objects are left unprotected because of the pixel-wise multiplication by the binary mask  $\bar{\mathbf{m}}_t$ . This operation can be removed so that the watermark signal  $\mathbf{w}_t$  spreads over the whole video frame. However, this would contradict the underlying philosophy of this embedding strategy: *a 3D point carries the same watermark sample all along the video scene*. As a result, alternative mechanisms have to be deployed to protect these objects if needed. Previous works have watermarked MPEG-4 video objects ac-

ording to their main directions [11], their animation parameters [78] or their texture [72].

### 7.1.2 Watermark Detection

On the detector side, the procedure is very similar to the embedding one. In a first step, warping parameters are computed for each frames of the video scene to be checked and the watermark  $\mathbf{w}_m$  is generated using the shared secret key. Next, the detector only checks whether the portion  $\mathbf{w}_t$  associated with each incoming frame  $\tilde{\mathbf{f}}_t$  has been effectively embedded in the background or not. This can be done using the following correlation score:

$$\rho(\tilde{\mathbf{f}}_t, \mathbf{w}_m) = \frac{\tilde{\mathbf{f}}_t \cdot \mathbf{w}_t^{(t)}}{m_t} \approx \frac{\epsilon\alpha}{m_t} (\tilde{\mathbf{m}}_t \otimes \mathbf{w}_t^{(t)}) \cdot \mathbf{w}_t^{(t)} = \epsilon\alpha, \quad (7.3)$$

where  $\cdot$  denotes the linear correlation,  $\epsilon$  equals 0 or 1 depending whether the video is watermarked or not and  $m_t$  is the percentage of pixels contained in the background of frame  $\tilde{\mathbf{f}}_t$ . A preprocessing step [33] can be added to remove host interference in Equation (7.3) and thus improve the detection statistics. The proposed correlation score should then be equal to  $\alpha$  if a watermark is present in the video frame, while it should be almost equal to zero if no watermark has been inserted. As a result, the computed score is compared to a threshold  $\tau_{\text{detect}}$  in order to assert the presence or absence of the watermark. The value given to this detection threshold determines the false positive and false negative probabilities and the value  $\alpha/2$  can be chosen for equal false positive and false negative probabilities.

In practice, successive video frames could also be exploited to assert whether a watermark is embedded in a video sequence or not. In this perspective, the different correlation scores are simply accumulated as follows:

$$P_w(\tilde{\mathbf{f}}_t, \mathbf{w}_m) = \frac{1}{2w+1} \sum_{\delta=-w}^w \rho(\tilde{\mathbf{f}}_{t+\delta}, \mathbf{w}_m) \approx \epsilon\alpha. \quad (7.4)$$

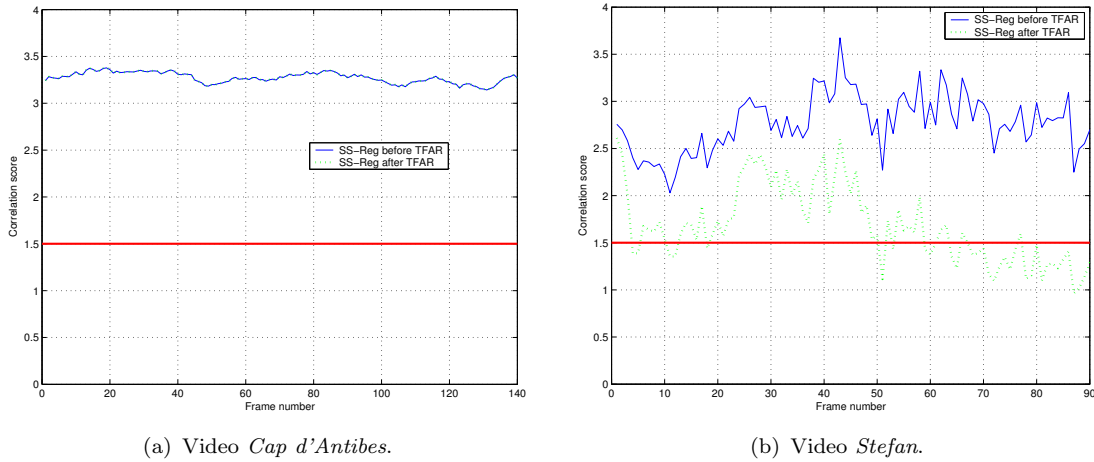
It should be noted that, when the temporal window covers the whole video sequence, such a detection procedure is equivalent to build the video mosaic of the scene and to compute the linear correlation with the watermark pattern  $\mathbf{w}_m$ . Considering many frames is commonly used [93] to enhance detection statistics. Indeed, some video processing, such as linear filtering, noise addition or lossy compression, are likely to introduce an interfering term in Equation (7.3). As a result, the correlation score is equal to  $\epsilon\alpha + n$ , where  $n$  can be considered as normally distributed with zero mean and unknown variance  $\sigma$ . This has a direct impact on the false positive and false negative probabilities. Accumulating successive scores as in Equation (7.4) allows to reduce the effect of the interfering term  $n$  since it divides its variance by a factor  $\sqrt{2w+1}$ .

## 7.2 System Analysis

The novelty of the proposed embedding strategy lies in the fact that camera motion is compensated before embedding the watermark. To the best knowledge of the author, such an approach is completely new even if some similarities can be exhibited with other works. For instance with the SLIDE algorithm, small watermark patches are embedded at some image dependent anchor locations [175]. One can expect that these anchor points remain the same from a 3D point of view all along a video sequence and thus be coherent with camera motion. However, tracking of anchor point has not been explicitly addressed in that paper. Another work of interest exploited block tracking to always embed the watermark signal in the same blocks in successive video frames [114]. Nevertheless, only a few blocks are considered for embedding which drastically reduces the embedding capacity. To validate the proposed watermarking strategy, the remainder of this section will survey different outcomes regarding some important properties in digital watermarking.

### 7.2.1 Enhanced Security

In Section 7.1, the very first motivation for considering motion compensation before embedding was to enhance performances with respect to TFAR. As a result, resilience against TFAR has to be verified to demonstrate the superiority of the new embedding strategy SS-Reg in comparison with SS or SS-1 strategies. To this end, two different video sequences have been considered for the experiments: on one side the synthetic sequence *Cap d'Antibes* and on the other side the real video scene *Stefan*. Both sequences have been previously considered in Subsection 6.1.2 to evaluate the performances of SS and SS-1 strategies with respect to TFAR. The basic difference between the two sequence is that *Cap d'Antibes* has been artificially created from a panoramic view. Therefore, displacements of the camera are perfectly known and there is no artifact due to lens imperfection for instance. On the other hand, registration parameters of the sequence *Stefan* have to be estimated and frame registration is not perfect. Both videos have been watermarked using the SS-Reg embedding strategy. The embedding strength has been set to 3 so that the embedding process introduces a distortion around 38 dB in terms of Peak Signal to Noise Ratio (PSNR). Next, the watermarked sequences have been submitted to TFAR with a temporal window size equal to 3 i.e. the background of each video frame is approximated using the previous and the next video frames. Finally, the detection score defined in Equation (7.3) is computed before and after TFAR and the results are plotted in Figure 7.2. This figure has to be compared with Figure 6.2 which depicts the performances of the SS and SS-1 embedding strategies with respect to BRA. At first sight, it is immediate to notice that whereas the detection score drops down to zero after TFAR with



**Figure 7.2:** Impact of TFAR on the SS-Reg system.

SS and SS-1 systems, it remains above the detection threshold  $\tau_{\text{detect}}$  when the SS-Reg embedding strategy is enforced. Nevertheless, this result is contrasted depending on the considered video sequence. With the video *Cap d'Antibes*, the curves are almost superimposed i.e. TFAR has no impact on the embedded watermark. This is due to the fact that the displacements between adjacent frames are perfectly known. In fact, it is a 2 pixels per frame horizontal translation from right to left. As a result, there is no interpolation when computing the frame to mosaic mapping. Alternatively, the situation is completely different with the video *Stefan*. In this case, camera motion is *estimated*. There are some approximations due to the selected motion model and several pixels corresponding to the same physical 3D point do not strictly overlap in the mosaic. Furthermore, pixel displacements are not necessarily integers and interpolations have to be computed. Both points have an impact on the performances of the proposed system. In fact, it seems to be more critical at the end of the video sequence. After examination, this seems to be mainly related with the variations of the zoom factor  $z$  in the motion model defined in Equation (6.2). This point will be further discussed in Subsection 7.3.1.

### 7.2.2 Video Capacity

The presented motion compensated video watermarking scheme has a zero bit capacity. It only gives an answer to the question: *is there a portion of the watermark  $\mathbf{w}_m$  in each video frame?* However, it should be possible to modify the embedding strategy so that some payload can be hidden in a video scene. In comparison with still images, a video sequence provides a larger number of digital samples which can be exploited to carry some hidden information. A common mistake consists then in asserting that a greater payload can be embedded. Such a claim is true if

there is no security requirement. For example, digital watermarking can be used for data hiding i.e. to embed some additional useful information in an invisible way. However, if the targeted application includes strong security specifications (copy control, fingerprinting), advanced hostile attacks such as TFAR are likely to occur and have to be addressed. As a result, the embedding strategy has to ensure that a given 3D point of the movie set always carries the same watermark sample in a video sequence. It is somewhat related with the notion of statistical invisibility introduced in previous work [176]. The proposed SS-Reg embedding strategy gives then some intuitive insight on how many bits can be *securely* embedded in a video sequence. Looking at Figure 7.1, the embedding procedure can be regarded as inserting a watermark into the mosaic representation of the video shot and subsequently exploiting this watermarked mosaic to replace the background in each video frame. In other words, the capacity is related with the dimensions of the mosaic i.e. with camera motion. If the camera is static, the mosaic image has the same dimensions as a video frame and a moderate payload can be embedded. On the other hand, as soon as the camera moves, new areas of the movie set are revealed and they can be used to hide a larger payload.

### 7.2.3 Watermark Visibility

As previously discussed in Section 4.3, evaluating the impact of distorting a signal as perceived by a human user is a great challenge. The amount and perceptibility of distortions, such as those introduced by lossy compression or digital watermarking, are indeed tightly related to the actual signal content. This has motivated the modeling of the human perception system to design efficient metrics. For example, when considering an image, it is now admitted that a low-frequency watermark is more visible than a high-frequency one or that a watermark is more noticeable in a flat area than in a texture one. The knowledge of such a behavior can then be exploited to perform efficient perceptual shaping. In the context of video, the Video Quality Experts Group (VQEG) [188] was formed in 1997 to devise objective methods for predicting video image quality. In 1999, they stated that no objective measurement system at test was able to replace subjective testing and that no objective model outperforms the others in all cases. This explains while the Peak Signal to Noise Ratio (PSNR) is still the most often used metric today to evaluate the visibility of a video watermark. However, from a subjective point of view, previous works [131, 197] have isolated two kinds of impairments which appear in video, when the embedding strength is increased, but not in still frames:

1. *Temporal flicker*: Embedding uncorrelated watermarks in successive video frames (SS strategy) usually results in annoying twinkle or flicker artifacts similar to the existing ones in video compression,

2. *Stationary pattern*: Embedding the same watermark pattern in all the video frames (SS-1 strategy) is visually disturbing since it gives the feeling that the scene has been filmed with a camera having a dirty lens when it pans across the movie set.

With the proposed motion compensated embedding strategy, different watermarks are still embedded in successive video frames. However, these differences are coherent with the camera motion and the user is no longer annoyed by flickering. In fact, the user has the feeling that the noise was already present in the filmed movie set and find it more *natural*. On the other hand, the proposed embedding strategy introduces a new kind of artifacts. All the embedded watermarks  $\mathbf{w}_t$  originate from the same watermark pattern  $\mathbf{w}$ . Nevertheless, they have been obtained using different warping parameters, and in particular different zoom factor  $z$ . As a result, the embedded watermarks have not the same frequency content: if the camera zooms in, the watermark slides towards low frequencies and thus becomes more visible. This issue will be briefly addressed in Subsection 7.3.1.

## 7.3 Further Studies


The introduced motion-compensated watermarking strategy has exhibited some interesting properties. Nevertheless, a few points can be improved to enhance the performances of the SS-Reg system. Two major issues will be addressed in the next subsections which basically appeared because of the choice of video mosaicing to produce motion compensated watermarks. First, Subsection 7.3.1 will explore in more details why the detection performances are sensible to interpolation and changes of zoom factor  $z$ . In particular, methods used for texture mapping in 3D computer graphics will be surveyed to address this issue. In the next subsection, a basic framework for local resynchronization is introduced. Indeed, nothing ensures that registration parameters estimated at the embedder and the detector will be the same since the input video sequences for video mosaicing at least differ by the embedded watermark. As a result, small misalignments are likely to occur and have to be compensated.

### 7.3.1 Detection Performances Improvement

When looking closely at the procedure depicted in Figure 7.1, the embedding process basically comes down to resampling the watermark pattern  $\mathbf{w}_m$ . For each pixel of the image, the associated position in the mosaic is computed and the corresponding sample value is retrieved. However, nothing ensures that this position will correspond to an existing watermark sample i.e. nothing ensures that the coordinates of the position within the mosaic will be integers. As a




result, it is necessary to perform interpolation. In the presented experiments, bilinear interpolation has been used to obtain the unknown watermark value at a given location according to its 4 nearest neighbors. Now the critical issue is that video mosaicing does not perform perfect registration. It only minimizes some error with respect to some motion model. When TFAR is performed, each pixel is reconstructed with the ones in the neighbor frames which are associated with the same position in the mosaic. Nevertheless, due to the imperfections of the registration process, those pixels may be not *exactly* the same, resulting in slightly different interpolated watermark sample values. Moreover this effect is emphasized when the zoom factor  $z$  increases. In this case, it can even happen that the pixel location in the mosaic moves from one 4-pixels neighborhood to another close one i.e. the watermark sample is interpolated using other samples.



**Figure 7.3:** Impact of the chip rate on detection performances.

To verify this statement, experiments have been performed with a watermark pattern  $\mathbf{w}_m$  supersampled by a chip rate  $c$ . In other words, the pattern is divided in blocks  $c \times c$  which contains the same watermark value. The video sequence *Stefan* has been watermarked with the SS-Reg system using 4 different chip-rate values: 1, 2, 3 and 4. Next the watermarked video sequences have been submitted to TFAR. Finally the ratio between the detection score computed after and before TFAR has been computed and the results have been plotted in Figure 7.3. As expected, the greater the chip rate is, the better the embedded watermark resists to TFAR. This is due to the fact that even if the considered 4-pixels neighborhood changes because of imperfect registration, the values of the samples considered for interpolation have fewer chances to be modified. However, this improvement in terms of detection performances has to be balanced with imperceptibility considerations. The greater the chip rate is, the more visible is

the embedded watermark because of its lower frequency content. An attentive watcher may even be able to spot the watermark blocks. Furthermore, the size of this blocking artifact is modified with respect to the zoom factor. The greater the zoom factor is, the bigger appear the watermark blocks even if, from a subjective point view, one may prefer to always have a watermark with the same frequency content.



**Figure 7.4:** Alternative strategies for watermark resampling.

Up to now, a pixel has been considered as a point. Each pixel in the video frame is associated with some coordinates in the mosaic which are then used to compute the interpolated watermark value. But, in reality, a discrete pixel represents an area rather than a point. The pixel value output by a digital camera is related with the number of photons which have hit a sensor covering a small area. As a result, as depicted in Figure 7.4, instead of interpolating the watermark value at a given location in the mosaic, one should integrate upon the projected area corresponding to the source pixel to properly reflect the information content being mapped onto the output pixel. This resampling issue is common in 3D computer graphics to perform antialiasing [195]. The projected area is referred to as *preimage* and the integral over this area is usually computed by combining supersampling and low-pass filtering. This approach has the advantage to produce a watermark which keeps the same frequency content. Furthermore, detection performances are likely to survive as long as the preimages corresponding to the same pixel in different video frames mainly overlap.

### 7.3.2 Local Resynchronization


Another concern with the SS-Reg system, which has not been addressed yet, is the stability of the registration parameters. It has been assumed for the moment that they have remained the same when they are computed either at the embedder side or at the detector side. However, in real life, some distortions are likely to occur in between. In particular, the embedded watermark is likely to alter a bit the computed registration parameters. Therefore, small misalignments

may happen and a resynchronization mechanism has to be introduced to cope with local distortions. As mentioned previously, no specific research has been conducted to design new motion estimation/compensation tools. Previous work by Professor Henri Nicolas has simply been reused i.e. registration parameters are not computed each time. Nevertheless, the next subsections will describe a possible framework to compensate for small distortions

### Basic Framework

Several approaches have already been proposed to cope with geometrical distortions and four major trends can be isolated. A possible solution is to make the original non-watermarked document available on the detector side. In this case, the undergone transformation can be estimated and inverted prior to watermark detection. However, such non-blind approaches require storing all the original documents which rapidly becomes intractable in practice when the size of the considered database increases. A second approach is to perform an exhaustive search. In this brute force perspective, each potential geometric transformation that might have been applied is inverted and the detector checks whether a watermark can be found or not. Nevertheless, computational cost of such an approach rapidly grows when the set of hypothetical geometrical transformation is enlarged. Moreover, performing successive watermark detection is likely to significantly increase the false positive probability [118]. Another common resynchronization technique consists in inserting an additional watermark which is often referred to as template, registration pattern or pilot watermark [108]. This template is then exploited as a reference to detect and compensate for geometric distortions. The main drawback of this technique is that it relies on the presence of local peaks e.g. in the frequency domain which can be easily detected and removed by a malicious party [190]. Finally another solution consists in embedding the watermark in a subspace which is immune to geometric distortions [2, 164]. In this case, the immunity against geometric distortions induces a reduction of the watermarking space dimension i.e. the embedding capacity is reduced.

The approach relying on pilot watermarks has been shown to give good performances. This has consequently motivated further research to get rid of the security weakness against estimation attacks [18, 58]. Once again, the idea is to have two different watermarks: one pilot watermark used for resynchronization and one payload watermark to encode the message to be hidden. However the pilot watermark no longer requires to exhibit specific patterns such as peaks in some domain. Both watermarks are either superimposed or finely interlaced. The basic assumption is that both watermarks will experience the same geometric distortion. Thus, once the detector has estimated the experienced distortion using the pilot watermark, it can exploit this knowledge to properly extract the embedded payload. Indeed, the detector knows that if the considered image  $\mathbf{i}$  has



**Figure 7.5:** Estimated optical flow with BM and EGM.

been watermarked, then it should contain the pilot watermark  $\mathbf{w}$  possibly locally distorted. As a result, the detector tries to estimate the most likely optical flow  $\Delta$  that the document has experienced. To do so, a simple Block Matching (BM) approach can be enforced. The image is divided in blocks  $\mathbf{i}_b$  and for each block the algorithm searches in the pilot watermark for the block in the neighborhood which has the highest matching score e.g. the highest correlation score. Therefore, the whole process can be regarded as minimizing the following cost function:

$$C(\mathbf{i}, \mathbf{w}, \Delta) = \sum_b C_{\text{match}}(\mathbf{i}_b, \mathbf{w}, \delta_b) = C_{\text{match}}(\mathbf{i}, \mathbf{w}, \Delta) \quad (7.5)$$

where  $b$  is some block index,  $\delta_b$  the displacement associated with block  $b$  and  $C_{\text{match}}(\cdot)$  some matching cost function. This function has a low (resp. high) value when the watermark block defined by  $\delta_b$  is very likely (resp. unlikely) to be embedded in the image block  $\mathbf{i}_b$ . Once the optical flow has been estimated, it is exploited to extract the message encoded by the payload watermark. Furthermore, a detection score is computed to assert whether a watermark is effectively present or not.

### Resynchronization Enhancement


The first important shortcoming of the presented resynchronization process is that it operates blindly in a best match fashion. Minimizing the cost function defined by Equation (7.5) basically comes down to minimizing each term in the sum. There is no constraint at all between the displacements  $\delta_b$  of neighbor

blocks and this is of course suboptimal. The visual quality of an image distorted by a geometric distortion is indeed determined by its homogeneity. The less homogenous the distortion is, the worse the visual quality is [42]. In other words, neighbor displacements are likely to be correlated to a certain extent so that the resulting overall distortion remains tolerable. However this property is not exploited during the resynchronization process based on BM. As a result, nothing ensures that the estimated optical flow will even belong to the set of possible transformations. For instance, a typical optical flow obtained with BM using blocks of size  $64 \times 64$  is depicted in Figure 7.5. The estimated flow has some non coherent displacements and corresponds to a geometrical transformation which cannot be applied in practice for visibility reasons. Thus, Elastic Graph Matching (EGM) has been considered to add a smoothness constraint in the BM framework [111]. This is done by introducing a rigidity parameter to prevent displacements of neighbor blocks from being incoherent. The estimation of the optical flow  $\Delta$  is then equivalent to minimizing the following cost function:

$$C(\mathbf{i}, \mathbf{w}, \Delta) = C_{\text{match}}(\mathbf{i}, \mathbf{w}, \Delta) + \lambda C_{\text{smooth}}(\Delta) \quad (7.6)$$

where the parameter  $\lambda$  controls the rigidity of the estimated optical flow. The cost function  $C_{\text{smooth}}(\cdot)$  measures in some sense the smoothness of the estimated optical flow. For instance, the sum of the squared distances between neighbor displacements  $\delta_b$  can be used considering only the four nearest neighbors. This smoothing cost function interferes with the BM process to enable blocks that are not the *best* matching ones to be still considered in case they are coherent with the current estimation of the optical flow  $\Delta$ . The optical flow is updated in an iterative fashion and Figure 7.5 illustrates the advantage of the EGM framework over BM. The few incoherent displacements obtained with BM are corrected with EGM.

The second shortcoming of the presented resynchronization process is that the block size has a great influence on the performances. On one hand, small blocks are likely not to contain enough watermark samples to enable a correct registration and compensate for local geometric distortions. On the other hand, considering large blocks prevents from estimating finely the geometric distortions. Therefore, one would like to reduce the size of the block to be able to compensate for more complex geometric distortions. However, Figure 7.6 depicts the pitfall of this approach. The estimated optical flow with blocks of size  $16 \times 16$  is almost random and payload extraction is no longer possible. This is a common shortcoming of methods relying on the minimization of some cost function. The iterative process gets trapped in a local minimum and the global minimum, which is usually the expected solution, can be missed. To circumvent this drawback, a multi-scales approach can be superimposed so that dense optical flows can be obtained. The basic idea is to start the resynchronization process with large blocks



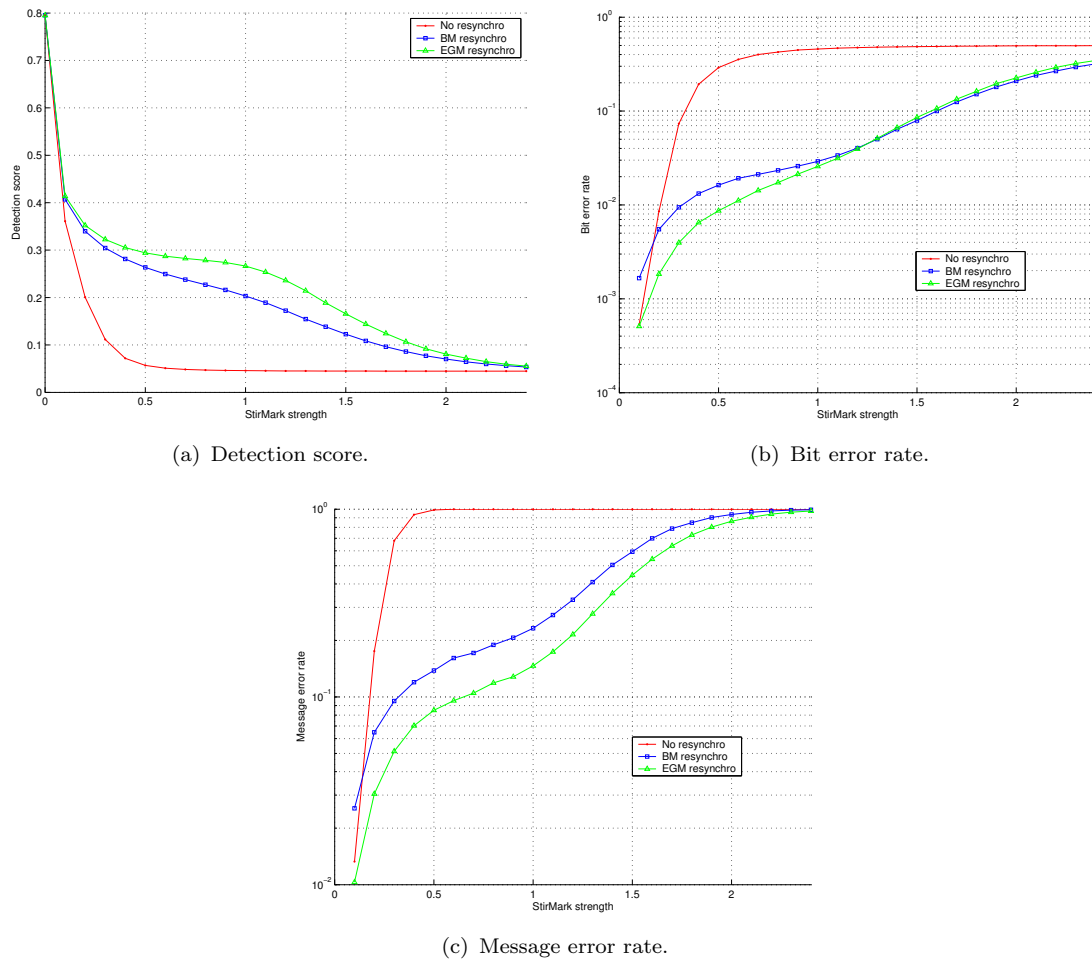
**Figure 7.6:** Influence of the multi-scales framework on the optical flow.

and then to successively consider smaller blocks. Large blocks enable to find a relevant initial estimation of the optical flow since many watermark samples are available for each block. Then, the block size can be slowly decreased to refine the optical flow by permitting more geometric distortions. It should be noted that the rigidity parameter  $\lambda$  has to be carefully updated. The same value cannot be used for  $64 \times 64$  blocks and  $16 \times 16$  blocks. Indeed, in the latter case, neighbor blocks are nearer and thus neighbor displacements should be more similar i.e. the rigidity constraint should be stronger. The parameter  $\lambda$  should consequently be set so that it grows larger when the block size decreases. Figure 7.6 shows that with such considerations, it is possible to obtain a dense and smooth optical flow, which should enable efficient payload extraction.

### Performances Evaluation

A database of 500 images of size  $512 \times 512$  has been considered for experiments. It contains snapshots, synthetic images, drawings and cartoons. All the images are first watermarked with a 64 bits message using the proprietary algorithm Eurémark described in Appendix B. Next, those watermarked images are submitted to the StirMark attack. StirMark is recognized to be a key benchmarking tools when local distortions are considered [151, 109]. It basically simulates a resampling process i.e. it introduces the same kind of distortions into an image as printing it on a high quality printer and then scanning it again with a high quality scanner. To this end, it combines a global bilinear transform, a global bending and random high frequency displacements. The attack is performed

with an increasing strength  $\alpha > 0$ . On the detector side, three resynchronization methods are surveyed: no resynchronization, BM based resynchronization and EGM resynchronization. For each method and for each attacked image, the detection score, the Bit Error Rate (BER) and the Message Error Rate (MER) are computed. This experiment is performed 25 times with alternative random embedding keys. It results in  $500 \times 25 = 12500$  curves which indicate the evolution of the detection score (or BER/MER) vs. the StirMark strength for a given image, a given embedding key and a given resynchronization method. All those curves are averaged and then reported in Figure 7.7.



**Figure 7.7:** Impact of TFAR on the SS-Reg system.

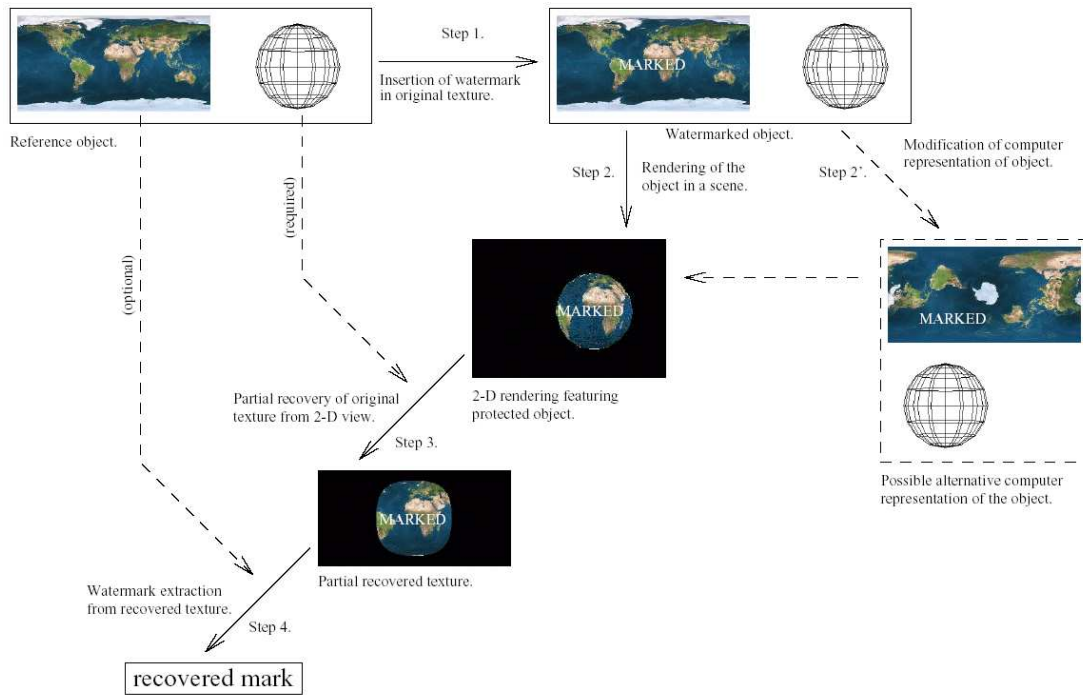
As expected, Eurémark is quickly defeated when no resynchronization is performed on the detector side. The resynchronization process improves significantly the performances of the algorithm. Furthermore, the novel EGM based resynchronization module appears to slightly outperform the previous one based on BM only. Following a common practice [109], the watermarking scheme is con-

sidered to be robust if at least 80% of the watermarks are correctly retrieved i.e. the MER is below 20%. The different schemes are respectively defeated for a StirMark strength equal to 0.2 with no resynchronization, 0.85 with BM resynchronization and 1.2 with EGM resynchronization. The improvement between BM and EGM resynchronization is due to two different aspects in the design of the resynchronization module. First, the rigidity constraint enables to correct incoherent estimated displacements as it has been depicted in Figure 7.5. Second, the multi-scales framework allows to better cope with small local distortions. For large StirMark strength, both schemes are defeated because the resynchronization procedure is limited to the size of the search window. Additionally, it is interesting to note that the curves for BM and EGM resynchronization do not vary regularly: it seems that there is a step somewhere in the middle. This reveals a weakness due to the fact that fractal coding is considered in Eurémark. Because of computational cost, the computation of the fractal cover is block-based. Thus, geometric distortions disturb the alignment of the blocks and the fractal cover is not computed using exactly the same blocks. It could be possible to get round this weakness by considering overlapping blocks during cover computation. But this comes of course with additional computational cost.

## 7.4 Discussion

In this chapter, a new watermark embedding strategy based on camera motion compensation has been presented. The main motivation was to gain some robustness against TFAR in comparison with previous common embedding strategies such as SS and SS-1 systems. In this perspective, the goal is to ensure that a physical 3D point of the filmed movie set always carries the same watermark sample wherever it appears in a frame of the considered video sequence. In other words, the embedding process tries to simulate a situation where a regular camera is filming a world which is already watermarked. The crucial point is then to be able to retrieve the embedded watermark. This can be related with previous work to protect 3D objects through texture mapping [72, 71]. This approach is described in Figure 7.8. Here, the watermark is embedded in the texture information rather than in the geometry of the considered 3D object which is the most commonly enforced embedding strategy. The motivation behind this drastic alternative is to be able to detect the watermark from 2D views of the watermarked object i.e. when 3D geometry information is no longer available. Those two problems - motion compensated watermarking and 3D object watermarking through texture mapping - are tightly connected. If it was possible to obtain a 3D representation of the movie set from a given video sequence, then using the presented watermarking technique on the 3D objects of the scene would produce a perfectly motion-compensated watermark after rendering with the same param-





**Figure 7.8:** 3D object watermarking through texture mapping.

eters. In this Ph.D. thesis, a practical implementation based on video mosaicing has been presented to obtain a motion-compensated watermark. This approach has proven that compensating motion camera has a valuable impact in terms of robustness against TFAR or perceptual invisibility. Nevertheless, video mosaicing is a quite computational demanding method and this is likely to prevent the introduction of such an approach in a commercial application where real-time is requested. As a result, further research needs to be conducted to find alternative ways of producing such motion-compensated watermarks.



---

## Similarities Inheritance

---

As presented in Section 6.2, for each signal block, Block Replacement Attacks (BRA) look for a linear combination of neighboring blocks resulting in a replacement block which is similar enough to the current block so that a substitution does not introduce strong visual artifacts. Since watermarking systems do not perform today anything specific to ensure that the embedded watermark is coherent with the self-similarities of the host signal, most of them are defeated by such attacks. Intuitively, to ensure that a watermark will survive BRA, the embedding process should guarantee that *similar signal blocks carry similar watermarks* or alternatively that *pixels with similar neighborhood carry watermark samples with close values*. In this perspective, assuming that it is possible to characterize the neighborhood in each point with a feature vector, signal coherent watermarking can be achieved if watermark samples are considered as the output of a linear form in this feature space as it is theoretically demonstrated in Section 8.1. A practical implementation of this approach using Gabor features is then described in Section 8.2. Next, in Section 8.3, a relationship with existing multiplicative watermarking schemes in the frequency domain is exhibited. Finally, several experiments are reported in Section 8.4 to investigate the performances of the proposed signal coherent watermarking schemes with respect to BRA.

### 8.1 Linear Watermarking with Neighborhood Characteristics

Let us assume for the moment that it is possible to associate to each pixel position  $\mathbf{p} = (x, y)$  with  $1 \leq x \leq X$  and  $1 \leq y \leq Y$  in the image  $\mathbf{i}$  a feature vector  $\mathbf{f}(\mathbf{i}, \mathbf{p})$

which characterizes *in some sense* the neighborhood of the image around this specific position. Thus, this function can be defined as follows:

$$\begin{aligned} \mathbf{f} : \mathcal{I} \times \mathcal{P} &\rightarrow \mathcal{F} \\ (\mathbf{i}, \mathbf{p}) &\mapsto \mathbf{f}(\mathbf{i}, \mathbf{p}) \end{aligned} \quad (8.1)$$

where  $\mathcal{I}$  is the image space,  $\mathcal{P} = [1 \dots X] \times [1 \dots Y]$  the position space and  $\mathcal{F}$  the feature space. From a very low-level perspective, generating a digital watermark can be regarded as associating a watermark value  $w(\mathbf{i}, \mathbf{p})$  to each pixel position in the image. However, if the embedded watermark is required to be immune against BRA, the following property should also be verified:

$$\mathbf{f}(\mathbf{i}, \mathbf{p}_0) \approx \sum_k \lambda_k \mathbf{f}(\mathbf{i}, \mathbf{p}_k) \Rightarrow w(\mathbf{i}, \mathbf{p}_0) \approx \sum_k \lambda_k w(\mathbf{i}, \mathbf{p}_k) \quad (8.2)$$

In other words, if at a given position  $\mathbf{p}_0$ , the local neighborhood is similar to a linear combination of neighborhoods at other locations  $\mathbf{p}_k$ , then the watermark sample  $w(\mathbf{p}_0)$  embedded at position  $\mathbf{p}_0$  should be close to the linear combination (with the same mixing coefficients  $\lambda_k$ ) of the watermark samples  $w(\mathbf{p}_k)$  at these locations. A simple way to obtain this property is to make the watermarking process be the composition of a feature extraction operation and a linear form  $\varphi$ .


Hence, one can write  $w = \varphi \circ \mathbf{f}$  where  $\varphi : \mathcal{F} \rightarrow \mathbb{R}$  is a linear form which takes  $F$ -dimensional feature vectors in input. Next, to completely define this linear form, it is sufficient to set the values  $\xi_f = \varphi(\mathbf{b}_f)$  for a given orthonormalized basis  $\mathcal{B} = \{\mathbf{b}_f\}$  of the feature space  $\mathcal{F}$ . Without loss of generality, one can consider the canonical basis  $\mathcal{O} = \{\mathbf{o}_f\}$  where  $\mathbf{o}_f$  is a  $F$ -dimensional vector filled with 0's except the  $f$ -th coordinate which is equal to 1. The whole secret of the algorithm is contained in the values  $\xi_f$  and they can consequently be pseudo-randomly generated using a secret key  $K$ . Now, if the values taken by the linear form on the unit sphere  $\mathcal{U}$  of this subspace are considered, the following probability density function is obtained:

$$f_{\varphi|\mathcal{U}}(w) = \frac{1}{\|\boldsymbol{\xi}\| \sqrt{\pi}} \frac{\Gamma(\frac{F}{2})}{\Gamma(\frac{F-1}{2})} \left[ 1 - \left( \frac{w}{\|\boldsymbol{\xi}\|} \right)^2 \right]^{\frac{F-3}{2}} \quad (8.3)$$

where  $\|\boldsymbol{\xi}\|^2 = \sum_{f=1}^F \xi_f^2$  and  $\Gamma(\cdot)$  is the Gamma function.

*Proof.* For any feature vector  $\mathbf{x} = \sum_{f=1}^F x_f \mathbf{o}_f$ , the associated watermark sample  $\varphi(\mathbf{x})$  is given by:

$$\varphi(\mathbf{x}) = \sum_{f=1}^F x_f \varphi(\mathbf{o}_f) = \sum_{f=1}^F x_f \xi_f = \mathbf{x} \cdot \boldsymbol{\xi} \quad (8.4)$$



**Figure 8.1:** Geometrical point of view of the linear form  $\varphi$ .

where  $\boldsymbol{\xi}$  is a vector containing the values  $\xi_f$  taken by the linear form  $\varphi$  on the canonical basis  $\mathcal{O}$ . When the feature vector  $\mathbf{x}$  lies on the unit sphere  $\mathcal{U}$ , the probability that the watermark value  $\varphi(\mathbf{x})$  falls within the range  $[w - dw/2, w + dw/2]$  is given by  $f_{\varphi|\mathcal{U}}(w)dw$  for small  $dw$ . Looking at Figure 8.1, it is easy to see that, from a geometrical point of view, this probability is equal to the proportion of points on the unit sphere whose projection on the vector  $\boldsymbol{\xi}$  falls within  $[x - dx/2, x + dx/2]$ . In other words, it is equal to the ratio between two surfaces:

$$f_{\varphi|\mathcal{U}}(w)dw = \frac{R_D(h)}{S_D(1)} \quad (8.5)$$

where  $S_D(1)$  is the surface of the  $D$ -dimensional unit sphere and  $R_D(h)$  the surface of the small gray ring in Figure 8.1. When this ring is unwrapped, it can be approximated by a  $D$ -dimensional rectangle having dimensions  $d\theta \times S_{D-1}(h)$  i.e.  $R_D(h) = d\theta \cdot S_{D-1}(h) = h^{D-2} \cdot d\theta \cdot S_{D-1}(1)$ . Now, the linear form  $\varphi$  takes its maximal value for the point on the unit sphere which is aligned with the vector  $\boldsymbol{\xi}$ , i.e. for  $\boldsymbol{\xi}_n = \boldsymbol{\xi}/\|\boldsymbol{\xi}\|$ , and this maximal value is equal to  $\varphi(\boldsymbol{\xi}_n) = \|\boldsymbol{\xi}\|$ . Furthermore, let us write each vector on the unit sphere as  $\mathbf{x} = x\boldsymbol{\xi}_n + \mathbf{x}^\perp$ , where  $\mathbf{x}^\perp$  is a vector orthogonal to  $\boldsymbol{\xi}$  i.e.  $\varphi(\mathbf{x}^\perp) = 0$ . As a result, it is easy to obtain:

$$\varphi(\mathbf{x}) = w = \varphi(x\boldsymbol{\xi}_n + \mathbf{x}^\perp) = x\|\boldsymbol{\xi}\| \quad (8.6)$$

Finally, using some basic trigonometry, it is straightforward to find that:

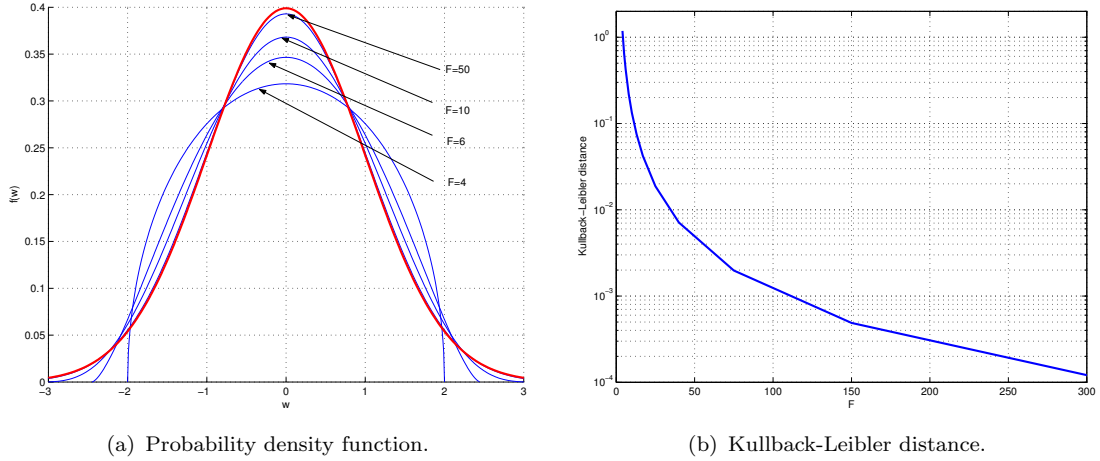
$$h = \sqrt{1 - x^2} = \sqrt{1 - \left(\frac{w}{\|\boldsymbol{\xi}\|}\right)^2} \quad (8.7)$$

$$dx = \frac{dw}{\|\boldsymbol{\xi}\|} = \sin \theta d\theta = \sqrt{1 - \left(\frac{w}{\|\boldsymbol{\xi}\|}\right)^2} d\theta \quad (8.8)$$

Merging all these results together, the following equation is obtained:

$$f_{\varphi|\mathcal{U}}(w) = \frac{1}{\|\boldsymbol{\xi}\|} \frac{S_{D-1}(1)}{S_D(1)} \left[1 - \left(\frac{w}{\|\boldsymbol{\xi}\|}\right)^2\right]^{\frac{F-3}{2}} \quad (8.9)$$

At this point, applying well-known results on the surface of  $D$ -dimensional hyperspheres, Equation (8.3) is immediate.  $\square$



**Figure 8.2:** Evolution of  $f_{\varphi|\mathcal{U}}$  when the feature space dimension increases.

When the dimension  $F$  of the feature space  $\mathcal{F}$  grows large, the probability density function defined in Equation (8.3) tends towards a truncated Gaussian distribution with zero mean and standard deviation  $\|\boldsymbol{\xi}\|/\sqrt{F}$ . Thus if the  $\xi_f$ 's are chosen to have zero mean and unit variance, this ensures that the values of the linear form restricted to the unit sphere  $\mathcal{U}$  are almost normally distributed with also zero mean and unit variance. Figure 8.2 depicts the evolution of the probability density function when the dimension  $F$  grows. In particular, the second plot highlights the fact that the Kullback-Leibler distance with the zero mean and unit variance Gaussian distribution rapidly decreases. At this point, keeping in mind that  $\varphi$  is linear and that the following equation is valid,

$$w(\mathbf{i}, \mathbf{p}) = \varphi\left(\|\mathbf{f}(\mathbf{i}, \mathbf{p})\| \frac{\mathbf{f}(\mathbf{i}, \mathbf{p})}{\|\mathbf{f}(\mathbf{i}, \mathbf{p})\|}\right) = \|\mathbf{f}(\mathbf{i}, \mathbf{p})\| \varphi(\mathbf{u}(\mathbf{i}, \mathbf{p})) \quad \text{with } \mathbf{u}(\mathbf{i}, \mathbf{p}) \in \mathcal{U} \quad (8.10)$$

it is straightforward to realize that the obtained watermark is equivalent to a Gaussian watermark with zero mean and unit variance multiplied by some local scaling factors. The more textured is the considered neighborhood, the more complicated it is to characterize and the greater the norm  $\|\mathbf{f}(\mathbf{i}, \mathbf{p})\|$  is likely to be. Looking back at Equation (8.10), it results that the watermark is amplified in textured area whereas it is attenuated in smooth ones. This can be regarded as some kind of perceptual shaping [189].

## 8.2 A Practical Implementation Using Gabor Features

To impose a linear relationship between watermark samples with respect to some characteristics of the neighborhood, it is first necessary to define the features which will be used to differentiate between neighborhoods i.e. it is required to define the feature extraction function  $\mathbf{f}(\cdot)$  mentioned in Equation (8.1). In this perspective, Gabor features are among the most popular ones and have been now used for a long time for a broad range of applications including image analysis and compression [37], texture segmentation [63], face authentication [60] and facial analysis [59]. Images are classically viewed either as a collection of pixels (spatial domain) or as a sum of sinusoids of infinite extent (frequency domain). But these representations are just two opposite extremes in a continuum of possible joint space/frequency representations. Indeed, frequency can be viewed as a local phenomenon that may vary with position throughout the image. Moreover, Gabor wavelets have also received an increasing interest in image processing since they are particularly close to 2-D receptive fields profiles of the mammalian cortical simple cells [160].

A Gabor Elementary Function (GEF)  $\mathbf{h}_{\rho,\theta}$  is defined by a radius  $\rho$  and an orientation  $\theta$  and the response of an input image  $\mathbf{i}$  to such a GEF can be computed as follows:

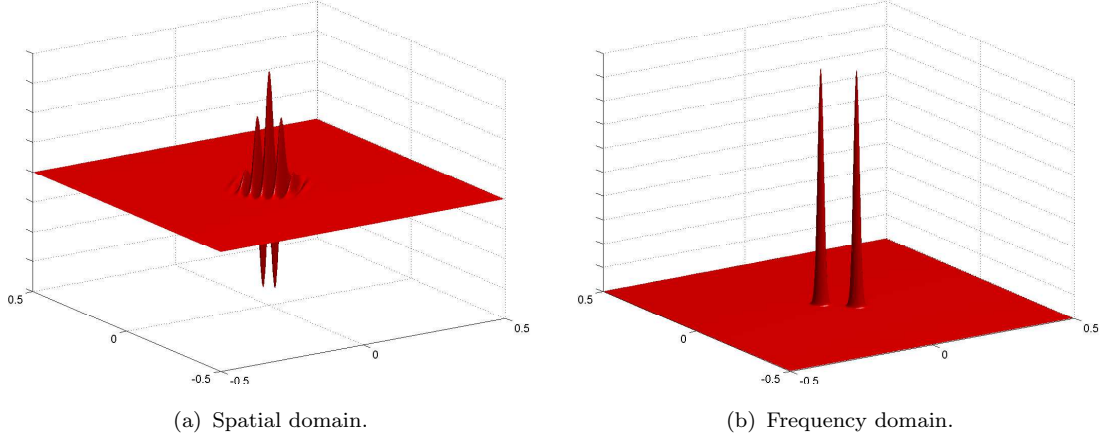
$$\mathbf{g}_{\rho,\theta} = \mathbf{i} * \mathbf{h}_{\rho,\theta} \quad (8.11)$$

where  $*$  denotes convolution and  $\mathbf{g}_{\rho,\theta}$  is the resulting filtered image. The GEF is a complex 2D sinusoid whose orientation and frequency are given by  $(\theta, \rho)$  restricted by a Gaussian envelope. For computational complexity reasons, Gabor filtering is usually performed in the Fourier domain since it then comes down to a simple multiplication with the following filter:

$$\mathbf{H}_{\rho,\theta}(u, v) = \exp \left[ -\frac{1}{2} \left( \left( \frac{u' - \rho}{\sigma_\rho} \right)^2 + \left( \frac{v'}{\sigma_\theta} \right)^2 \right) \right]$$

with  $\begin{pmatrix} u' \\ v' \end{pmatrix} = \mathbf{R}_\theta \begin{pmatrix} u \\ v \end{pmatrix} = \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix} \begin{pmatrix} u \\ v \end{pmatrix}$  (8.12)

where  $\sigma_\rho$  and  $\sigma_\theta$  characterize the bandwidth of the GEF. In other words,  $\mathbf{H}_{\rho,\theta}$  is a 2D Gaussian that is shifted  $\rho$  frequency units along the frequency  $u$ -axis and rotated by an angle  $\theta$ . Thus, it acts as a bandpass filter with a center frequency controlled by  $\rho$  and  $\theta$  and a bandwidth regulated by  $\sigma_\rho$  and  $\sigma_\theta$ . To obtain real valued features  $\mathbf{g}_{\rho,\theta}$  in the spatial domain, GEFs are paired as follows  $\mathbf{H}_{\rho,\theta} \leftarrow \mathbf{H}_{\rho,\theta} + \mathbf{H}_{\rho,\theta+\pi}$ . As depicted in Figure 8.3, such a GEF pair comes down in the spatial domain to real-valued 2D cosine waveform whose orientation and frequency are defined by  $(\theta, \rho)$  restricted by a Gaussian envelope.



**Figure 8.3:** GEF pair visualization.

A single GEF pair associates to each pixel  $\mathbf{p}$  of the image a single feature value  $\mathbf{g}_{\rho,\theta}(\mathbf{i}, \mathbf{p})$ . As a result, the idea is now to design a filter bank of such GEF pairs to obtain for each pixel a multi-dimensional feature vector  $\mathbf{g}(\mathbf{i}, \mathbf{p}) = \{\mathbf{g}_{\rho_{i,j}, \theta_{i,j}}(\mathbf{i}, \mathbf{p})\}$  with  $1 \leq i \leq M$  and  $1 \leq j \leq N$ . Based on previous work [60], the different parameters of the GEF pairs are computed as follows:

$$\rho_{i,j} = \rho_{\min} + b \frac{(s+1)s^{i-1} - 2}{s-1} \quad (8.13)$$

$$\sigma_{\rho_{i,j}} = tbs^{i-1} \quad (8.14)$$


$$\theta_{i,j} = \frac{(j-1)\pi}{N} \quad (8.15)$$

$$\sigma_{\theta_{i,j}} = t \frac{\pi \rho_{i,j}}{2N} \quad (8.16)$$

$$b = \frac{\rho_{\max} - \rho_{\min}}{2} \left( \frac{s-1}{s^M - 1} \right) \quad (8.17)$$

The whole filter bank is specified by the 6 parameters  $M$ ,  $N$ ,  $\rho_{\min}$ ,  $\rho_{\max}$ ,  $s$  and  $t$ . The first two parameters determine respectively the number of orientations and frequencies in the filter bank. The next two ones specify the bandwidth within which the GEFs are bound. The parameter  $s$  controls how much the radial





**Figure 8.4:** Graphical representation in the Fourier domain of the GEFs levelset for value  $1/e$  with  $M = 8$ ,  $N = 4$ ,  $s = 2$  and  $t = 1$ .

bandwidth increases when the radius increases. For instance, when it is set to 2, frequency bands are distributed in octave steps with a frequency bandwidth which doubles at each step. Finally, the parameter  $t$  sets the value at which neighboring filters intersect. As an example, with  $t = 1$ , they cross at equal value  $1/e$  along their principal axis. Figure 8.4 depicts how GEFs are scattered throughout a specified frequency ring in the Fourier domain.

In each pixel position  $\mathbf{p}$ , the resulting  $MN$ -dimensional vector  $\mathbf{g}(\mathbf{i}, \mathbf{p})$  can be regarded as the local power spectrum of the image and thus be used to characterize the neighborhood. It should be noted that if the Gabor filter bank is properly designed, it is possible to impose higher constraints. For instance, if the fractal approach depicted in Figure 6.4 is enforced, neighborhoods which are the same modulo a small set of geometrical operations, e.g. 8 isometries and downsampling by a factor 2, are required to carry the same watermark samples to achieve robustness [158]. Such constraints need to be taken into account to define the kernel of the linear form  $\varphi$  i.e. the non null vectors  $\mathbf{v}$  for which  $\varphi(\mathbf{v}) = 0$ . However, more constraints induce a lower dimensional subspace for watermarking which can rapidly become critical.

### 8.3 Relationship with Multiplicative Watermarking Schemes in the Frequency Domain

Since the values  $\xi_f$  of the linear form  $\varphi$  are defined on the canonical basis  $\mathcal{O}$  when Gabor features are considered, the watermark sample obtained at position  $\mathbf{p}$  is simply given by:

$$\mathbf{w}(\mathbf{i}, \mathbf{p}) = \sum_{f=1}^F \xi_f \mathbf{g}_f(\mathbf{i}, \mathbf{p}) \quad (8.18)$$

where  $\mathbf{g}_f(\mathbf{i}, \mathbf{p})$  is the  $f$ -th coordinate of the  $F$ -dimensional Gabor feature vector  $\mathbf{g}(\mathbf{i}, \mathbf{p})$ . In other words, the watermark is a linear combination of different Gabor responses  $\mathbf{g}_f$ . However, when  $M$  and  $N$  grow, more and more Gabor responses need to be computed which can be quickly computationally prohibitive. Hopefully, when the Fourier domain is considered, the watermark can be computed as follows:

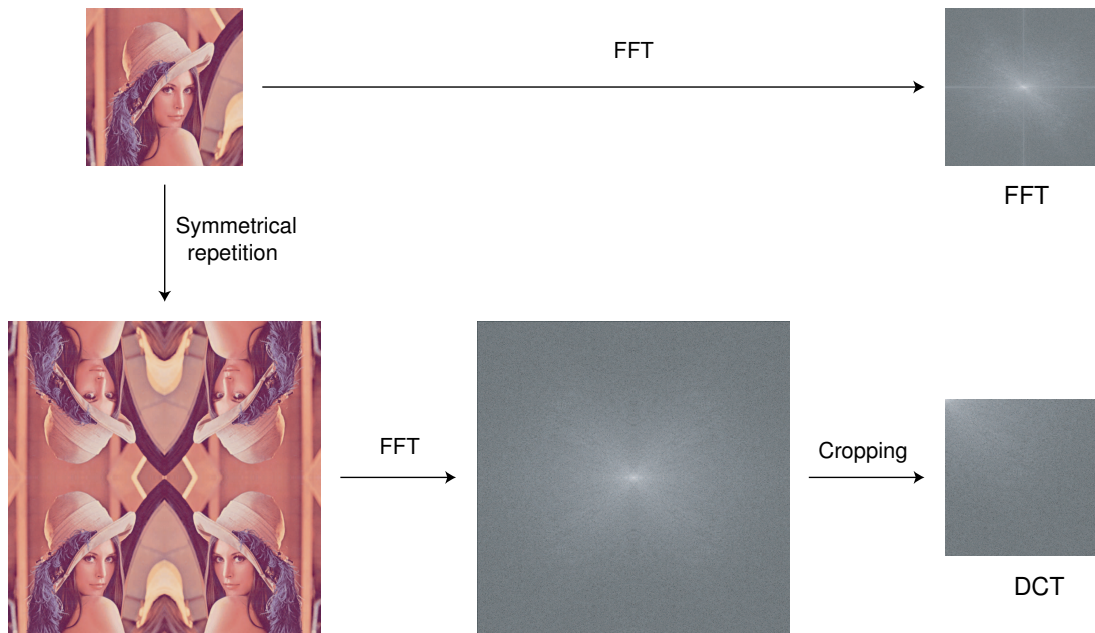
$$\begin{aligned} \mathbf{W}(\mathbf{i}, \mathbf{q}) &= \sum_{\mathbf{p} \in \mathcal{P}} \left( \sum_{f=1}^F \xi_f \mathbf{g}_f(\mathbf{i}, \mathbf{p}) \right) \omega_{\mathbf{p}, \mathbf{q}} \\ &= \sum_{f=1}^F \xi_f \left( \sum_{\mathbf{p} \in \mathcal{P}} \mathbf{g}_f(\mathbf{i}, \mathbf{p}) \omega_{\mathbf{p}, \mathbf{q}} \right) = \sum_{f=1}^F \xi_f \mathbf{G}_f(\mathbf{i}, \mathbf{q}) \\ &= \sum_{f=1}^F \xi_f \mathbf{H}_f(\mathbf{q}) \mathbf{I}(\mathbf{q}) = \mathbf{H}(K, \mathbf{q}) \mathbf{I}(\mathbf{q}) \end{aligned} \quad (8.19)$$

$$\text{with } \mathbf{H}(K, \mathbf{q}) = \sum_{f=1}^F \xi_f \mathbf{H}_f(\mathbf{q})$$

where  $\omega_{\mathbf{p}, \mathbf{q}} = \exp[-j2\pi((x-1)(u-1)/X + (y-1)(v-1)/Y)]$ , capital letters indicate FFT-transformed variables and  $\mathbf{q} = (u, v)$  denotes a frequency position with  $1 \leq u \leq U$  and  $1 \leq v \leq V$ . In other words, the watermark can be generated in one row in the Fourier domain by computing  $\mathbf{H}$  and such an approach is likely to significantly reduce the computational cost.

Looking closely at Equation (8.19), it is straightforward to realize that the watermark generation process comes down to a simple multiplication between the image spectrum  $\mathbf{I}$  and some pseudo-random signal  $\mathbf{H}(K)$ . In other words, it really looks similar to basic well-known multiplicative embedding schemes in the frequency domain [31, 7]. When the bandwidth of a GEF is close to 0, the 2D Gaussian in the Fourier domain tends toward a Dirac impulse centered at coordinates  $(\rho, \theta)$  i.e. it tends toward an infinite sinusoid in the spatial domain. Therefore, multiplicative embedding in the FFT domain<sup>1</sup> is equivalent to imposing a linear relationship on the watermark samples according to the neighborhood

<sup>1</sup>In this thesis, multiplicative embedding in the FFT domain means that the *complex* FFT



**Figure 8.5:** Relationship between FFT and DCT.

which is characterized by its response to infinite sinusoids. Under this new light, FFT multiplicative watermarks can be seen as a special case of the Gabor watermarks introduced in Section 8.2 and should therefore be also coherent with the host signal i.e. they should resist to BRA. Next, it is useful to remind that DCT coefficients are simply FFT coefficients of some periodic image [120]. When the FFT of an image is computed, it is implicitly assumed that the input signal is periodic both horizontally and vertically. However this may create artificial discontinuities at the border of the image and it is revealed by a *cross artifact* in the FFT domain. A simple way to eliminate these artificial discontinuities is to repeat the signal in a symmetrical fashion as depicted in Figure 8.5 before applying the FFT. Now, it is immediate to assert that DCT multiplicative watermarks [31] should also be signal coherent and resist to BRA. At this point, it is interesting to note that multiplicative watermarking in the frequency domain was initially motivated by contrast masking properties: larger coefficients can convey a larger watermark value without compromising invisibility [69]. This can be related with the natural perceptual shaping of signal coherent watermarks exhibited in Equation (8.10).

---

coefficients are multiplied by pseudo-random values. It is slightly different from the algorithm described in [7] where only the *magnitude* of the FFT coefficients were watermarked.

## 8.4 Investigations

The major claim in this chapter is that a watermark whose samples have inherited the same linear relationships as the neighborhoods of the host signal should not be affected by BRA. An embedding scheme using Gabor features has been designed in Section 8.2 so that the generated watermark exhibits this property. Moreover, it has been shown in Section 8.3 that previous embedding schemes based on multiplicative embedding in a frequency space, either FFT or DCT, is also likely to resist BRA. It is consequently pertinent to check whether or not these identified watermarks are degraded by such attacks in comparison with other current watermarks e.g. additive SS watermarks in the spatial domain. To this end, large-scale experiments have been conducted. The experimental protocol is first detailed in Subsection 8.4.1. Next, the influence of the number of GEFs considered to characterize the local neighborhood has been investigated in Subsection 8.4.2 with respect to the immunity against BRA. Finally, the performances of the three different proposed signal coherent watermarks have been compared in Subsection 8.4.3.

### 8.4.1 Protocol

A watermark with zero mean and unit variance  $\mathbf{w}(K, \mathbf{i})$  is embedded in the input image  $\mathbf{i}$  to obtain a watermarked image  $\mathbf{i}_w$  according to the following embedding rule:

$$\mathbf{i}_w = \mathbf{i} + \alpha \mathbf{w}(K, \mathbf{i}) \quad (8.20)$$

where  $K$  is a secret key used to generate the watermark and  $\alpha$  an embedding strength equal to 3 so that the embedding process results in a distortion about 38 dB in terms of Peak Signal to Noise Ratio (PSNR). Four different watermark generation processes will be surveyed during the experiments:

**SS:** The embedded watermark is completely independent of the host content i.e.  $\mathbf{w}(K, \mathbf{i}) = \mathbf{r}(K)$  where  $\mathbf{r}(K)$  is a fixed pseudo-random pattern which is generated using the secret key  $K$  and which is normally distributed with zero mean and unit variance.

**Gabor:** The generation process considers Gabor features to make the watermark inherit the self-similarities of the host signal. As discussed in Subsection 8.3, the watermark is generated in the Fourier domain using Equation (8.19) i.e.  $\mathbf{W}(K, \mathbf{i}) = \mathbf{H}(K) \mathbf{I}$ . Inverse FFT is then performed to come back to the spatial domain and the resulting watermark is globally scaled to have unit variance.

**FFT:** The watermark is generated in the Fourier domain as follows  $\mathbf{W}(K, \mathbf{i}) = \mathbf{r}(K) \mathbf{I}$  where  $\mathbf{r}(K)$  is a fixed pseudo-random pattern which is symmetric

with respect to the center of the Fourier domain and which has value 0 at the DC coefficient position. This property has to be verified so that the resulting watermark is real-valued with zero mean after inverse transform. Once again, inverse FFT is performed to come back to the spatial domain and the resulting watermark is scaled to have unit variance. This algorithm can be regarded as an extension of the previous one when the GEFs are reduced to Dirac impulses in the frequency domain.

**DCT:** The watermark is generated in the frequency domain using the following formula  $\hat{\mathbf{W}}(K, \mathbf{i}) = \mathbf{r}(K) \hat{\mathbf{I}}$  where “capital hat” denotes the DCT transform and  $\mathbf{r}(K)$  is a normally distributed pseudo-random pattern which has value 0 at the DC coefficient position. Inverse DCT is then performed to come back to the spatial domain and the resulting watermark is scaled to have unit variance.

Next, the watermarked image  $\mathbf{i}_w$  is attacked using the latest version of BRA described in Table 6.4. In the experiments,  $8 \times 8$  blocks have been considered with an overlap of 4 pixels and the search window size has been set to  $64 \times 64$ . Furthermore, the two thresholds  $\tau_{\text{low}}$  and  $\tau_{\text{high}}$  have been set equal to the same value  $\tau_{\text{target}}$ . As a result, the replacement block is obtained by considering more or less eigenblocks so that the distortion with the original signal block is as close as possible to the target value  $\tau_{\text{target}}$ . This threshold can be used as an attacking strength which can be modified during experiments.

On the detector side, the only concern is to know whether or not the embedded watermark has survived. Therefore, non-blind detection can be considered and the residual correlation is computed as follows:


$$d(\mathbf{i}, \tilde{\mathbf{i}}_w) = (\tilde{\mathbf{i}}_w - \mathbf{i}) \cdot \mathbf{w}(K, \tilde{\mathbf{i}}_w) \quad (8.21)$$

where  $\tilde{\mathbf{i}}_w$  is the attacked image and  $\cdot$  denotes the linear correlation operation. To anticipate future blind detection, the detector generates the watermark using the attacked image instead of the original image. This has no impact for SS since it is content independent, but this may have one with signal coherent watermarks. The residual correlation should be equal to  $\alpha$  if the watermark has survived while it should drop down to 0 when the watermark signal has been completely washed out. As a result, the presence of the watermark can be asserted by comparing the residual correlation  $d(\mathbf{i}, \tilde{\mathbf{i}}_w)$  with a detection score  $\tau_{\text{detect}}$  which can be set to  $\alpha/2$  for equal false positive and false negative probabilities.

### 8.4.2 Influence of the Number of GEFs

A database of 500 images of size  $512 \times 512$  has been considered for experiments. It contains snapshots, synthetic images, drawings and cartoons. All the images

are first watermarked using either the SS or the Gabor watermark generation process. In this subsection, the influence of the Gabor filterbank design will be investigated. Therefore, the parameters have been set as follows:  $N = 16$ ,  $\rho_{\min} = 0.01$ ,  $\rho_{\max} = 0.45$ ,  $s = 2$  and  $t = 1.5$ . Moreover, the number of orientations  $M$  has been successively set to 2, 4, 8, 16 and 32. Next, each watermarked image is submitted to BRA with varying attacking strength  $\tau_{\text{target}}$  to obtain a distortion vs. residual correlation curve. Finally, all the curves associated with a given watermarking method are averaged to depict the statistical behavior of this scheme against BRA. Those results have been gathered in Figure 8.6. It should




**Figure 8.6:** Number of GEFs vs. performances against BRA.

be reminded that the goal of the attacker is to decrease the residual correlation while maintaining the image quality. First of all, it should be noted that the proposed Gabor watermark generation process clearly outperforms the SS one when the resilience against BRA is considered. Indeed, the residual correlation never goes below 2.5 with Gabor watermarks while it already drops below the detection threshold  $\tau_{\text{detect}} = 1.5$  for a distortion of 40 dB when SS watermarks are considered. Furthermore, even if more images should be considered to allow a pertinent comparison, one can already assert that the number of GEFs used to characterize the local neighborhood has not a drastic impact on the immunity of the watermark against BRA. On the other hand, it is important to increase the number of GEF pairs so that watermarks generated with different secret keys  $K$  are as little correlated as possible and thus decrease the false positive probability. Of course, increasing the number of GEF pairs also raises the computational load and a trade-off has to be found.

### 8.4.3 Signal Coherent Watermarks Comparison

In a second row, FFT and DCT systems have been surveyed to check whether or not they produce BRA-immune signal coherent watermarks as predicted in Section 8.3. The same database of images has been considered and all the images have been first watermarked using one of the watermarking systems under study i.e. SS, Gabor, FFT or DCT. In this case, the Gabor filter bank has been configured as follows:  $M = 32$ ,  $N = 16$ ,  $\rho_{\min} = 0.01$ ,  $\rho_{\max} = 0.45$ ,  $s = 2$  and  $t = 1.5$ . This results in 4 collections of 500 watermarked images each. Next, each watermarked image is submitted to BRA with varying attacking strength  $\tau_{\text{target}}$  to obtain a distortion vs. residual correlation curve. Finally, all the curves associated with a given watermarking method are averaged to depict the statistical behavior of this scheme against BRA. These results have been gathered in Figure 8.7. First of all, experimental results clearly show that Gabor, FFT and



**Figure 8.7:** Comparison of the impact of BRA with the 4 watermarking schemes under study: whereas non coherent watermarks (SS) are washed out when the attacking strength increases, coherent watermarks (Gabor/FFT/DCT) survive.

DCT watermarks share the same immunity against BRA. As a matter of fact, the residual correlation never goes below 2.5 with signal coherent watermarks (Gabor, FFT or DCT) while it already drops below the detection threshold  $\tau_{\text{detect}} = 1.5$  for a distortion of 40 dB when SS watermarks are considered. Moreover, even if experiments at a larger scale should be carried out for a pertinent comparison, some kind of *ranking* appears amongst the signal coherent watermarking schemes. The observation that FFT behaves better than Gabor may be explained by the fact that the first algorithm is an extension of the second one. Therefore, the FFT curve would give some bound for the achievable performances with the Ga-

bor scheme for different filter bank configurations. Finally, the superiority of DCT over FFT might be due to the properties of the DCT which ensure that the watermark will not be embedded in *fake* image frequencies revealed by the Fourier transform [69].

## 8.5 Discussion

Block Replacement Attacks (BRA) are recognized to be among the most critical signal manipulations against watermarking systems today. Typically, these attacks exploit the fact that *similar blocks do not carry similar watermarks* to confuse the watermark detector. In this chapter, a novel watermarking strategy has been investigated to remove this weak link. It basically aims at making the embedded watermark inherit the self-similarities of the host signal. Gabor features are extracted in each pixel position to characterize the neighborhood and are exploited to export linear relationships between neighborhoods to watermark samples. Moreover, previous multiplicative embedding schemes in the frequency domain [31, 7] have been shown to also produce signal-coherent watermarks even if such a property has never been foreseen before.

From a more general point of view, signal coherent watermarking can be considered as some kind of informed watermarking [35, 66]. Indeed, digital watermarking can be seen as moving a point in a high dimensional media space to a nearby location i.e. introducing a small displacement in a random direction. The introduced framework only stipulates that the host signal self-similarities have to be considered to resist BRA and that in this case some of the possible directions are now prohibited. Nevertheless it is still necessary to explore how former works [7, 26] can be used to design a blind detector for signal coherent watermarks. Furthermore, security investigations have to be conducted to determine whether or not an attacker can gain some knowledge about the imposed watermarking structure. As demonstrated in Chapter 5, using a redundant watermarking structure can indeed lead to security pitfalls.



---

## Conclusion

---

Ten years after its infancy, digital watermarking is still considered as a young technology. Despite the fact that it has been introduced for security-related applications such as copyright protection or fingerprinting, almost no study has been conducted to evaluate the security of watermarking algorithms i.e. to assert the survival of embedded watermarks in a hostile environment. This may explain in part the failure of the few trials to insert digital watermarking in content distribution frameworks [64, 167]. Therefore, a security analysis based on collusion attacks has been conducted in this Ph.D. thesis and possible countermeasures have been proposed as summarized in Section 9.1. Indeed, although this security issue has been neglected for years in the watermarking community, it is now becoming a hot topic. This is revealed by the recent sessions dedicated to this subject: special session on *Watermarking Security* in SPIE'05, special session on *Media Security* in WIAMIS'05, special session entitled *Watermarking Security: Where Do We Stand?* in IWDW'05. Of course, this Ph.D. thesis does not reveal an ideal secure watermarking algorithm and tracks for future work are given in Section 9.2. Security plays a key role in many applications and it has to be carefully evaluated so that digital watermarking can eventually be introduced in DRM frameworks released to a large audience.

### 9.1 Summary and Contributions

This Ph.D. thesis has highlighted the fact that the lack of security evaluation in digital watermarking system has led to critical pitfalls against statistical analysis, also referred to as collusion attacks. Such attacks typically consider several

watermarked documents and combine them to produce unwatermarked content. This threat is all the more relevant when digital video is considered since each individual video frame can be regarded as a single watermarked document by itself. Indeed, video watermarking is still regarded most of the time as an extension of watermarking for still images and frame-by-frame embedding strategies are commonly enforced.

In this perspective, two alternative collusion strategies have been investigated. The first one basically assumes that a redundant watermarking structure has been used. In this case, an attacker can eavesdrop the watermarking channel, i.e. examine a collection of watermarked documents, to isolate some suspicious and unexpected patterns. In particular, it is possible to approximate a watermark which has been redundantly embedded, or to estimate a small set of secret watermark patterns using vector quantization, or even to identify a low-dimensional subspace within which embedded watermarks are bounded thanks to space reduction techniques. Once the attacker has acquired such a knowledge, it is usually relatively simple to design an efficient remodulation attack to confuse the watermark detector. Although simple additive watermarks have been considered in this analysis, such estimate-and-remodulate attacks can also be relevant to defeat other watermarking techniques such as the Scalar Costa Scheme (SCS) for instance.

However, even if using a redundant watermarking structure is not secure, using completely independent watermarks is not the solution either. If an attacker can collect similar documents carrying uncorrelated watermarks, averaging them will sum the watermark samples to zero. Moreover, since video content exhibits both temporal and spatial redundancy, efficient attacks can be designed which basically replace each part of the signal with a similar one taken from another location in the same signal or with a combination of similar parts. In particular, temporal redundancy between successive video frames can be exploited to estimate each single frame with the neighbor ones using Temporal Frame Averaging after Registration (TFAR). Additionally, spatial redundancy can also be considered to design powerful Block Replacement Attack (BRA) inspired from fractal coding. Such an approach has been demonstrated to be more critical than common signal processing primitives such as JPEG compression and Gaussian filtering.

Once these security pitfalls have been isolated, two main countermeasures have been proposed in this Ph.D. thesis to circumvent them. First, video mosaicing has been considered to produce motion-compensated watermarks and thus resist to TFAR. The underlying idea is that successive video frames are multiple 2D projections of the same 3D movie set. Therefore, motion-compensated watermarking forces a physical 3D point of the movie set to always carry the same watermark sample whenever it is projected onto the screen. In other words, everything goes as if the camera is filming a scene which is already watermarked. Even

if the mosaicing process may be too computationally expensive to be deployed in a real-life application, this approach has led to interesting results, in particular regarding the imperceptibility of the embedded watermark and this is likely to encourage further research in this direction to produce motion-compensated watermarks by other means.

Next, the notion of similarity inheritance has been introduced to combat BRA. Such attacks exploit indeed the fact that similar signal blocks do not carry similar watermarks. The idea is consequently to examine some features which characterize in some sense the local neighborhood to force host samples with similar neighborhoods to carry close watermark values. To this end, a linear form can be defined in the feature space so that the watermark inherits the linear relationships between neighborhoods. In this Ph.D. thesis, a practical implementation using Gabor features has been presented. Using Gabor features has also revealed an unexpected relationship with earlier multiplicative embedding schemes in the frequency domain i.e. these schemes also produce signal coherent watermarks and thus resist to BRA. This is likely to renew the interest for such schemes despite the fact they were almost abandoned for the last few years.

This work resulted in one book chapter,

1. G. Doërr and J.-L. Dugelay, *Video Watermarking: Overview and Challenges*, Chapter 42 in *Handbook of Video Databases: Design and Applications*, by B. Furht and O. Marques (editors), CRC Press, 2003

three international journal articles,

1. G. Doërr and J.-L. Dugelay, *A Guide Tour of Video Watermarking*, in *Signal Processing: Image Communication*, vol. 18, no. 4, pp. 263–282, 2003.
2. G. Doërr and J.-L. Dugelay, *Security Pitfalls of Frame-by-Frame Approaches to Video Watermarking*, in *IEEE Transactions on Signal Processing, Supplement on Secure Media*, vol. 52, no. 10, pp. 2955–2964, 2004.
3. G. Doërr, J.-L. Dugelay and D. Kirovski, *On the Need for Signal-Coherent Watermarks*, accepted for publication in the *IEEE Transactions on Multimedia*.

1 national journal article,

1. G. Doërr and J.-L. Dugelay, *Problématique de la Collusion en Tatouage Vidéo*, submitted for publication in *Traitement du Signal*, 2005.

12 international conference/workshop papers,

1. C. Rey, G. Doërr, J.-L. Dugelay and G. K. Csurka, *Toward Generic Image Dewatermarking?*, in *Proceedings of the IEEE International Conference on Image Processing*, vol. III, pp. 633–636, 2002.

2. G. Doërr and J.-L. Dugelay, *New Intra-Video Collusion Attack Using Mosaicing*, in Proceedings of the IEEE International Conference on Multimedia and Expo, vol. II, pp. 505–508, 2003.
3. G. Doërr and J.-L. Dugelay, *Secure Video Watermarking Via Embedding Strength Modulation*, in Proceedings of the Second International Workshop on Digital Watermarking, LNCS 2939, pp. 340–354, 2003.
4. G. Doërr and J.-L. Dugelay, *Secure Background Watermarking Based on Video Mosaicing*, in Security, Steganography and Watermarking of Multimedia Contents VI, Proceedings of SPIE 5306, pp. 304–314, 2004.
5. G. Doërr and J.-L. Dugelay, *Danger of Low-Dimensional Watermarking Subspaces*, in Proceedings of the IEEE International Conference On Acoustics, Speech and Signal Processing, vol. III, pp. 93–96, 2004.
6. G. Doërr, J.-L. Dugelay and L. Grangé, *Exploiting Self-Similarities to Defeat Digital Watermarking Systems - A Case Study on Still Images*, in Proceedings of the ACM Multimedia and Security Workshop, pp. 133–142, 2004.
7. G. Doërr and J.-L. Dugelay, *Collusion Issue in Video Watermarking*, in Security, Steganography and Watermarking of Multimedia Content VII, Proceedings of SPIE 5681, pp. 685–696, 2005.
8. G. Doërr, C. Rey and J.-L. Dugelay, *Watermark Resynchronization based on Elastic Graph Matching*, in Proceedings of the International Conference on Sciences of Electronic, Technologies of Information and Telecommunications, 2005.
9. G. Doërr and J.-L. Dugelay, *Signal Coherent Watermarking in Video*, in Proceedings of the 6th International Workshop on Image Analysis for Multimedia Interactive Services, 2005.
10. G. Doërr and J.-L. Dugelay, *How to Combat Block Replacement Attacks?*, in Pre-proceedings of the 7th Information Hiding Workshop, pp. 137–151, 2005.
11. G. Doërr and J.-L. Dugelay, *A Countermeasure to Resist Block Replacement Attacks*, accepted for publication in the IEEE International Conference on Image Processing, 2005.
12. G. Doërr and J.-L. Dugelay, *Countermeasures for Collusion Attacks Exploiting Host Signal Redundancy*, accepted for publication in the 4th International Workshop on Digital Watermarking, 2005.

and one national conference/workshop paper:

1. G. Doërr et J.-L. Dugelay, *Problématique de la Collusion en Tatouage Vidéo - Attaques et Ripostes*, submitted for publication in COMpression et REprésentation des Signaux Audiovisuels, 2005.

## 9.2 Tracks for Future Work

It is now almost the end of this manuscript. It does not mean that an *ideal* watermarking system has been found. It only means that after 3 years and a half this Ph.D. thesis has to be stopped even if several open issues remain. A few tracks for future work are presented below:

**Motion compensated watermarking.** In Chapter 7, motion compensated watermarking has been considered as a possible mean to circumvent Temporal Frame Averaging after Registration (TFAR). However, the proposed implementation relies on video mosaicing which induces a high computational cost. Therefore, it may be interesting to investigate alternative ways to produce such motion compensated watermarks. One possible approach could be to exploit the motion prediction/compensation module which is present in many video codecs. Alternatively, one can also investigate whether multiplicative watermarking in 3D transforms also produce successive watermarks which are coherent with the camera motion.

**Further studies with signal coherent watermarks.** In Chapter 8, a signal coherent watermark is produced by defining a linear form in some feature space. This can be seen as imposing some kind of watermarking structure. However, it has also been demonstrated in Chapter 5 that such structures might be estimated by hostile attackers. Therefore, such signal coherent watermarks should be carefully surveyed to know wheter or not a statistical analysis can reveal some critical information about the watermarking system.

**Improve Scalar Costa Scheme (SCS).** In Section 5.4, it has been shown that SCS also leaks some information which can be exploited to defeat the system. Recent advances have proposed to use adaptive quantization step sizes [146, 117] to make SCS robust against valumetric scaling which is recognized to be the Achille's heel of such systems. Nevertheless, even if those modifications slightly complicate the task of the attacker, they do not make the algorithm more secure since they are not key dependent. A dithering term can be introduced to enhance the security of SCS but it should be made content-dependent so that the dithering sequence differs

from one document to the other. A possible way to achieve this may be to use a signal coherent watermark for dithering.

**Improve trellis dirty paper watermarking.** Trellis dirty paper watermarking [136] is today one of the most secure algorithm with respect to estimate-and-remodulate attacks. However, it is likely to be weak against Block Replacement Attacks (BRA) since it relies on an additive embedding rule. It could be consequently interesting to modify the embedding process to use a multiplicative rule and thus produce a signal coherent watermark.

**Generic security metric.** Recent works have considered information theory to define a security metric based on equivocation as given in Equation (5.51). However, this only evaluates the information leakage about the secret watermark i.e. it only addresses eavesdropping attacks. In other words, it neglects attacks such as BRA or TFAR which directly estimate the original non-watermarked document. In such cases, it may be also relevant to consider the following equation:

$$H(\mathbf{o}|\mathbf{o}_1, \dots, \mathbf{o}_N) = H(\mathbf{o}) - I(\mathbf{o}; \mathbf{o}_1, \dots, \mathbf{o}_N) \quad (9.1)$$

where  $\mathbf{o}$  is an original document and  $\{\mathbf{o}_i\}$  a set of  $N$  different watermarked version of it. Next, it would be useful to define some unified metric which address both issues.

**Independent Component Analysis (ICA).** ICA is now receiving a growing interest in the signal processing community. As depicted in Figure 3.4, this transform produces different components which are related with the semantic meaning of the considered video sequence. This opens avenues to design new video watermarking algorithms which are coherent with the content of the video.

**New applications for digital watermarking.** Chapter 2 has given a rapid overview of the possible applications for digital video watermarking. Nevertheless, it is now commonly admitted in the watermarking community that no algorithm exhibits enough good performances to be released in a hostile environment. Nevertheless, a large number of applications such as data hiding do not have any security requirement. In this context, a significant effort is devoted to identify the killer application for digital watermarking.

# A

---

## Résumé Français (Version Longue)

---

### A.1 Introduction

La fin du XXème siècle a vu le monde lentement basculer de l'analogique au numérique. De nos jours, les équipements numériques (lecteurs CD/DVD, ordinateurs, assistants personnels, baladeurs) sont de plus en plus répandus. Cependant, cette formidable révolution technique n'a pas été assez encadrée en termes de protection des droits numériques, et tout particulièrement en termes de protection des droits d'auteur. Alors que chaque génération de copies analogiques introduisait une dégradation supplémentaire, les copies numériques sont parfaites. De plus, les réseaux d'échange de fichiers pair à pair permettent d'échanger facilement de très grands volumes de données multimédia. Bien entendu, cette situation a rapidement suscité l'inquiétude des fournisseurs de contenus qui ont vu leurs ventes chuter de façon significative. Ces derniers sont donc particulièrement attentifs à toute nouvelle technologie qui permettrait d'améliorer la gestion des droits numériques et d'empêcher la redistribution illégale de contenus multimédia protégés par des droits d'auteur. Dans cette optique, le tatouage numérique a été introduit au début des années 90 comme un mécanisme de sécurité complémentaire au cryptage. En effet, tôt ou tard, les données cryptées doivent être décryptées pour les rendre accessibles aux utilisateurs. à ce moment précis, les données numériques ne sont plus protégées par le cryptage et peuvent être éventuellement copiées et redistribuées à grande échelle.

Le tatouage numérique a donc été introduit comme une seconde ligne de défense. L'idée de base consiste à protéger un document numérique en enfouissant un signal codant de l'information de façon robuste et imperceptible [35]. Il existe

un compromis entre trois paramètres conflictuels: la capacité, l'imperceptibilité et la robustesse. La capacité est la quantité d'information insérée dans un document, c'est à dire le nombre de bits codés par le signal tatouage enfoui. En fonction de l'application, le nombre de bits à cacher peut varier. Si quelques bits suffisent à mettre en place un service de contrôle de copie, il est en revanche nécessaire de cacher beaucoup d'information pour permettre l'authentification de documents multimédia. Par ailleurs, le processus de tatouage va inévitablement modifier le signal hôte et introduire une certaine distorsion. La contrainte d'imperceptibilité impose que cette distorsion reste complètement indétectable par un observateur / auditeur. Dans ce but, les caractéristiques du système audio-visuel humain peuvent être exploitées. Par exemple, le signal de tatouage étant souvent considéré comme du bruit, il sera moins perceptible dans les zones texturées d'une image que dans les zones unies. Ainsi, amplifier (resp. atténuer) le signal de tatouage dans les zones texturées (resp. uniformes) permet de diminuer la visibilité du tatouage. Enfin, le tatouage doit être construit de telle sorte qu'il résiste à la plus large palette possible d'opérations qu'un utilisateur puisse effectuer. Cette robustesse face aux traitements usuels du signal (filtrage, compression avec pertes, quantification) est souvent quantifiée en ayant recours à des bancs de test.

Néanmoins, en dépit de nombreux efforts pour optimiser ce compromis complexe entre ces trois paramètres, les quelques tentatives pour introduire un signal de tatouage dans des systèmes de distribution de contenus [167, 64] se sont révélées être des échecs plus ou moins retentissants. L'un des éléments qui explique ces revers est que peu de travaux se sont intéressés à la survie du tatouage face une intelligence malveillante. Ainsi, même si le tatouage numérique a été introduit à l'origine pour des applications vouées à être déployées dans un environnement hostile (contrôle de copie, suivi de copies, etc.), la problématique de la sécurité a été quasiment ignorée. Par conséquent, la section A.2 s'efforce dans un premier temps de définir de façon pertinente la notion de sécurité dans le domaine du tatouage numérique et en particulier d'établir une distinction avec le concept de robustesse. De plus, les attaques par collusion sont introduites comme un moyen possible pour évaluer la sécurité en vidéo. Ainsi, la section A.3 dresse ensuite un vaste panorama d'attaques par collusion et passe en revue par la même occasion les différentes faiblesses des algorithmes de tatouage vidéo communément utilisés actuellement. Une fois ces menaces clairement identifiées, de nouvelles stratégies de tatouage sont introduites dans la section A.4 qui s'efforcent de rendre le signal de tatouage cohérent avec la redondance spatio-temporelle du signal hôte vidéo. Finalement, les différents résultats seront résumés dans la section A.5 et des pistes de recherche seront exposées.



## A.2 Problématique de la sécurité

Quand bien même le tatouage numérique a toujours été étiqueté comme une technologie ayant trait à la sécurité, il n'a jamais été vraiment clair ce à quoi ce terme *sécurité* renvoyait. Au tout début, cela était plus ou moins relié avec le fait qu'une clé secrète est nécessaire pour insérer/extraire le tatouage. Ainsi, une analogie directe avec les principes de Kerckhoffs [97] qui régissent la sécurité en cryptographie. Un exemple très connu est par exemple qu'un algorithme rendu public ne doit pas pouvoir être "cassé" du moment que la clé demeure secrète. Pendant une très longue période, la communauté a pensé que casser un algorithme de tatouage se résumait à effacer le signal de tatouage. Cependant, des utilisateurs n'ayant pas accès à la clé secrète ne devrait pas être en mesure de détecter, estimer, écrire ou modifier le tatouage enfoui [92]. De même, une hypothèse courante en tatouage est que l'attaqueur a accès à un unique document tatoué. Mais en pratique, de nombreuses autres situations sont possibles [8]; par exemple, l'attaqueur peut avoir une collection de documents tatoués, des paires de documents originaux/tatoués, etc. Devant cette situation confuse, les sous-sections qui suivent s'efforceront de donner une définition de la sécurité dans le contexte du tatouage.

### A.2.1 Confiance dans un environnement hostile

Dans de nombreuses applications de tatouage, il est nécessaire d'avoir confiance en l'information transportée par le canal de tatouage. C'est en effet souvent sur cette information que repose le modèle économique d'une application. Dans le contexte d'une application de suivi de copies, le fournisseur de contenu possède un document multimédia de grande valeur qu'il veut distribuer à un large public. Par conséquent, à chaque fois qu'il vend une copie de ce document à un consommateur, il insère un tatouage qui code l'identité du consommateur. Par la suite, si une copie pirate est trouvée, il suffit d'extraire le tatouage pour identifier l'identité de la personne qui n'a pas respecté ses engagements et lancer les poursuites appropriées. L'ensemble du système de protection repose sur la capacité d'identifier à l'aide du tatouage les consommateurs qui font des copies. Afin que ce système fonctionne, il est donc nécessaire qu'une personne n'ayant pas accès à la clé secrète ne puisse pas effacer ou modifier le tatouage insérer. De la même façon, dans une application de contrôle de copie, le tatouage est inséré pour autoriser ou non la copie d'un document. Là encore, le tatouage joue un rôle crucial dans le système de protection. Si un attaqueur est capable d'effacer le tatouage, alors il peut copier "librement" les documents qu'il a déprotégés sans reverser le moindre centime aux auteurs.

En revanche, de leur côté, les consommateurs voient le tatouage comme une protection qui les dérange: il les empêche de copier leurs données numériques

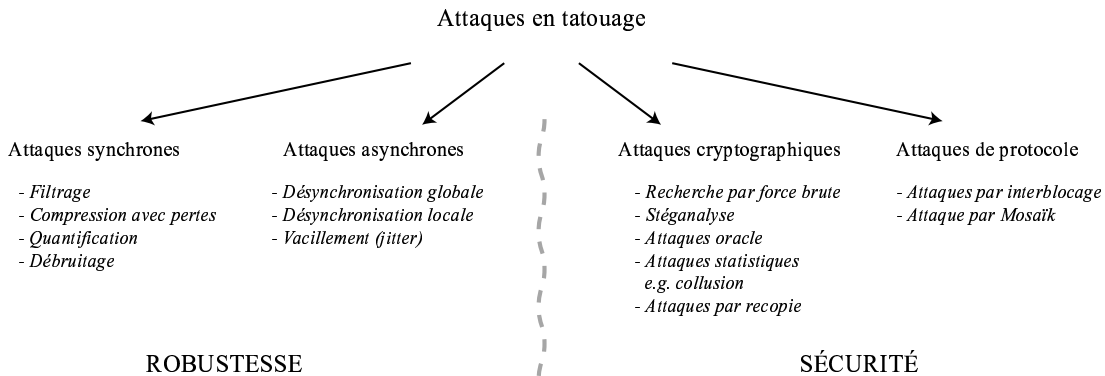
comme ils le souhaitent, il permet de retrouver l'identité des personnes qui ont créé et distribué des copies illégales... Par conséquent, ces utilisateurs sont susceptibles de déployer des stratégies d'attaque très hostiles pour faire tomber cette protection gênante. Ils ne vont pas se contenter d'appliquer des traitements usuels tels que filtrage ou compression avec pertes dans l'espoir d'altérer le signal de tatouage. Ils vont plutôt essayer de rassembler le plus d'information possible sur le système de protection pour mettre au point de nouvelles attaques dédiées. C'est ce genre de comportement qui est d'intérêt lorsqu'on parle de sécurité en tatouage numérique. Ainsi, la notion de sécurité est intimement liée avec le besoin de confiance dans un environnement hostile. D'un côté, la pérennité du modèle économique nécessite de devoir faire confiance à l'information codée par le tatouage. De l'autre côté, les utilisateurs perçoivent la protection apportée par le tatouage comme une gêne et s'efforcent de casser la fiabilité du système par divers moyens.

Il faut néanmoins noter que de nombreuses applications du tatouage n'ont aucune spécification en termes de sécurité. C'est en particulier le cas pour les applications où le tatouage inséré ajoute un service supplémentaire (qualité supérieure, correction d'erreur, information d'indexation). Dans ce cas, les utilisateurs n'ont aucun intérêt à enlever le tatouage et il n'est pas nécessaire de s'inquiéter d'un potentiel comportement malicieux.

### A.2.2 Robustesse et Sécurité

Sécurité et robustesse demeurent encore de nos jours des concepts flous qu'il est difficile de distinguer dans le contexte du tatouage numérique. Il convient de donner quelques éléments simples qui permettent de faire la distinction entre ces deux notions. Le premier élément est sans aucun doute l'*environnement*. Comme cela a été mentionné dans la précédente sous-section, parler de sécurité revient à faire l'hypothèse implicite que le système de tatouage évolue dans un environnement hostile. La robustesse s'intéresse plutôt à la survie du tatouage lorsque les documents protégés sont soumis à des traitements courants. Ainsi, un utilisateur qui compresse avec pertes des documents tatoués ne peut pas être assimilé à une menace contre la sécurité du système, même si cette opération est susceptible d'altérer le signal de tatouage. Le point principal ici est que l'utilisateur n'a pas l'intention d'enlever le tatouage mais cherche juste à réduire la taille de ses données pour faciliter leur stockage/transmission. En d'autres termes, il utilise de façon aveugle des opérations existantes de traitement du signal. Ceci est radicalement différent d'un attaqueur hostile dont la stratégie est souvent divisée en deux étapes. Dans un premier temps, il va rassembler autant d'éléments d'information que possible sur le système de tatouage; dans un second temps, il va exploiter ce savoir pour mettre au point de nouvelles attaques dédiées qui mettront à mal le système. On peut donc dire que le *type de traitement* est un second élément de

distinction: générique pour la robustesse et spécialisé pour la sécurité. Enfin, le dernier point a trait à l'*impact des attaques*. Les opérations usuelles de traitement du signal sont seulement susceptibles d'empêcher le détecteur d'extraire le tatouage. En revanche, les attaques contre la sécurité du système peuvent aussi aboutir éventuellement à la détection non autorisée du tatouage, à son estimation, à sa modification ou bien même à l'insertion d'un nouveau tatouage dans un document non protégé.



**Figure A.1:** Classification robustesse/sécurité des attaques couramment utilisées en tatouage numérique.

Une fois ces trois différences établies, il peut être utile de lister les attaques couramment utilisées en tatouage numérique et de les ventiler entre robustesse et sécurité comme cela est illustré dans la Figure A.1. Cette distinction étend les classifications précédemment considérées en tatouage [151, 192]. Dans la partie gauche, les attaques relatives à la robustesse sont séparées en deux catégories. Les attaques synchrones incluent toutes les opérations usuelles telles que filtrage, compression avec pertes, quantification, débruitage qui modifient la valeur des échantillons du signal et qui sont donc susceptibles d'altérer le signal de tatouage. De l'autre côté, les attaques asynchrones regroupent tous les traitements qui modifient la position des échantillons. Par conséquent, la convention de synchronisation entre le tatoueur et le détecteur devient caduque. Ainsi, même si ces traitements ne suppriment pas effectivement le signal de tatouage, le détecteur n'est plus capable d'extraire le tatouage. Un exemple très connu en image fixe est l'attaque StirMark [151, 173] qui introduit localement des déplacements aléatoires de faible amplitude.


Du côté sécurité, une première catégorie d'attaque s'efforce de mettre en défaut le protocole autour du système de tatouage plutôt que d'attaquer directement le signal lui-même. Dans le cadre d'une application de protection des droits d'auteur, si un document contient deux tatouages, la plupart des algorithmes ne permettent pas de dire lequel a été inséré en premier. Il y a interblocage et

personne ne peut revendiquer la paternité du document [36]. Par ailleurs, les applications de suivi de copies exploitent des robots qui inspectent les sites Internet pour vérifier s'ils hébergent ou non illégalement des documents propriétaires. Une façon simple de faire échouer ces robots et de diviser les documents, par exemple une image, en morceaux et de juxtaposer ceux-ci lors de l'affichage. Si les morceaux sont assez petits, il est impossible de détecter le tatouage dans l'un d'entre eux [1]. La seconde catégorie d'attaques ayant trait à la sécurité vise à obtenir des renseignements sur le signal de tatouage lui-même, l'approche la plus simple (et le plus souvent coûteuse) étant de rechercher de façon exhaustive la clé qui a été utilisée. De son côté, la stéganalyse a pour objectif d'isoler les caractéristiques des algorithmes de tatouage afin d'obtenir un outil disant par exemple si un document est tatoué ou non [23]. Dans certaines applications, typiquement le contrôle de copie, le public a accès à un détecteur (oracle). Un attaquant peut alors considérer ce détecteur comme une boîte noire et modifier faiblement les données numériques de façon itérative jusqu'à ce que la copie soit autorisée [125]. Comme cela sera détaillé dans le reste de l'article, plusieurs documents peuvent être considérés et combinés pour mettre en défaut le système de tatouage. Enfin un exemple d'insertion non autorisée de tatouage est l'attaque par recopie [110, 83]: le tatouage est estimé à partir d'un document tatoué et réinséré dans un document non protégé.

### A.2.3 Attaques par collusion

La collusion est une stratégie d'attaque connue depuis un certain temps en cryptographie. Une clique d'utilisateurs malicieux se rassemble et met en commun ses informations/connaissances sur le système de protection, quelles qu'elles soient, pour générer des données non protégées. Ce type de comportement a été mentionné pour la première fois lorsque des protocoles ont été mis au point pour diviser un secret entre plusieurs individus sans qu'aucun d'entre eux n'ait accès à l'ensemble du secret [133]. Un exemple typique est le partage de secret pour contrôler des actions critiques telles que l'ouverture de la porte d'un coffre fort particulier à la banque. Le client et le responsable de la banque ont tous les deux une clé et les deux sont nécessaires pour ouvrir le coffre. Si une partie du secret (clé) manque, la porte du coffre reste fermée. à plus grande échelle, plusieurs clés contenant une partie du secret sont distribuées et il est nécessaire de rassembler au moins  $k$  clés différentes pour avoir accès à l'intégralité du secret. Dans ce contexte, les attaquants sont un groupe de  $u$  utilisateurs qui cherchent à construire de fausses clés ou à reconstruire l'intégralité du secret quand bien même  $u < k$ . On retrouve aussi cette problématique de la collusion dans des schémas de distribution dynamique de clés [67] pour les sessions de audio/vidéo conférences, vidéo à la demande, etc.

En tatouage numérique, les attaques par collusion ont été mentionnées pour la



**Figure A.2:** Collusion en tatouage numérique: Plusieurs utilisateurs rassemblent plusieurs documents tatoués et les combinent pour produire des documents ne contenant plus aucun tatouage.

première fois dans le contexte du suivi de copies [200]. Dans ce cas, les fournisseurs de contenus veulent distribuer un faible nombre de contenus à une très large audience. Ils désirent par conséquent avoir les moyens de pister une copie pirate jusqu'à la personne à l'origine de cette fuite. Dans ce but, au lieu de distribuer par exemple le même film à tous les consommateurs, des copies sensiblement différentes sont assignées à chacun d'entre eux. Ainsi, chaque consommateur a une copie unique portant son propre tatouage. Si un utilisateur isolé rend sa copie disponible sur Internet, on peut alors l'identifier en utilisant le tatouage. Par conséquent, les attaquants sont tentés de se regrouper pour combiner leur différentes copies afin de générer un nouveau document qui ne contiendrait plus de tatouage comme illustré dans la Figure A.2. Il existe principalement deux stratégies de collusion en tatouage:

1. soit les documents sont analysés pour estimer certaines propriétés du signal de tatouage qui pourraient être utilisé dans un second temps pour retirer le signal de tatouage,
2. soit les documents sont combinés pour estimer directement le document original non tatoué.

Des parades ont déjà été proposées dans la littérature. Par exemple, on peut exploiter des codes ayant certaines propriétés pour assurer que lorsque des documents tatoués sont combinés, certaines parties du tatouage restent intactes [16]. Ces parties résiduelles sont alors examinées pour isoler et identifier de façon certaine au moins un des individus dans la clique des attaquants.

### A.3 Collusion en vidéo

L'évaluation de la sécurité est devenue aujourd'hui une problématique majeure dans le domaine du tatouage numérique. Le but est d'anticiper les comportements hostiles des utilisateurs afin d'introduire à temps des parades appropriées. Dans ce contexte, les attaques par collusion doivent être considérées sérieusement. Dans ce type d'approche, plusieurs utilisateurs se rassemblent pour accumuler différents documents tatoués. Ils les combinent ensuite pour obtenir des documents qui ne contiennent plus aucun signal de tatouage. Le tatouage de vidéo numérique se résume souvent à des approches image par image [48] comme écrit ci-dessous:

$$\check{\mathbf{f}}_t = \mathbf{f}_t + \alpha \mathbf{w}_t, \quad \mathbf{w}_t \sim \mathcal{N}(0, 1) \quad (\text{A.1})$$

où  $\mathbf{f}_t$  est la trame vidéo originale à l'instant  $t$ ,  $\check{\mathbf{f}}_t$  sa version tatouée,  $\alpha$  la force de tatouage et  $\mathbf{w}_t$  le signal de tatouage qui est distribué suivant une loi gaussienne à moyenne nulle et à variance unité. Par conséquent, chaque trame peut être considérée comme un document tatoué individuellement [174, 54]. En d'autres termes, un attaquer *isolé* peut mettre au point une attaque par collusion en considérant les différentes trames d'une vidéo tatouée. Il n'est plus nécessaire de devoir de se rassembler à plusieurs. Les sous-sections qui suivent donnent un aperçu des différentes attaques par collusion possibles en vidéo lorsqu'on suit ce raisonnement.

#### A.3.1 Estimer une structure redondante

Lorsque les tatouages enfouis dans différentes trames ne sont pas complètement indépendants, une stratégie de collusion est d'examiner les différentes trames tatouées et d'identifier des structures suspectes statistiquement redondantes. Ces fuites d'information peuvent être vues comme une empreinte statistique déposée par l'algorithme de tatouage. Du point de vue d'un attaquer, la situation idéale serait d'avoir accès au canal de tatouage directement. Cependant, en pratique, les trames vidéo originales ne sont pas disponibles et il est impossible d'obtenir une estimation parfaite du tatouage en calculant  $\mathcal{E}_o(\check{\mathbf{f}}_t) = \check{\mathbf{f}}_t - \mathbf{f}_t$  pour chaque trame. à défaut, du fait de la nature habituellement haute fréquence du tatouage, une estimation approximative peut être obtenue par des méthodes de débruitage ou bien, plus simplement, en calculant la différence entre chaque trame tatouée et sa version filtrée passe-bas:

$$\mathcal{E}(\check{\mathbf{f}}_t) = \check{\mathbf{f}}_t - \mathcal{L}(\check{\mathbf{f}}_t) = \tilde{\mathbf{w}}_t \quad (\text{A.2})$$

où  $\tilde{\mathbf{w}}_t$  est l'estimation du tatouage enfoui à l'instant  $t$  et  $\mathcal{L}(\cdot)$  un filtre passe-bas comme par exemple un filtre moyenneur  $5 \times 5$ . Maintenant, ayant à sa disposition une collection de tatouages bruités, l'attaquer doit trouver une structure redondante secrète qui puisse être exploitée par la suite pour retirer le signal de tatouage.

### Estimer un unique tatouage

Pour s'affranchir simplement de la contrainte de synchronisation temporelle, une solution consiste à enfouir toujours le même tatouage de référence  $\mathbf{r}$  dans toutes les trames de la vidéo [93]. De plus, si l'algorithme de détection est linéaire, alors accumuler dans le temps des scores de détection calculés à différents instants est équivalent à faire une unique détection en utilisant l'accumulation temporelle des trames vidéo. En d'autres termes, il n'est pas nécessaire de lancer la procédure de détection pour chaque trame, ce qui peut être utile pour traiter la vidéo en temps réel. En revanche, d'un point de vue sécurité, toujours enfouir le même tatouage rend le signal de référence  $\mathbf{r}$  statistiquement visible. En effet, si chaque estimation  $\tilde{\mathbf{w}}_t$  est individuellement trop bruitée pour menacer la pérennité de l'algorithme de détection, les combiner permet de raffiner de façon significative l'estimation finale  $\tilde{\mathbf{r}}$  du tatouage. Une approche simple consiste par exemple à moyenner des différentes estimations comme suit :

$$\tilde{\mathbf{r}} = \frac{1}{T} \sum_t \tilde{\mathbf{w}}_t \quad (\text{A.3})$$


où  $T$  est le nombre de trames utilisées pour la collusion [82, 174]. Cette estimation peut alors être remodulée pour effacer de manière efficace le signal de tatouage dans chaque trame [191]. L'ensemble de cette attaque par Estimation du Tatouage et Remodulation (ETR) est illustré dans la Figure A.3. Il est important de noter que le processus de raffinement de l'estimation du tatouage est d'autant plus efficace que les trames vidéo utilisées pour la collusion sont différentes. Ainsi cette attaque par ETR est plus pertinente dans des scènes dynamiques ou lorsque des trames clefs de la séquence vidéo sont considérées. Par ailleurs, plus on combine un nombre important d'estimations individuelles, meilleure est l'estimation finale  $\tilde{\mathbf{r}}$  du tatouage. Ces deux remarques sont aussi valides pour les autres attaques présentées dans cette section.

### Estimer une collection de tatouages

Une parade immédiate face à la menace d'une attaque par ETR est d'utiliser plus qu'un seul et unique tatouage de référence. Dans cette perspective, pour chaque trame vidéo, le tatouage qui est enfoui va être choisi parmi une bibliothèque de  $N$  tatouages de référence  $\{\mathbf{r}_i\}$  comme noté ci-dessous :

$$\begin{aligned} \forall t \quad \mathbf{w}_t &= \mathbf{r}_{\Phi(t)} \\ \text{avec} \quad &\begin{cases} \mathbf{r}_i \cdot \mathbf{r}_j = \delta_i^j, & 1 \leq i, j \leq N \\ \text{P}(\Phi(t) = i) = 1/N, & 1 \leq i \leq N \end{cases} \end{aligned} \quad (\text{A.4})$$

où  $\cdot$  représente l'opérateur de corrélation linéaire et  $\delta$  le symbole de Kronecker. Cette stratégie de tatouage recouvre un grand nombre d'algorithmes en allant



**Figure A.3:** Attaque par Estimation du Tatouage et Remodulation (ETR): différentes estimations du tatouage obtenues avec différentes trames vidéo sont combinées pour raffiner l'estimation du tatouage de référence. Ensuite, cette estimation est remodulée pour enlever le signal de tatouage dans chaque trame vidéo.

d'une succession périodique des tatouages  $\mathbf{r}_i$  à une succession complètement aléatoire [122]. Comme les différents tatouages de référence sont émis de façon équiprobable, une attaque par ETR est vouée à l'échec. En effet, si un attaqueur moyenne plusieurs estimations  $\tilde{\mathbf{w}}_t$ , il obtient la moyenne des tatouages de référence et ce signal ne peut pas être utilisé pour retirer ensuite le signal de tatouage. Dans ce cas, le gain en termes de sécurité repose sur l'hypothèse qu'un attaqueur est incapable de construire des ensembles de trames vidéo portant le même tatouage de référence  $\mathbf{r}_i$ . Néanmoins, chaque estimation  $\tilde{\mathbf{w}}_t$  peut être considérée comme un vecteur dans un espace de grande dimension qui est censé approximer un des tatouages de référence. Par conséquent, une quantification vectorielle permet d'isoler  $N$  amas de vecteurs  $\mathcal{R}_i$  dont les centroïdes  $\tilde{\mathbf{r}}_i$  sont de bonnes estimations des tatouages de référence secrets. Cette approche peut être implantée de façon simple en utilisant un algorithme des  $k$ -moyennes et une stratégie de division-fusion pour éviter une initialisation aléatoire [53]. Une fois que les tatouages de référence ont été estimés, l'attaqueur teste chaque trame vidéo pour identifier quel tatouage est présent et effectue une remodulation pour l'enlever. On peut remarquer que la précédente attaque par ETR est un cas particulier de cette approche par Quantification des Estimations de Tatouage et Remodulation (QETR) pour  $N = 1$ .



### Estimer un sous espace de tatouage

Les attaques par ETR et QETR exploitent la même faille de sécurité pour vaincre les systèmes de tatouage. Lorsque les tatouages enfouis dans chaque trame sont vus comme des vecteurs dans un espace à grande dimension, les stratégies d'insertion vues précédemment introduisent des points d'accumulation dans l'espace qui peuvent facilement être identifiés de façon aveugle. Pour éviter ce piège, on peut enfouir dans chaque trame une combinaison de tatouages de référence:

$$\forall t \quad \mathbf{w}_t = \sum_{i=1}^N \frac{\lambda_i(t)}{\sqrt{\sum_{j=1}^N \lambda_j(t)^2}} \mathbf{r}_i \quad (\text{A.5})$$

où les  $\lambda_i(t)$  sont  $N$  coefficients de mixage variant dans le temps. Comme les tatouages de référence ont une variance unité, les tatouages successifs  $\mathbf{w}_t$  décrivent une trajectoire sur la sphère unité. Si cette trajectoire ne présente pas de points d'accumulation, alors une attaque QETR est vouée à l'échec. Cependant, une faiblesse subsiste du fait que le nombre  $N$  de tatouages considérés est habituellement largement inférieur à la dimension de l'espace  $D$  du média considéré ( $N \ll D$ ). En d'autres termes, le signal de tatouage est borné au sein d'un sous-espace de faible dimension  $\mathcal{R} = \text{vect}(\mathbf{r}_i)$ . Il est alors possible d'utiliser des techniques de réduction de dimensions telles que l'Analyse par Composantes Principales (ACP) pour obtenir une estimation  $\tilde{\mathcal{R}}$  de ce sous-espace en considérant les différentes estimations  $\tilde{\mathbf{w}}_t$ . Ensuite, pour chaque trame vidéo, supprimer l'énergie présente dans ce sous-espace permet de retirer le signal de tatouage [50]. Bien entendu, afin d'obtenir une bonne estimation du sous-espace de tatouage  $\mathcal{R}$ , il est nécessaire de prendre en compte un nombre d'estimations  $\tilde{\mathbf{w}}_t$  d'autant plus grand que sa dimension  $N$  est grande.

De récents travaux sont venus renforcer ces résultats expérimentaux en adoptant une démarche basée sur la théorie de l'information pour quantifier l'étendue des fuites d'information [20]. Dans ce but, l'ignorance à propos du système est mesurée en utilisant l'entropie conditionnelle:

$$H(K|\mathbf{d}_1, \dots, \mathbf{d}_T) = H(K) - I(K; \mathbf{d}_1, \dots, \mathbf{d}_T) \quad (\text{A.6})$$

où  $\{\mathbf{d}_i\}$  est un ensemble de documents tatoués et  $K$  le secret à estimer. Ainsi, les fuites d'information sont assimilées à l'information mutuelle entre les documents tatoués et le secret. Lorsque l'entropie conditionnelle  $H(K|\mathbf{d}_1, \dots, \mathbf{d}_T)$  tombe à zéro, la totalité du secret du système a été dévoilée.

#### A.3.2 Combiner différents tatouages

Si une structure redondante de tatouage est susceptible d'être facilement estimée par un attaqueur, insérer des tatouages complètement indépendants dans des


trames successives n'est pas non plus la solution. En effet, en s'appuyant sur le fait que la somme de plusieurs échantillons de tatouage indépendants est souvent égale à zéro, un attaquer peut alors mettre au point des attaques par collusion très efficaces. Désormais le but n'est plus d'identifier une structure secrète pour enlever ensuite le tatouage mais plutôt d'estimer directement le document original non tatoué. Bien sûr, pour des raisons de fidélité, ces documents doivent être assez similaires pour être combinés sans dégrader de façon perceptible le document considéré. Les contenus vidéo présentent assez de redondance pour mettre en oeuvre de telles stratégies d'attaque.

### Compensation de mouvement

L'une des toutes premières méthodes de tatouage vidéo considère le contenu vidéo comme un signal monodimensionnel et ajoute simplement un signal de tatouage pseudo aléatoire [79]. D'un point de vue image, cela revient à toujours enfouir un tatouage différent dans chaque trame vidéo. Dans une séquence vidéo avec peu de mouvement, les trames successives sont fortement corrélées et peuvent être moyennées dans le temps sans endommager de façon notable la qualité de la vidéo. Cependant, comme les tatouages successifs sont indépendants, cette opération de moyennage temporel diminue de façon très significative l'énergie du tatouage  $\mathbf{w}_t$  présente dans la trame vidéo  $\mathbf{f}_t$ . Cette stratégie doit être légèrement modifiée lorsque la séquence vidéo contient des éléments dynamiques tels que des mouvements de caméra et/ou des objets en mouvements. En particulier, le mouvement de la caméra doit être compensé afin de permettre un Moyennage Temporel après Recalage (MTR) [49]. Comme l'illustre la Figure A.4, cette attaque consiste à estimer l'arrière-plan de chaque trame en utilisant les trames voisines. Cela est possible car les trames successives d'une séquence vidéo sont différentes vues du même décor de cinéma, ou encore différentes projections 2D de la même scène 3D. Les objets en mouvement étant plus difficiles à estimer, ils sont conservés tels quels. Cette attaque peut aussi être vue comme un moyennage temporel suivant l'axe du mouvement. Quoi qu'il en soit, du fait que la plupart des algorithmes de tatouage ne prêtent pas attention à l'évolution de la structure de la scène pendant l'enfouissement du tatouage, le MTR parvient à éliminer le signal qui a été introduit. Enfin, il est utile de noter que l'utilisation de mosaïques vidéo pour compresser efficacement l'arrière-plan comme préconisé dans le standard MPEG-4 aurait un impact similaire au MTR sur le tatouage [105].

### Autosimilarités

S'il est aisé d'admettre qu'une séquence vidéo est redondante dans le temps, il est moins immédiat de remarquer que chaque trame vidéo présente aussi une certaine redondance spatiale. Ces autosimilarités ont déjà été utilisées pour concevoir des algorithmes de compression efficaces [68]. Ainsi, à la façon du codage




**Figure A.4:** Moyennage Temporel après Recalage (MTR): pour chaque trame vidéo: une fois les objets en mouvement isolés (a), les trames voisines sont recalées (b) et combinées pour obtenir une estimation du fond de la trame courante (c). Ensuite, les objets vidéo manquants sont réinsérés.

fractal, un attaquer peut mettre au point une Attaque par Remplacement de Bloc (ARB) comme illustré dans la Figure A.5 qui consiste à remplacer chaque bloc de l'image par un autre pris ailleurs dans l'image qui est similaire au bloc d'origine modulo une transformation géométrique et photométrique [158]. De façon différente, l'attaquer peut choisir de combiner plusieurs blocs de sorte que le bloc obtenu soit assez similaire pour être échangé sans menacer de détruire la qualité visuelle de l'image [102]. évidemment, il existe un compromis entre l'efficacité de l'attaque et son impact perceptuel. Plus (resp. moins) le bloc candidat au remplacement est similaire au bloc à remplacer, moins (resp. plus) l'attaque est susceptible d'être efficace. Ce constat a motivé l'introduction d'un schéma adaptatif pour combiner un nombre variable de blocs en fonction de la nature du bloc considéré [57]. Il est en effet nécessaire de combiner plus (resp. moins) de blocs pour approximer de façon satisfaisante un bloc texturé (resp. uni). Comme aujourd'hui les algorithmes de tatouage ignorent les autosimilarités du signal, les ARB parviennent la plupart du temps à altérer de façon critique le signal de tatouage.

## A.4 Tatouage cohérent avec le signal

D'un côté, une structure de tatouage redondante utilisée pour tatouer des documents différents peut être estimée. D'un autre côté, des tatouages indépendants insérés dans des documents (ou des parties de documents) similaires peuvent être effacés par simple moyennage. Ce constat conduit intuitivement à essayer de respecter une règle d'enfouissement qui assure que *les tatouages insérés dans deux documents sont aussi corrélés que les documents eux-mêmes*. Différentes approches ont déjà été proposées pour remplir ce cahier des charges: rendre le



**Figure A.5:** Attaque par Remplacement de Bloc (ARB): chaque bloc est remplacé par un bloc pris à une autre position qui lui est similaire à une transformation géométrique et photométrique près.

tatouage dépendant des trames vidéos [82], utiliser des signatures numériques binaires des trames vidéos pour générer des tatouages qui sont aussi corrélés que ces signatures [70, 41], enfouir le tatouage à des endroits dépendants du contenu des trames vidéos [174]. Néanmoins, aucune de ces solutions ne s'est révélée vraiment satisfaisante. En considérant plus particulièrement les faiblesses soulignées dans la sous-section A.3.2, on s'aperçoit que le signal de tatouage doit être cohérent avec le contenu de la séquence vidéo, cohérent avec le mouvement de la caméra d'une part (sous-section A.4.1) et cohérent avec les autosimilarités du signal d'autre part (sous-section A.4.2).


#### A.4.1 Gérer le mouvement de la caméra

Pour une scène vidéo donnée, l'arrière plan de trames successives peut être assimilé à différentes projections 2D d'un même décor 3D. Fondamentalement, le MTR exploite le fait que les algorithmes de tatouage ne prennent pas en compte le mouvement de la caméra. Par conséquent, un point du décor 3D qui est projeté à plusieurs endroits dans des trames vidéo différentes est associé à des échantillons de tatouage non corrélés. Ainsi, moyenniser les trames recalées permet d'enlever le tatouage. Une riposte possible à cette faiblesse est de renseigner le tatoueur en termes de mouvements de la caméra et définir une stratégie de tatouage qui force chaque point 3D du décor à toujours être associé avec le même échantillon de tatouage, où qu'il soit visible dans la scène vidéo. Comme illustré dans la Figure A.6, cette tactique peut être implantée en ayant recours aux mosaïques vidéo [51]. Pour chaque trame, des paramètres de recalage  $\theta_t$  sont calculés pour définir la position de la trame dans la mosaïque. Un tatouage de référence  $\mathbf{r}$  de la taille de la mosaïque est construit. La portion  $\mathbf{r}_t$  associée à chaque trame vidéo est récupérée et recalée pour obtenir le signal  $\mathbf{r}_t^{(\theta_t)}$  à enfouir dans chaque trame. Par ailleurs, les objets en mouvement ne sont pas tatoués pour suivre la philosophie: *un point 3D porte toujours le même échantillon de tatouage tout au*

long de la scène. L'ensemble du processus peut s'écrire:

$$\mathbf{w}_t = \mathbf{m}_t \otimes \mathbf{r}_t^{(\theta_t)} \quad (\text{A.7})$$

où  $\mathbf{m}_t$  est un masque binaire qui distingue les objets en mouvement de l'arrière-plan et  $\otimes$  représente la multiplication pixel à pixel. De son côté, le détecteur vérifie simplement si la portion de tatouage  $\mathbf{r}_t^{(\theta_t)}$  est effectivement présente dans chaque trame ou pas. Des précédents travaux en mosaïque vidéo [143, 142] ont été exploités pour implanter cette stratégie de tatouage et démontrer sa supériorité en termes de résistance au MTR [51].



**Figure A.6:** Tatouage cohérent avec le mouvement de la caméra: la partie du signal de tatouage qui est associée avec la trame vidéo courante est identifiée et recalée. Ensuite, elle est enfouie dans l'arrière-plan de la trame vidéo.

De plus, le fait de gérer le mouvement de la caméra au moment de l'insertion a aussi donné des résultats intéressants vis à vis de l'invisibilité du tatouage. évaluer l'impact de la distorsion induite par le tatouage comme perçue par un utilisateur humain est toujours et encore un grand défi en vidéo. Le comité VQEG [188] qui est chargé de définir des méthodes pour évaluer la qualité visuelle d'une vidéo a statué en 1999 qu'aucune des métriques testées n'était meilleure que les autres dans tous les cas, et surtout qu'aucune d'entre elles pouvait remplacer une inspection visuelle subjective. C'est la raison pour laquelle l'évaluation de la visibilité d'un tatouage se résume souvent à une recherche manuelle d'artéfacts visuels. Deux principaux défauts ont été isolés en vidéo [131, 197]:

1. *Le scintillement:* Enfouir des tatouages indépendants dans des trames vidéo successives introduit souvent un scintillement désagréable, comme un bruit de capteur;
2. *La persistance:* Enfouir le même tatouage de référence dans toutes les trames de la vidéo produit un motif fixe dérangerant visuellement parlant, comme si la scène avait été filmée par une caméra ayant des lentilles sales.

Si la stratégie de compensation de mouvement proposée est mise en oeuvre, les tatouages insérés dans chaque trame sont différents. Néanmoins, ils n'introduisent plus de scintillement car le tatouage est cohérent avec le mouvement de la caméra. En fait, si le tatouage est assez amplifié pour être visible, on a la sensation que la caméra filme une scène qui est bruitée. En d'autres termes, en plus d'assurer une meilleure sécurité, la stratégie proposée simule un monde utopique où le décor de la vidéo serait tatoué avant d'être filmé et cela dérange beaucoup moins le système visuel humain.

#### A.4.2 Hériter des autosimilarités

Si la compensation de mouvement permet de rendre le tatouage cohérent avec la redondance temporelle du signal vidéo, il ne résout pas le problème des ARB. Ces attaques profitent du fait que les algorithmes de tatouage ne tiennent pas compte des autosimilarités du signal. Par conséquent, des blocs similaires (modulo une transformation géométrique et photométrique ou une combinaison linéaire) ne sont pas tatoués de façon similaire. Intuitivement, s'il était possible d'assurer que des blocs similaires portent des tatouages similaires, les ARB devraient être inefficaces. Formulé d'une autre manière, le but est donc d'imposer que *des pixels ayant des voisinages similaires portent des échantillons de tatouage ayant des valeurs proches* i.e. de faire hériter le signal de tatouage des autosimilarités du signal porteur [55]. En admettant qu'il est possible de définir le voisinage d'un pixel à la position  $\mathbf{p}$  dans une trame  $\mathbf{f}$  à l'aide d'un vecteur caractéristique  $\mathbf{v}(\mathbf{f}, \mathbf{p})$ , cela revient à écrire:

$$\mathbf{v}(\mathbf{f}, \mathbf{p}_o) \approx \sum_k \lambda_k \mathbf{v}(\mathbf{f}, \mathbf{p}_k) \Rightarrow w(\mathbf{f}, \mathbf{p}_o) \approx \sum_k \lambda_k w(\mathbf{f}, \mathbf{p}_k) \quad (\text{A.8})$$

où  $w(\mathbf{f}, \mathbf{p})$  est la valeur de tatouage insérée à la position  $\mathbf{p}$  dans la trame  $\mathbf{f}$ . Pour obtenir cette propriété, il suffit de définir la fonction de tatouage  $w(\cdot)$  comme étant une forme linéaire  $\varphi(\cdot)$  dans l'espace  $\mathcal{V}$  des vecteurs caractéristiques. Cette forme linéaire est complètement définie par les valeurs  $w_i$  qu'elle prend sur une base orthonormée. C'est là que peut être injecté du secret dans le système en utilisant la clé secrète pour générer ces valeurs qui déterminent la forme linéaire. Une implantation de cette riposte exploitant des ondelettes de Gabor pour caractériser le voisinage en chaque pixel a montré de bonnes performances vis à vis des ARB [55].

Le fait d'utiliser des filtres de Gabor a permis d'établir un lien intéressant avec des algorithmes de tatouage existants qui enfouissent un signal pseudo-aléatoire de façon multiplicative dans un domaine fréquentiel. En effet, lorsque le tatouage est exprimé dans le domaine de Fourier, on obtient la relation suivante [56]:

$$\mathbf{W} = \mathbf{H}(K)\mathbf{I}, \quad \text{avec } \mathbf{H}(K) = \sum_{i=1}^N w_i \mathbf{H}_i \quad (\text{A.9})$$

où  $\mathbf{W}$  (resp.  $\mathbf{I}$ ) est la transformée de Fourier du tatouage (resp. de l'image) et  $\mathbf{H}_i$  est un des filtres de Gabor utilisés pour caractériser le voisinage. Il est à noter que des filtres symétriques par rapport à l'origine dans le domaine fréquentiel ont été utilisés afin d'obtenir un vecteur caractéristique  $\mathbf{v}(\mathbf{f}, \mathbf{p})$  à valeur réelles. Si on fait tendre la bande passante d'un filtre de Gabor vers 0, on obtient alors deux pics de Dirac symétriques par rapport en centre. Et dans ce cas, le schéma proposé revient à un enfouissement multiplicatif dans le domaine de Fourier [7]. En d'autres termes, un tatouage obtenu en multipliant le spectre de l'image avec un signal pseudo-aléatoire symétrique par rapport au centre présente les mêmes autosimilarités que le signal hôte. De même, il est possible de montrer qu'une multiplication dans le domaine DCT [31] produit aussi un tatouage qui a hérité des autosimilarités du signal. Ces deux points ont d'ailleurs été vérifiés expérimentalement en vérifiant la résistance de tels tatouages face aux ARB [56].

Un point intéressant à ce moment est de se rappeler que ces schémas multiplicatifs ont été considérés pour des raisons de masquage perceptuel dû au contraste: des coefficients de fortes valeurs peuvent transporter des valeurs de tatouage plus grandes sans compromettre l'invisibilité du tatouage [69]. D'un autre côté, en utilisant la linéarité de la fonction de tatouage du système qui a été proposé, on peut écrire la relation suivante:

$$\mathbf{w}(\mathbf{f}, \mathbf{p}) = \|\mathbf{v}(\mathbf{f}, \mathbf{p})\| \varphi \left( \frac{\mathbf{v}(\mathbf{f}, \mathbf{p})}{\|\mathbf{v}(\mathbf{f}, \mathbf{p})\|} \right) \quad (\text{A.10})$$

Le vecteur  $\mathbf{u}(\mathbf{f}, \mathbf{p}) = \mathbf{v}(\mathbf{f}, \mathbf{p})/\|\mathbf{v}(\mathbf{f}, \mathbf{p})\|$  qui est passé en argument de la forme linéaire  $\varphi(\cdot)$  est sur la sphère unité. Sous certaines hypothèses, il est possible de montrer que  $\varphi(\mathbf{u}(\mathbf{f}, \mathbf{p}))$  suit une distribution gaussienne de moyenne nulle et de variance unité [47]. En d'autres termes, l'échantillon de tatouage est amplifié ou atténué en fonction de la valeur de la norme du vecteur caractéristique  $\mathbf{v}(\mathbf{f}, \mathbf{p})$ . C'est là encore une technique couramment utilisée pour mettre en forme le tatouage afin de réduire son impact perceptuel [189].

## A.5 Conclusion

Longtemps négligée, la sécurité est devenue récemment une problématique majeure dans le domaine du tatouage numérique. Cela est lié au fait que la plupart des applications visées, comme la protection des droits d'auteur ou le suivi de copie, sont vouées à être déployées dans des environnements hostiles, c'est à dire où des attaquants malicieux s'attaquent délibérément au système. Dans cet article, deux principales stratégies de collusion ont été passées en revue: soit des documents différents ont été tatoués avec la même structure de tatouage et le but est d'estimer cette structure de tatouage; soit le même document a été tatoué de différentes façons et le but est d'approximer directement le document original. Ces attaques qui combinent différents documents tatoués sont d'autant plus

critiques en vidéo que chaque trame peut être vue comme un document tatoué distinct. Deux ripostes ont donc été introduites pour rendre le tatouage cohérent avec le signal vidéo. L'une s'appuie sur la compensation de mouvement pour tenir compte de la redondance temporelle, l'autre considère les autosimilarités pour gérer la redondance spatiale.

Les implantations proposées ne sont pas optimales et peuvent être améliorées. Néanmoins, cette démarche *sécurité face aux attaques par collusion* a permis de jeter un éclairage original sur le tatouage vidéo. Tout d'abord, l'importance de tenir compte du signal porteur a été soulignée lorsque le tatouage doit résister à des attaques hostiles type collusion. En particulier, l'utilisation de traitements vidéo tels que les mosaïques, la segmentation d'objets s'est avérée une approche pertinente pour générer des tatouages plus performants. De plus, le cheminement pour obtenir un tatouage cohérent avec le signal a permis de donner un regain d'intérêt pour d'anciens schémas de tatouage multiplicatifs dans le domaine fréquentiel. Enfin, quand bien même la progression de cette étude était guidée par une recherche de sécurité accrue, des résultats très intéressants ont été obtenus en termes d'invisibilité du tatouage inséré. En vue de mettre au point une nouvelle génération de tatoueurs vidéo, il pourrait être utile de réfléchir comment les outils usuels en vidéo peuvent être exploités dans le cadre du tatouage. De même, il pourrait être fructueux de penser comment intégrer une méthode de tatouage dans un système vidéo complet de sorte que le signal inséré n'interfère pas avec les fonctionnalités d'indexation/compression.



# B

---

## Euremark

---

Eurécom watermarking algorithm [61, 62] exploits fractal image coding theory [68] and in particular the notion of self-similarities. The image is considered as a collection of local similarities modulo an affine photometric compensation and a pool of geometric transformations. The underlying idea is then to use invariance properties of fractal coding such as invariance to affine photometric transformations to ensure watermark robustness. Furthermore, the extraction process is performed in a blind fashion i.e. the original image is not required.

### B.1 Watermark embedding

The embedding process can be divided in three different steps. First, a *fractal approximation*  $\mathbf{i}_o^{\text{IFS}}$  of the original image  $\mathbf{i}_o$  is computed. Second, the payload is properly formatted and encrypted to obtain the watermark  $\mathbf{w}$  to be embedded. Finally, the watermark is merged with the cover  $\mathbf{i}_o^c = \mathbf{i}_o - \mathbf{i}_o^{\text{IFS}}$  according to a sign rule.

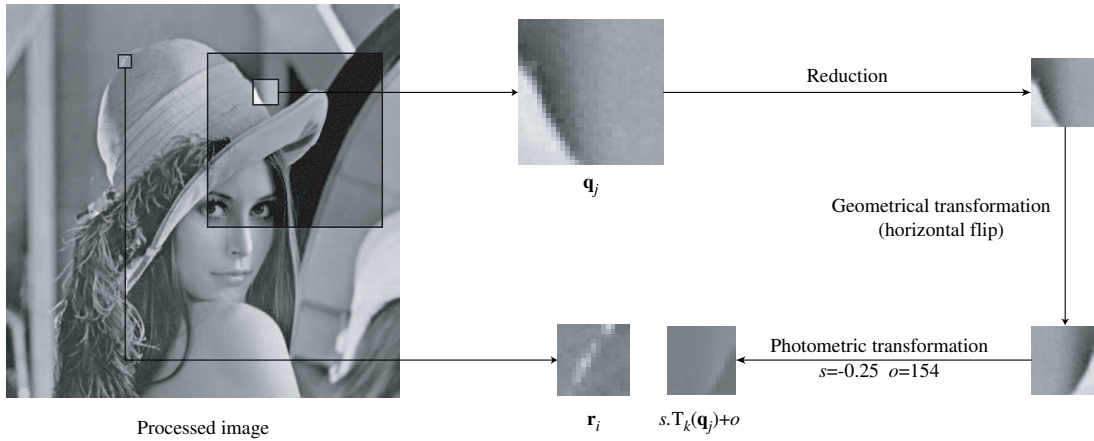
#### B.1.1 Cover generation

The input image is scanned block by block. Those blocks  $\mathbf{r}_i$  are labeled as *range blocks* and have a dimension  $r \times r$  e.g.  $8 \times 8$  pixels. The goal is then to find for each block a *domain block*  $\mathbf{d}_i$  taken from a pool of blocks which is similar according to the Mean Square Error (MSE) criterion defined below:

$$\text{MSE}(\mathbf{r}, \mathbf{d}) = \frac{1}{r^2} \sum_{\mathbf{p} \in [1,r] \times [1,r]} (\mathbf{r}(\mathbf{p}) - \mathbf{d}(\mathbf{p}))^2 \quad (\text{B.1})$$

where  $\mathbf{p}$  is a bi-dimensional spatial index used to address the pixels of the blocks  $\mathbf{r}$  and  $\mathbf{d}$ . By analogy with fractal image coding theory, for each range block, a search window is defined and the blocks  $\mathbf{q}_j$  lying in it are collected to initialize a codebook. Each block is then scaled to match the dimensions  $r \times r$  of the range blocks. Next, the codebook is enlarged by building  $k$  geometrically transformed blocks  $T_k(\mathbf{q}_j)$  e.g. identity, 4 flips and 3 rotations. An affine photometric compensation is then performed for each transformed block to minimize the Mean Square Error with the range block  $\mathbf{r}_i$  i.e a photometric scaling  $s$  and an offset  $o$  are computed to minimize  $\text{MSE}(s.T_k(\mathbf{q}_j) + o, \mathbf{r}_i)$ . Finally, the range block  $\mathbf{r}_i$  is substituted by the transformed block  $s.T_k(\mathbf{q}_j) + o$  which has the lowest MSE. The whole matching process is depicted in Figure B.1. The cover  $\mathbf{i}_o^c$  is simply obtained by computing the signed difference between the original image and its fractal approximation:

$$\mathbf{i}_o^c = \mathbf{i}_o - \mathbf{i}_o^{\text{IFS}} \quad (\text{B.2})$$



**Figure B.1:** Self-similarities: an example of association between range and domain blocks modulo an affine photometric compensation and a pool of geometric transformations.

### B.1.2 Watermark formatting

The payload to be hidden (a string or a logo) is first converted into a binary mark<sup>1</sup>. Then it is duplicated to ensure robustness against small modifications of the cover. On one hand, the binary mark is over-sampled by a scaling factor to produce a low-frequency watermark more resilient to low-pass filtering and lossy

<sup>1</sup>An Error Correcting Code (ECC), typically a block turbo code [157], can be inserted before any other formatting to further improve robustness against photometric attacks.

compression. On the other hand, this over-sampled mark is tiled horizontally and vertically up to the size of the image. This spatial repetition enables to compensate loss of information due to local image manipulations. At this point, the final binary watermark  $\mathbf{w}$  is obtained by encrypting the over-sampled tiled binary mark with a binary over-sampled pseudo-random sequence using a XOR operator. The XOR operation removes repetitive patterns and thus reduces the psycho-visual impact of the watermark. Nevertheless, using an over-sampled sequence permits to retain the low-frequency nature of the encrypted binary mark. Additionally, the XOR operation secures the hidden payload, typically against collusion attacks.

### B.1.3 Modulation

Modulating the watermark  $\mathbf{w}$  with the cover  $\mathbf{i}_o^c$  basically consists in zeroing some cover samples depending on their sign and the corresponding watermark bit to hide. More formally the following rules are applied:

$$\mathbf{i}_o^w(\mathbf{p}) = \begin{cases} \mathbf{i}_o^c(\mathbf{p}), & \text{if } \mathbf{w}(\mathbf{p}) = 1 \text{ and } \mathbf{i}_o^c(\mathbf{p}) > 0 \\ & \text{or } \mathbf{w}(\mathbf{p}) = 0 \text{ and } \mathbf{i}_o^c(\mathbf{p}) < 0 \\ 0, & \text{otherwise} \end{cases} \quad (\text{B.3})$$

where  $\mathbf{i}_o^w$  is the watermarked cover. It should be noted that, in average, only one pixel out of two is modified. Furthermore, for visibility reasons, high valued samples should not be zeroed. A threshold  $\tau_{\text{high}}$  is consequently introduced to discard systematically high valued samples as follows:

$$\mathbf{i}_o^w(\mathbf{p}) = \mathbf{i}_o^c(\mathbf{p}) \text{ if } |\mathbf{i}_o^c(\mathbf{p})| > \tau_{\text{high}} \quad (\text{B.4})$$

Finally, the watermarked cover is added to the fractal approximation to produce the watermarked image  $\mathbf{i}_w = \mathbf{i}_o^{\text{IFS}} + \mathbf{i}_o^w$ . By default, the threshold  $\tau_{\text{high}}$  is chosen so that the embedding process results in a distortion of 38 dB in terms of Peak Signal to Noise Ratio (PSNR).

## B.2 Watermark extraction

The extraction process is somewhat dual to the embedding. In a first step, a fractal approximation is computed. Then the embedded payload is retrieved according to some extraction rules and a detection score is computed.

### B.2.1 Cover extraction

As during the embedding process, a fractal approximation  $\mathbf{i}_w^{\text{IFS}}$  of the watermarked image is computed and the associated cover  $\mathbf{i}_w^c = \mathbf{i}_w - \mathbf{i}_w^{\text{IFS}}$  is extracted. A basic

assumption is that fractal coding is stable enough so that  $\mathbf{i}_w^{\text{IFS}} \approx \mathbf{i}_o^{\text{IFS}}$  and thus  $\mathbf{i}_w^c \approx \mathbf{i}_o^w$ . This cover is then decoded according to the following rule to obtain a ternary watermark  $\tilde{\mathbf{w}}$ :

$$\tilde{\mathbf{w}}(\mathbf{p}) = \begin{cases} 1, & \text{if } \tau_{\text{low}} < \mathbf{i}_w^c(\mathbf{p}) < \tau_{\text{high}} \\ -1, & \text{if } -\tau_{\text{high}} < \mathbf{i}_w^c(\mathbf{p}) < -\tau_{\text{low}} \\ 0, & \text{otherwise} \end{cases} \quad (\text{B.5})$$

Only samples whose magnitude is between the thresholds  $\tau_{\text{low}}$  and  $\tau_{\text{high}}$  are considered as carrying information related to the watermark. High valued samples are discarded since they are likely not to have been considered for watermarking during the embedding process. Furthermore, low valued samples are neglected since they might result from the non perfect cover stability ( $\mathbf{i}_w^c \neq \mathbf{i}_o^w$ ).

## B.2.2 Payload extraction

The binary pseudo-random sequence used during embedding is regenerated using the shared secret key. Its values are then mapped from  $\{0,1\}$  to  $\{1,-1\}$  and the resulting antipodal binary sequence is multiplied with the ternary watermark  $\tilde{\mathbf{w}}$  to invert the XOR operation performed during embedding. Next, the following quantities are computed for each payload bit:

$$d_k = \sum_{\mathbf{p} \in \mathcal{R}_k} \tilde{\mathbf{w}}(\mathbf{p}) \quad \text{and} \quad s_k = \sum_{\mathbf{p} \in \mathcal{R}_k} |\tilde{\mathbf{w}}(\mathbf{p})| \quad (\text{B.6})$$

where  $\mathcal{R}_k$  is the set of positions where the  $k^{\text{th}}$  bit has been duplicated. The value  $s_k$  indicates how many positions have been considered as carrying information related to the watermark and  $d_k$  the difference between position voting for 1 and those voting for 0. The final value of the  $k^{\text{th}}$  payload bit  $b_k$  can then be determined with a simple *majority vote* as follows:

$$b_k = \begin{cases} 0, & \text{if } d_k < 0 \\ 1, & \text{if } d_k \geq 0 \end{cases} \quad (\text{B.7})$$

Right now, whatever image is given in input, a sequence of bit is extracted. The following score is consequently computed:

$$\rho = \frac{\sum_{k=1}^K |d_k|}{\sum_{k=1}^K s_k} \quad (\text{B.8})$$

where  $K$  is the number of payload bits. When all the positions associated with a given bit are voting for the same bit value (watermarked image),  $d_k = \pm s_k$  and  $\rho = 1$ . On the contrary, if the positions vote evenly for 0 and 1 (non watermarked image), then  $d_k = 0$  and  $\rho = 0$ . As a result, the detection score  $\rho$  can be compared to a threshold  $\tau_{\text{detect}}$  to assert whether a watermark has been effectively embedded or not.

# BIBLIOGRAPHY

- [1] 2Mosaic. <http://www.petitcolas.net/fabien/watermarking/2mosaic>.
- [2] M. Alghoniemy and A. Tewfik. Geometric distortion correction through image normalization. In *Proceedings of the IEEE International Conference on Multimedia and Expo*, volume III, pages 1291–1294, August 2000.
- [3] M. Barni. What is the future for watermarking? (part I). *IEEE Signal Processing Magazine*, 20(5):55–59, September 2003.
- [4] M. Barni. What is the future for watermarking? (part II). *IEEE Signal Processing Magazine*, 20(6):53–57, November 2003.
- [5] M. Barni and F. Bartolini. *Watermarking Systems Engineering – Enabling Digital Assets Security and Other Applications*. Marcel Dekker, Inc., 2004.
- [6] M. Barni, F. Bartolini, R. Caldelli, A. De Rosa, and A. Piva. A robust watermarking approach for raw video. In *Proceedings of the 10th International Packet Video Workshop*, May 2000.
- [7] M. Barni, F. Bartolini, A. De Rosa, and A. Piva. A new decoder for optimum recovery of nonadditive watermarks. *IEEE Transactions on Image Processing*, 10(5):755–766, May 2001.
- [8] M. Barni, F. Bartolini, and T. Furon. A general framework for robust watermarking security. *Signal Processing, Special Section on Security of Data Hiding Technologies*, 83(10):2069–2084, October 2003.
- [9] F. Bartolini, A. Manetti, A. Piva, and M. Barni. A data hiding approach for correcting errors in H.263 video transmitted over a noisy channel. In *Proceedings of the IEEE Fourth Workshop on Multimedia Signal Processing*, pages 65–70, October 2001.
- [10] F. Bartolini, A. Tefas, M. Barni, and I. Pitas. Image authentication techniques for surveillance applications. *Proceedings of the IEEE*, 89(10):1403–1418, October 2001.

- [11] P. Bas and B. Macq. A new video-object watermarking scheme robust to object manipulation. In *Proceedings of the IEEE International Conference on Image Processing*, volume II, pages 526–529, October 2001.
- [12] J. Bloom, I. Cox, T. Kalker, J.-P. Linnartz, M. Miller, and C. Traw. Copy protection for DVD video. *Proceedings of the IEEE*, 87(7):1267–1276, July 1999.
- [13] Y. Bodo, N. Laurent, and J.-L. Dugelay. Watermarking video, hierarchical embedding in motion vectors. In *Proceedings of the IEEE International Conference on Image Processing*, volume II, pages 739–742, September 2003.
- [14] Y. Bodo, N. Laurent, and J.-L. Dugelay. A comparative study of different modes of perturbation for video watermarking based on motion vectors. In *Proceedings of the 12th European Signal Processing Conference*, pages 1501–1504, September 2004.
- [15] Y. Bodo, N. Laurent, C. Laurent, and J.-L. Dugelay. Video water-scrambling: Towards a video protection scheme based on the disturbance of motion vectors. *EURASIP Journal on Applied Signal Processing*, 2004(14):2224–2237, October 2004.
- [16] D. Boneh and J. Shaw. Collusion secure fingerprinting for digital data. *IEEE Transaction on Information Theory*, 44(5):1897–1905, September 1998.
- [17] I. Brown, C. Perkins, and J. Crowcroft. Watercasting: Distributed watermarking of multicast media. In *Proceedings of the First International Workshop on Networked Group Communication*, volume 1736 of *Lecture Notes in Computer Science*, pages 286–300, November 1999.
- [18] R. Caldelli, A. De Rosa, R. Becarelli, and M. Barni. Coping with local geometric attacks by means of optic-flow-based resynchronization for robust watermarking. In *Security, Steganography and Watermarking of Multimedia Contents VII*, volume 5681 of *Proceedings of SPIE*, pages 164–174, January 2005.
- [19] M. Carli, D. Bailey, M. Farias, and S. Mitra. Error control and concealment for video transmission using data hiding. In *Proceedings of the IEEE Fifth International Symposium on Wireless Personal Multimedia Communications*, volume II, pages 812–815, October 2002.
- [20] F. Cayre, C. Fontaine, and T. Furon. Watermarking security, part I: Theory. In *Security, Steganography and Watermarking of Multimedia Contents VII*, volume 5681 of *Proceedings of SPIE*, pages 746–757, January 2005.

- [21] F. Cayre, C. Fontaine, and T. Furon. Watermarking security, part II: Practice. In *Security, Steganography and Watermarking of Multimedia Contents VII*, volume 5681 of *Proceedings of SPIE*, pages 758–768, January 2005.
- [22] Certimark. <http://www.certimark.org>.
- [23] R. Chandramouli, M. Kharrazi, and N. Memon. Image steganography and steganalysis: Concepts and practice. In *Proceedings of the Second International Workshop on Digital Watermarking*, volume 2939 of *Lecture Notes in Computer Science*, pages 35–49, March 2004.
- [24] Checkmark. <http://watermarking.unige.ch/checkmark>.
- [25] B. Chen and G. Wornell. Quantization index modulation: A class of provably good methods for digital watermarking and information embedding. *IEEE Transactions on Information Theory*, 47(4):1423–1443, May 2001.
- [26] Q. Cheng and T. Huang. Robust optimum detection of transform domain multiplicative watermarks. *IEEE Transactions on Signal Processing*, 51(4):906–924, April 2003.
- [27] B. Chor, A. Fiat, and M. Naor. Tracing traitors. In *Proceedings of the 14th Annual International Cryptology Conference on Advances in Cryptology*, volume 839 of *Lecture Notes in Computer Science*, pages 257–270, August 1994.
- [28] B. Chor, A. Fiat, M. Naor, and B. Pinkas. Tracing traitors. *IEEE Transactions on Information Theory*, 46(3):893–910, May 2000.
- [29] H. Cohen. *A Course in Computational Algebraic Number Theory*. Springer-Verlag, 1993.
- [30] M. Costa. Writing on dirty paper. *IEEE Transactions on Information Theory*, 29(3):439–441, May 1983.
- [31] I. Cox, J. Kilian, T. Leighton, and T. Shamoan. Secure spread spectrum watermarking for multimedia. *IEEE Transactions on Image Processing*, 6(12):1673–1687, December 1997.
- [32] I. Cox and M. Miller. A review of watermarking and the importance of perceptual modeling. In *Human Vision and Electronic Imaging II*, volume 3016 of *Proceedings of SPIE*, pages 92–99, February 1997.
- [33] I. Cox and M. Miller. Preprocessing media to facilitate later insertion of a watermark. In *Proceedings of the International Conference on Digital Signal Processing*, volume I, pages 67–70, July 2002.

- [34] I. Cox and M. Miller. Facilitating watermark insertion by preprocessing media. *EURASIP Journal on Applied Signal Processing*, 2004(14):2081–2092, October 2004.
- [35] I. Cox, M. Miller, and J. Bloom. *Digital Watermarking*. Morgan Kaufmann Publishers, 2001.
- [36] S. Craver, N. Memon, B.-L. Yeo, and M. Yeung. Resolving rightful ownerships with invisible watermarking techniques: Limitations, attacks, and implications. *Journal on Selected Areas in Communications*, 16(4):573–586, May 1998.
- [37] J. Daugman. Complete discrete 2-D Gabor transforms by neural network for image analysis and compression. *IEEE Transactions on Acoustics, Speech and Signal Processing*, 36(7):1169–1179, July 1988.
- [38] F. Deguillaume, G. Csurka, and T. Pun. Countermeasures for unintentional and intentional video watermarking attacks. In *Security and Watermarking of Multimedia Contents II*, volume 3971 of *Proceedings of SPIE*, pages 346–357, January 2000.
- [39] F. Deguillaume, G. Csurka, J. Ó Ruanaidh, and T. Pun. Robust 3D DFT video watermarking. In *Security and Watermarking of Multimedia Contents*, volume 3657 of *Proceedings of SPIE*, pages 113–124, January 1999.
- [40] D. Delannay, J.-F. Delaigle, B. Macq, and M. Barlaud. Compensation of geometrical transformations for watermark extraction in the digital cinema application. In *Security and Watermarking of Multimedia Contents III*, volume 4314 of *Proceedings of SPIE*, pages 149–157, January 2001.
- [41] D. Delannay and B. Macq. A method for hiding synchronization marks in scale and rotation resilient watermarking schemes. In *Security and Watermarking of Multimedia Contents IV*, volume 4675 of *Proceedings of SPIE*, pages 548–554, January 2002.
- [42] D. Delannay, I. Setyawan, R. Lagendijk, and B. Macq. Relevant modeling and comparison of geometric distortions in watermarking systems. In *Application of Digital Image Processing XXV*, volume 4790 of *Proceedings of SPIE*, pages 200–210, July 2002.
- [43] G. Depovere, T. Kalker, J. Haitzma, M. Maes, L. De Strycker, P. Termont, J. Vandewege, A. Langell, C. Alm, P. Normann, G. O’Reilly, B. Howes, H. Vaanholt, R. Hintzen, P. Donnelly, and A. Hudson. The VIVA project: Digital watermarking for broadcast monitoring. In *Proceedings of the IEEE International Conference on Image Processing*, volume II, pages 202–205, October 1999.



- [44] J. Dittmann, M. Stabenau, and R. Steinmetz. Robust MPEG video watermarking technologies. In *Proceedings of ACM Multimedia*, pages 71–80, September 1998.
- [45] J. Dittmann, M. Steinebach, I. Rimac, S. Fisher, and R. Steinmetz. Combined audio and video watermarking: Embedding content information in multimedia data. In *Security and Watermarking of Multimedia Contents II*, volume 3971 of *Proceedings of SPIE*, pages 455–464, January 2000.
- [46] J. Dittmann, A. Steinmetz, and R. Steinmetz. Content-based digital signature for motion pictures authentication and content fragile watermarking. In *Proceedings of the IEEE International Conference on Multimedia Computing and Systems*, volume II, pages 209–213, June 1999.
- [47] G. Doërr. *Security Issue and Collusion Attacks in Video Watermarking*. PhD thesis, Université de Nice Sophia-Antipolis, France, June 2005.
- [48] G. Doërr and J.-L. Dugelay. A guide tour of video watermarking. *Signal Processing: Image Communication, Special Issue on Technologies for Image Security*, 18(4):263–282, April 2003.
- [49] G. Doërr and J.-L. Dugelay. New intra-video collusion attack using mosaicing. In *Proceedings of the IEEE International Conference on Multimedia and Expo*, volume II, pages 505–508, July 2003.
- [50] G. Doërr and J.-L. Dugelay. Danger of low-dimensional watermarking subspaces. In *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume III, pages 93–96, May 2004.
- [51] G. Doërr and J.-L. Dugelay. Secure background watermarking based on video mosaicing. In *Security, Steganography and Watermarking of Multimedia Contents VI*, volume 5306 of *Proceedings of SPIE*, pages 304–314, January 2004.
- [52] G. Doërr and J.-L. Dugelay. Secure video watermarking via embedding strength modulation. In *Proceedings of the Second International Workshop on Digital Watermarking*, volume 2939 of *Lecture Notes in Computer Science*, pages 340–354, March 2004.
- [53] G. Doërr and J.-L. Dugelay. Security pitfalls of frame-by-frame approaches to video watermarking. *IEEE Transactions on Signal Processing, Supplement on Secure Media*, 52(10):2955–2964, October 2004.
- [54] G. Doërr and J.-L. Dugelay. Collusion issue in video watermarking. In *Security, Steganography and Watermarking of Multimedia Contents VII*, volume 5681 of *Proceedings of SPIE*, pages 685–696, January 2005.

- [55] G. Doërr and J.-L. Dugelay. A countermeasure to resist block replacement attacks. In *Accepted for publication in the IEEE International Conference on Image Processing*, September 2005.
- [56] G. Doërr and J.-L. Dugelay. How to combat block replacement attacks? In *Pre-Proceedings of the 7th Information Hiding Workshop*, pages 137–151, June 2005.
- [57] G. Doërr, J.-L. Dugelay, and L. Grangé. Exploiting self-similarities to defeat digital watermarking systems - a case study on still images. In *Proceedings of the ACM Multimedia and Security Workshop*, pages 133–142, September 2004.
- [58] G. Doërr, C. Rey, and J.-L. Dugelay. Watermark resynchronization based on elastic graph matching. In *Proceedings of the International Conference on Sciences of Electronic, Technologies of Information and Telecommunications*, March 2005.
- [59] G. Donato, M. Bartlett, J. Hager, P. Ekman, and T. Sejnowski. Classifying facial actions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(10):974–989, October 1999.
- [60] B. Duc, S. Fisher, and J. Bigün. Face authentication with Gabor information on deformable graphs. *IEEE Transactions on Image Processing*, 8(4):504–516, April 1999.
- [61] J.-L. Dugelay. Method for hiding binary data in a digital image. Pending Patent PCT/FR99/00485 (EURECOM 09-PCT), March 1999.
- [62] J.-L. Dugelay and S. Roche. Process for marking a multimedia document, such an image, by generating a mark. Pending Patent EP 99480075.3 (EURECOM 11/12 EP), July 1999.
- [63] D. Dunn, W. Higgins, and J. Wakeley. Texture segmentation using 2-D Gabor elementary functions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(2):130–149, February 1994.
- [64] DVD Copy Control Association. <http://www.dvdcca.org>.
- [65] J. Eggers, R. Bäuml, R. Tzschoppe, and B. Girod. Scalar costa scheme for information embedding. *IEEE Transactions on Signal Processing*, 51(4):1003–1019, April 2003.
- [66] J. Eggers and B. Girod. *Informed Watermarking*. Kluwer Academic Publishers, 2002.

- [67] A. Eskicioglu. Multimedia security in group communications: Recent progress in key management, authentication and watermarking. *ACM Multimedia Systems, Special Issue on Multimedia Security*, 9(3):239–248, September 2003.
- [68] Y. Fisher. *Fractal Image Compression: Theory and Applications*. Springer-Verlag, 1994.
- [69] J. Foley and G. Legge. Contrast masking in human vision. *Journal of the Optical Society of America*, 70(12):1458–1470, December 1980.
- [70] J. Fridrich and M. Goljan. Robust hash functions for digital watermarking. In *Proceedings of the International Conference on Information Technology: Coding and Computing*, pages 178–183, March 2000.
- [71] E. Garcia. *Tatouage d’Objets 3D Basé sur la Texture*. PhD thesis, Université de Nice Sophia-Antipolis, France, July 2004.
- [72] E. Garcia and J.-L. Dugelay. Texture-based watermarking of 3D video objects. *IEEE Transactions on Circuits and Systems for Video Technology*, 13(8):853–866, August 2003.
- [73] M. Gharavi-Alkhansari and T. Huang. A fractal-based image block-coding algorithm. In *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume V, pages 345–348, April 1993.
- [74] A. Giannoula and D. Hatzinakos. Compressive data hiding for video signals. In *Proceedings of the IEEE International Conference on Image Processing*, volume I, pages 529–532, September 2003.
- [75] C. Griwodz, O. Merkel, J. Dittmann, and R. Steinmetz. Protecting VoD the easier way. In *Proceedings of the ACM Multimedia Conference*, pages 21–28, September 1998.
- [76] J. Haitisma and T. Kalker. A watermarking scheme for digital cinema. In *Proceedings of the IEEE International Conference on Image Processing*, volume II, pages 487–489, October 2001.
- [77] Ö. Harmancı, M. Kucukgoz, and K. Mihçak. Temporal synchronization of watermarked video using image hashing. In *Security, Steganography and Watermarking of Multimedia Contents VII*, volume 5681 of *Proceedings of SPIE*, pages 370–380, January 2005.
- [78] F. Hartung, P. Eisert, and B. Girod. Digital watermarking of MPEG-4 facial animation parameters. *Computers & Graphics*, 22(4):425–435, July 1998.

- [79] F. Hartung and B. Girod. Watermarking of uncompressed and compressed video. *Signal Processing*, 66(3):283–301, May 1998.
- [80] F. Hartung, J. Su, and B. Girod. Spread spectrum watermarking: Malicious attacks and counterattacks. In *Security and Watermarking of Multimedia Contents*, volume 3657 of *Proceedings of SPIE*, pages 147–158, January 1999.
- [81] Herodotus. *The Histories*. Penguin Books, 1996.
- [82] M. Holliman, W. Macy, and M. Yeung. Robust frame-dependent video watermarking. In *Security and Watermarking of Multimedia Contents II*, volume 3971 of *Proceedings of SPIE*, pages 186–197, January 2000.
- [83] M. Holliman and N. Memon. Counterfeiting attack on oblivious block-wise independent invisible watermarking schemes. *IEEE Transactions on Image Processing*, 9(3):432–441, March 2000.
- [84] C.-Y. Hsu and C.-S. Lu. Geometric distortion-resilient image hashing system and its application scalability. In *Proceedings of the ACM Multimedia and Security Workshop*, pages 81–92, September 2004.
- [85] M. Irani, P. Anandan, J. Bergen, R. Kumar, and S. Hsu. Mosaic representations of video sequences and their applications. *Signal Processing: Image Communication*, 8(4):327–351, May 1996.
- [86] M. Johnson and K. Ramchandran. Dither-based secure image hashing using distributed coding. In *Proceedings of the IEEE International Conference on Image Processing*, volume II, pages 751–754, September 2003.
- [87] N. Johnson, Z. Duric, and S. Jajodia. *Information Hiding: Steganography and Watermarking – Attacks and Countermeasures*. Elsevier Academic Press, 2001.
- [88] I. Jolliffe. *Principal Component Analysis*. Springer-Verlag, 1986.
- [89] F. Jordan, M. Kutter, and T. Ebrahimi. Proposal of a watermarking technique for hiding/retrieving data in compressed and decompressed video. In *JTC1/SC29/WG11 MPEG97/M2281*. ISO/IEC, July 1997.
- [90] JPEG Standard. Digital compression and coding of continuous-tone still images. In *JTC1/SC29/WG1 10918-1*. ISO/IEC, February 1994.
- [91] D. Kahn. *The Codebreakers : The Comprehensive History of Secret Communication from Ancient Times to the Internet*. Scribner, 1996.

- [92] T. Kalker. Considerations on watermarking security. In *Proceedings of the IEEE Fourth Workshop on Multimedia Signal Processing*, pages 201–206, October 2001.
- [93] T. Kalker, G. Depovere, J. Haitsma, and M. Maes. A video watermarking system for broadcast monitoring. In *Security and Watermarking of Multimedia Contents*, volume 3657 of *Proceedings of SPIE*, pages 103–112, January 1999.
- [94] T. Kalker and J. Haitsma. Efficient detection of a spatial spread-spectrum watermarking in MPEG video streams. In *Proceedings of the IEEE International Conference on Image Processing*, volume I, pages 434–437, September 2000.
- [95] S. Katzenbeisser and F. Petitcolas. *Information Hiding: Techniques for Steganography and Digital Watermarking*. Artech House, 1999.
- [96] D. Kelly. Motion and vision II: Stabilized spatiotemporal threshold surface. *Journal of the Optical Society of America*, 69(10):1340–1349, October 1979.
- [97] A. Kerckhoffs. La cryptographie militaire. *Journal des sciences militaires*, IX:5–83, January 1883.
- [98] D. Kilburn. Dirty linen, dark secrets. *Adweek*, 38(40):35–40, November 1997.
- [99] S.-W. Kim, S. Suthaharan, H.-K. Lee, and K. Rao. Perceptually tuned robust watermarking scheme for digital video using motion entropy masking. In *Proceedings of the IEEE International Conference on Consumer Electronics*, pages 104–105, June 1999.
- [100] D. Kirovski and H. Malvar. Robust covert communication over a public audio channel using spread spectrum. In *Proceedings of the Fourth International Workshop on Information Hiding*, volume 2137 of *Lecture Notes in Computer Science*, pages 354–368, April 2001.
- [101] D. Kirovski, H. Malvar, and Y. Yacobi. Multimedia content screening using a dual watermarking and fingerprinting system. In *Proceedings of the Tenth ACM International Conference on Multimedia*, pages 372–381, November 2002.
- [102] D. Kirovski and F. Petitcolas. Blind pattern matching attack on watermarking systems. *IEEE Transactions on Signal Processing*, 51(4):1045–1053, April 2003.

- [103] D. Kirovski and F. Petitcolas. Replacement attack on arbitrary watermarking systems. In *Proceedings of the ACM Digital Rights Management Workshop*, volume 2696 of *Lecture Notes in Computer Science*, pages 177–189, July 2003.
- [104] P. Kocher, J. Jaffe, B. Jun, C. Laren, and N. Lawson. Self-protecting digital content. Cryptography Research Inc. White Paper, April 2003.
- [105] R. Koenen. MPEG-4 overview. In *JTC1/SC29/WG11 N4668*. ISO/IEC, March 2002.
- [106] S. Kozat, R. Venkatesan, and K. Mhçak. Robust perceptual image hashing via matrix invariants. In *Proceedings of the IEEE International Conference on Image Processing*, pages 3443–3446, October 2004.
- [107] M. Kucukgoz, Ö. Harmancı, K. Mhçak, and R. Venkatesan. Robust video watermarking via optimization algorithm for quantization of pseudo-random semi-global statistics. In *Security, Steganography and Watermarking of Multimedia Contents VII*, volume 5681 of *Proceedings of SPIE*, pages 363–369, January 2005.
- [108] M. Kutter. Watermarking resisting to translation, rotation and scaling. In *Multimedia Systems and Applications*, volume 3528 of *Proceedings of SPIE*, pages 423–431, November 1998.
- [109] M. Kutter and F. Petitcolas. A fair benchmark for image watermarking systems. In *Security and Watermarking of Multimedia Contents*, volume 3657 of *Proceedings of SPIE*, pages 226–239, January 1999.
- [110] M. Kutter, S. Voloshynovskiy, and A. Herrigel. Watermark copy attack. In *Security and Watermarking of Multimedia Contents II*, volume 3971 of *Proceedings of SPIE*, pages 371–380, January 2000.
- [111] M. Lades, J. Vorbrüggen, J. Buhmann, J. Lange, C. Malsburg, R. Würtz, and W. Konen. Distortion invariant object recognition in the dynamic link architecture. *IEEE Transactions on Computers*, 42(3):300–311, March 1993.
- [112] G. Langelaar and R. Lagendijk. Optimal differential energy watermarking of DCT encoded images and video. *IEEE Transactions on Image Processing*, 10(1):148–158, January 2001.
- [113] G. Langelaar, R. Lagendijk, and J. Biemond. Real-time labelling of MPEG-2 compressed video. *Journal of Visual Communication and Image Representation*, 9(4):256–270, August 1998.

- [114] C.-H. Lee, H.-S. Oh, and H.-K. Lee. Adaptive video watermarking using motion information. In *Security and Watermarking of Multimedia Contents II*, volume 3971 of *Proceedings of SPIE*, pages 209–216, January 2000.
- [115] F. Lefèbvre, B. Macq, and J.-D. Legat. RASH: RAdon Soft Hash algorithm. In *Proceedings of the European Signal Processing Conference*, volume I, pages 299–302, September 2002.
- [116] J. Lewis. Power to the peer. *LAWeekly*, May 2002.
- [117] Q. Li and I. Cox. Using perceptual models to improve fidelity and provide invariance to volumetric scaling for quantization index modulation watermarking. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, volume II, pages 1–4, March 2005.
- [118] J. Lichtenauer, I. Setyawan, T. Kalker, and R. Lagendijk. Exhaustive geometrical search and the false positive watermark detection probability. In *Security and Watermarking of Multimedia Contents V*, volume 5020 of *Proceedings of SPIE*, pages 203–214, January 2003.
- [119] J. Lichtenauer, I. Setyawan, and R. Lagendijk. Hiding correlation-based watermark templates using secret modulation. In *Security, Steganography and Watermarking of Multimedia Contents VI*, volume 5306 of *Proceedings of SPIE*, pages 501–512, January 2004.
- [120] J. Lim. *Two-Dimensional Signal and Image Processing*. Prentice Hall International Editions, 1989.
- [121] E. Lin and E. Delp. Temporal synchronization in video watermarking. In *Security and Watermarking of Multimedia Contents IV*, volume 4675 of *Proceedings of SPIE*, pages 478–490, January 2002.
- [122] E. Lin and E. Delp. Temporal synchronization in video watermarking. *IEEE Transactions on Signal Processing, Supplement on Secure Media*, 52(10):3007–3022, October 2004.
- [123] E. Lin, C. Podilchuk, T. Kalker, and E. Delp. Streaming video and rate scalable video: What are the challenges for watermarking? In *Security and Watermarking of Multimedia Contents III*, volume 4314 of *Proceedings of SPIE*, pages 116–127, January 2001.
- [124] Y. Linde, A. Buzo, and R. Gray. An algorithm for vector quantizer design. *IEEE Transactions on Communications*, 28(1):84–95, January 1980.

- [125] J.-P. Linnartz. The ticket concept for copy control based on signal embedding. In *Proceedings of the Fifth European Symposium on Research in Computer Security*, volume 1485 of *Lecture Notes in Computer Science*, pages 257–274, September 1998.
- [126] J.-P. Linnartz and M. van Dijk. Analysis of the sensitivity attack against electronic watermarks in images. In *Proceedings of the Second International Workshop on Information Hiding*, volume 1525 of *Lecture Notes in Computer Science*, pages 258–272, April 1998.
- [127] J. Lixvar. Watermarking requirements for boeing digital cinema. In *Security and Watermarking of Multimedia Contents V*, volume 5020 of *Proceedings of SPIE*, pages 546–552, January 2003.
- [128] C.-S. Lu. Wireless multimedia error resilience via a data hiding technique. In *Proceedings of the IEEE Fifth Workshop on Multimedia Signal Processing*, pages 316–319, December 2002.
- [129] C.-S. Lu, J.-R. Chen, H.-Y. Liao, and K.-C. Fan. Real-time MPEG2 video watermarking in the VLC domain. In *Proceedings of the IEEE International Conference on Pattern Recognition*, volume II, pages 552–555, August 2002.
- [130] J. Lubin, J. Bloom, and H. Cheng. Robust, content-dependent, high-fidelity watermark for tracking in digital cinema. In *Security and Watermarking of Multimedia Contents V*, volume 5020 of *Proceedings of SPIE*, pages 536–545, January 2003.
- [131] W. Macy and M. Holliman. Quality evaluation of watermarked video. In *Security and Watermarking of Multimedia Contents II*, volume 3971 of *Proceedings of SPIE*, pages 486–500, January 2000.
- [132] M. Maes, T. Kalker, J. Haisma, and G. Depovere. Exploiting shift invariance to obtain high payload in digital image watermarking. In *Proceedings of the IEEE International Conference on Multimedia Computing and Systems*, volume I, pages 7–12, June 1999.
- [133] A. Menezes, P. van Oorschot, and S. Vanstone. *Handbook of Applied Cryptography*. CRC Press, 1996.
- [134] J. Meng and S. Chang. Tools for compressed-domain video indexing and editing. In *Storage and Retrieval for Image and Video Database*, volume 2670 of *Proceedings of SPIE*, pages 180–191, March 1996.
- [135] K. Mihçak and R. Venkatesan. New iterative geometric methods for robust perceptual image hashing. In *Proceedings of the ACM Workshop on Security and Privacy in Digital Rights Management*, pages 13–21, November 2001.



- [136] M. Miller, G. Doërr, and I. Cox. Applying informed coding and informed embedding to design a robust, high capacity watermark. *IEEE Transactions on Image Processing*, 13(6):792–807, June 2004.
- [137] B. Mobasseri and D. Cinalli. Watermarking of compressed multimedia using error-resilient VLCs. In *Proceedings of the IEEE Fifth Workshop on Multimedia Signal Processing*, pages 320–323, December 2002.
- [138] B. Mobasseri and D. Cinalli. Reversible watermarking using two-way decodable codes. In *Security, Steganography and Watermarking of Multimedia Contents VI*, volume 5306 of *Proceedings of SPIE*, pages 397–404, January 2004.
- [139] B. Mobasseri, M. Sieffert, and R. Simard. Content authentication and tamper detection in video. In *Proceedings of the IEEE International Conference on Image Processing*, volume I, pages 458–461, September 2000.
- [140] V. Monga and B. Evans. Robust perceptual image hashing using feature points. In *Proceedings of the IEEE International Conference on Image Processing*, pages 677–680, October 2004.
- [141] D. Mukherjee, J. Chae, and S. Mitra. A source and channel coding approach to data hiding with applications to hiding speech in video. In *Proceedings of the IEEE International Conference on Image Processing*, volume I, pages 348–352, October 1998.
- [142] H. Nicolas. New methods for dynamic mosaicing. *IEEE Transactions on Image Processing*, 10(8):1239–1251, August 2001.
- [143] H. Nicolas and C. Labit. Motion and illumination variation estimation using a hierarchy of models: Application to image sequence coding. *Journal of Visual Communication and Image Representation*, 6(4):303–316, December 1995.
- [144] X. Niu, M. Schmucker, and C. Busch. Video watermarking resisting to rotation, scaling and translation. In *Security and Watermarking of Multimedia Contents IV*, volume 4675 of *Proceedings of SPIE*, pages 512–519, January 2002.
- [145] G. Øien, S. Lepsøy, and T. Ramstad. An inner product space approach to image coding by contractive transformations. In *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume III, pages 2773–2776, May 1991.

- [146] J. Oostveen, T. Kalker, and M. Staring. Adaptive quantization watermarking. In *Security, Steganography and Watermarking of Multimedia Contents VI*, volume 5306 of *Proceedings of SPIE*, pages 296–303, January 2004.
- [147] OpenWatermark. <http://www.openwatermark.org>.
- [148] Optimark. <http://poseidon.csd.auth.gr/optimark>.
- [149] A. Patrizio. Why the DVD hack was a cinch. *Wired*, November 1999.
- [150] F. Pérez-González, F. Balado, and J. Hernández-Martin. Performance analysis of existing and new methods for data hiding with known-host information in additive channels. *IEEE Transactions on Signal Processing*, 51(4):960–980, April 2003.
- [151] F. Petitcolas, R. Anderson, and M. Kuhn. Attacks on copyright marking systems. In *Proceedings of the Second International Workshop on Information Hiding*, volume 1525 of *Lecture Notes in Computer Science*, pages 219–239, April 1998.
- [152] F. Petitcolas and D. Kirovski. The blind pattern matching attack on watermarking systems. In *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume IV, pages 3740–3743, May 2002.
- [153] R. Pickholtz, D. Schilling, and L. Millstein. Theory of spread spectrum communications - a tutorial. *IEEE Transactions on Communications*, 30(5):855–884, May 1982.
- [154] A. Piva, R. Caldelli, and A. De Rosa. A DWT-based object watermarking system for MPEG-4 video streams. In *Proceedings of the IEEE International Conference on Image Processing*, volume III, pages 5–8, September 2000.
- [155] L. Qiao and K. Nahrstedt. Watermarking methods for MPEG encoded video: Towards resolving rightful ownership. In *Proceedings of the IEEE International Conference on Multimedia Computing and Systems*, pages 276–285, June 1998.
- [156] C. Rey. *Tatouage d’Images: Gain en Robustesse et Intégrité des Images*. PhD thesis, Université d’Avignon, France, February 2003.
- [157] C. Rey, K. Amis, J.-L. Dugelay, R. Pyndiah, and A. Picart. Enhanced robustness in image watermarking using block turbo codes. In *Security and Watermarking of Multimedia Contents V*, volume 5020 of *Proceedings of SPIE*, January 2003.

- [158] C. Rey, G. Doërr, J.-L. Dugelay, and G. Csurka. Toward generic image dewatermarking? In *Proceedings of the IEEE International Conference on Image Processing*, volume III, pages 633–636, September 2002.
- [159] C. Rey and J.-L. Dugelay. A survey of watermarking algorithms for image authentication. *EURASIP Journal on Applied Signal Processing*, 2002(6):613–621, June 2002.
- [160] D. Ringach. Spatial structure and symmetry of simple-cell receptive fields in macaque primary visual cortex. *Journal of Neurophysiology*, 88(1):455–463, July 2002.
- [161] D. Roberts. Security camera video authentication. In *Proceedings of the IEEE Tenth Digital Signal Processing Workshop*, pages 125–130, October 2002.
- [162] D. Robie and R. Mersereau. Video error correction using steganography. *EURASIP Journal on Applied Signal Processing*, 2002(2):164–173, February 2002.
- [163] S. Roweis. EM algorithms for PCA and SPCA. *Neural Information Processing Systems*, 10:626–632, July 1998.
- [164] J. Ó Ruanaidh and T. Pun. Rotation, scale and translation invariant digital image watermarking. *Signal Processing*, 68(3):303–317, May 1998.
- [165] P. Sallee. Model-based steganography. In *Proceedings of the Second International Workshop on Digital Watermarking*, volume 2939 of *Lecture Notes in Computer Science*, pages 154–167, October 2003.
- [166] A. Sankar. Experiments with a gaussian merging-splitting algorithm for HMM training for speech recognition. In *Proceedings of DARPA Speech Recognition Workshop*, pages 99–104, February 1998.
- [167] Secure Digital Music Initiative. <http://www.sdmi.org>.
- [168] H. Sencar, M. Ramkumar, and A. Akansu. *Data Hiding Fundamentals and Applications – Content Security in Digital Multimedia*. Elsevier Academic Press, 2004.
- [169] G. Simmons. The prisoners’ problem and the subliminal channel. In *Proceedings of CRYPTO*, pages 51–67, 1983.
- [170] P. Smaragdis and M. Casey. Audio/visual independent components. In *Proceedings of the 4th International Symposium on Independent Component Analysis and Blind Signal Separation*, pages 709–714, April 2003.

- [171] A. Smolic, M. Lorei, and T. Sikora. Adaptive kalman-filtering for prediction and global motion parameter tracking of segments in video. In *Proceedings of the Picture Coding Symposium*, March 1996.
- [172] J. Song and R. Liu. A data embedded video coding scheme for error-prone channels. *IEEE Transactions on Multimedia*, 3(4):415–423, December 2001.
- [173] Stirmark. <http://www.petitcolas.net/fabien/watermarking/stirmark>.
- [174] K. Su, D. Kundur, and D. Hatzinakos. A novel approach to collusion resistant video watermarking. In *Security and Watermarking of Multimedia Contents IV*, volume 4675 of *Proceedings of SPIE*, pages 491–502, January 2002.
- [175] K. Su, D. Kundur, and D. Hatzinakos. Spatially localized image-dependent watermarking for statistical invisibility and collusion resistance. *IEEE Transactions on Multimedia*, 7(1):52–66, February 2005.
- [176] K. Su, D. Kundur, and D. Hatzinakos. Statistical invisibility for collusion-resistant digital video watermarking. *IEEE Transactions on Multimedia*, 7(1):43–51, February 2005.
- [177] J. Sun, J. Liu, and H. Hu. Data hiding in independent components of video. In *Proceedings of 5th International Conference on Independent Component Analysis and Blind Signal Separation*, pages 970–976, September 2004.
- [178] S.-W. Sun and P.-C. Chang. Video watermarking synchronization based on profile statistics. *IEEE Aerospace and Electronics Magazine*, 19(5):21–25, May 2004.
- [179] Z. Sun and M. Tekalp. Trifocal motion modeling for object-based video compression and manipulation. *IEEE Journal on Circuits and Systems for Video Technology*, 8(5):667–685, May 1998.
- [180] A. Swaminathan, Y. Mao, and M. Wu. Image hashing resilient to geometric and filtering operations. In *Proceedings of the IEEE Sixth Workshop on Multimedia Signal Processing*, pages 355–358, September 2004.
- [181] M. Swanson, B. Zhu, and A. Tewfik. Data hiding for video-in-video. In *Proceedings of the IEEE International Conference on Image Processing*, volume II, pages 676–679, October 1997.
- [182] M. Swanson, B. Zhu, and A. Tewfik. Multiresolution scene-based video watermarking using perceptual models. *IEEE Journal on Selected Areas in Communications*, 16(4):540–550, May 1998.

- [183] R. Szeliski and H.-Y. Shum. Creating full view panoramic image mosaics and environment maps. In *Proceedings of the International Conference on Computer Graphics and Interactive Techniques*, pages 251–258, April 1997.
- [184] P. Termont, L. De Stycker, J. Vandewege, M. Op de Beeck, J. Haitsma, T. Kalker, M. Maes, and G. Depovere. How to achieve robustness against scaling in a real-time digital watermarking system for broadcast monitoring. In *Proceedings of the IEEE International Conference on Image Processing*, volume 1, pages 407–410, September 2000.
- [185] J. van Hateren. Spatiotemporal contrast sensitivity for early vision. *Vision Research*, 33(2):257–267, January 1992.
- [186] A. van Leest, J. Haitsma, and T. Kalker. On digital cinema and watermarking. In *Security and Watermarking of Multimedia Contents V*, volume 5020 of *Proceedings of SPIE*, pages 526–535, January 2003.
- [187] R. Vankatesan, S.-M. Koon, M. Jacobowski, and P. Moulin. Robust image hashing. In *Proceedings of the IEEE International Conference on Image Processing*, volume III, pages 664–666, September 2000.
- [188] Visual Quality Expert Group (VQEG). <http://www.vqeg.org>.
- [189] S. Voloshynovskiy, A. Herrigel, N. Baumgärtner, and T. Pun. A stochastic approach to content adaptive digital image watermarking. In *Proceedings of the Third International Workshop on Information Hiding*, volume 1768 of *Lecture Notes in Computer Science*, pages 211–236, September 1999.
- [190] S. Voloshynovskiy, A. Herrigel, and Y. Rystar. The watermark template attack. In *Security and Watermarking of Multimedia Contents III*, volume 4314 of *Proceedings of SPIE*, pages 394–405, January 2001.
- [191] S. Voloshynovskiy, S. Pereira, A. Herrigel, N. Baumgärtner, and T. Pun. Generalized watermarking attack based on watermark estimation and perceptual remodulation. In *Security and Watermarking of Multimedia Contents II*, volume 3971 of *Proceedings of SPIE*, pages 358–370, January 2000.
- [192] S. Voloshynovskiy, S. Pereira, V. Iquise, and T. Pun. Attack modeling: Towards a second generation watermarking benchmark. *Signal Processing*, 81(6):1177–1214, June 2001.
- [193] Watermark Evaluation Testbed (WET). Contact Professor E. Delp, Purdue University, USA.

- [194] A. Watson. DCT quantization matrices optimized for individual images. In *Human Vision, Visual Processing and Digital Display IV*, volume 1913 of *Proceedings of SPIE*, pages 202–216, September 1993.
- [195] A. Watt. *3D Computer Graphics*. Addison-Wesley, third edition, 2000.
- [196] P. Wayner. *Disappearing Cryptography – Information Hiding: Steganography & Watermarking*. Morgan Kaufmann Publishers, 2002.
- [197] S. Winkler, E. Gelasca, and T. Ebrahimi. Towards perceptual metrics for video watermark evaluation. In *Applications of Digital Image Processing*, volume 5203 of *Proceedings of SPIE*, pages 371–378, August 2003.
- [198] D. Wu, Y. Hou, and Y.-Q. Zhang. Transporting real-time video over the Internet: Challenges and approaches. *Proceedings of the IEEE*, 88(12):1855–1875, December 2000.
- [199] M. Wu and B. Liu. Attacks on digital watermarks. In *Proceedings of 33rd Asilomar Conference on Signals, Systems, and Computers*, volume II, pages 1508–1512, October 1999.
- [200] M. Wu, W. Trappe, J. Wang, and R. Liu. Collusion-resistant fingerprinting for multimedia. *IEEE Signal Processing Magazine*, 21(2):15–27, March 2004.
- [201] P. Yin, B. Liu, and H. Yu. Error concealment using data hiding. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing*, volume III, pages 1453–1456, May 2001.
- [202] J. Zhang, J. Li, and L. Zhang. Video watermark technique in motion vector. In *Proceedings of the 14th Brazilian Symposium on Computer Graphics and Image Processing*, pages 179–182, October 2001.
- [203] L.-H. Zhang, H.-T. Wu, and C.-L. Hu. A video watermarking algorithm based on 3D Gabor transform. *Journal of Software*, 15(8):1252–1258, August 2004.