Security and Robustness Constraints for Spread-Spectrum Tardos Fingerprinting

Benjamin Mathon, Patrick Bas, François Cayre, Benoît Macq

SECURITY AND ROBUSTNESS CONSTRAINTS FOR SPREAD-SPECTRUM TARDOS FINGERPRINTING

Benjamin Mathon [1,3]*, Patrick Bas* [2]*, François Cayre* [1]*, Benoît Macq* [3]

[1] GIPSA-Lab, Grenoble-INP
961 rue de la Houille Blanche
BP 46
38402 Grenoble Cedex
France

[2] LAGIS, Ecole Centrale Lille
Avenue Paul Langevin
BP 48
59651 Villeneuve d'Ascq Cedex
France

[3] TELE, UCL
Place du Levant 2
Bâtiment Stévin
1348 Louvain-la-Neuve
Belgium

ABSTRACT

This paper presents a practical analysis of the impact of robustness and security on Tardos' collusion-secure fingerprinting codes using spread-spectrum watermarking modulations. In this framework, we assume that the coalition has to face an embedding scheme of given security level and consequently has to suffer a probability of wrongly estimating their embedded symbols. We recall the Worst Case Attack associated to this probability, e.g. the optimal attack which minimises the mutual information between the sequence of a colluder and the pirated one. For a given achievable rate of the Tardos' fingerprinting model, we compare the Improved Spread-Spectrum embedding versus a new secure embedding (called ρ-Circular Watermarking) considering the AWGN channel. We show that secure embeddings are more immune to decoding errors than non-secure ones while keeping the same fingerprinting capacity.

***Index Terms*—** Traitor tracing, Watermarking, Security

1. INTRODUCTION

Traitor tracing using fingerprinting codes is nowadays a relevant solution in order to be able to trace users or devices who deliver over the Internet copyrighted contents without authorisation. One practical way to develop a traitor tracing scheme is to use watermarking in order to embed fingerprinting codes in the content [1] in a robust way. However, if this solution is adopted, one has also to take into account the possible security of the embedding method. The design of fingerprinting codes has been studied in various works [2, 3, 4, 5, 6] and the security analysis of watermarking schemes has also been investigated [7, 8] but outside of the fingerprinting framework. In the sequel, we study the impact of the use of secure embeddings in the context of fingerprinting using spread-spectrum methods and present dedicated collusion strategies.

2. NOTATIONS

We first list the conventions used in this article. Functions are noted in roman fonts, sets in calligraphy fonts and variables in italic fonts. Vectors and matrices are set in bold fonts, vectors are written in small letters and matrices in capital ones. $\mathbf{x}(i)$ is the i-th component of a vector $\mathbf{x}$. As for the C programming language, all indexes start from 0. We write $(\mathbf{x}(0)\dots\mathbf{x}(m-1))$ the content of a vector $\mathbf{x}$ of length m. $\sigma^2_{\mathbf{x}}$ represents the unbiased variance of a signal $\mathbf{x}$ and $\langle .|. \rangle$ denotes the usual scalar product. $\|\mathbf{x}\|$ denotes the Euclidean norm of a vector $\mathbf{x}$. If n is an integer, $[n]$ denotes the set $[\![0; n-1]\!]$. $\#\mathcal{A}$ is the cardinality of the set $\mathcal{A}$, $\mathrm{span}\,(\mathcal{A})$ represents the vector space spanned by $\mathcal{A}$. $\mathcal{M}_{n,m}(\mathbb{K})$ denotes the set of n-by-m matrices whose components are in the division ring $\mathbb{K}$.

3. TARDOS PROBABILISTIC FINGERPRINT CODES

We recall here the basic principles of Gabor Tardos' traitor tracing codes.

3.1. Fingerprints generation

Gabor Tardos' probabilistic codes [9] are used to generate fingerprints of size N_m over an alphabet Σ of size $\#\Sigma = q$ considering n users. The construction for $\Sigma = \mathbb{F}_2$ is the following: we first generate N_m values $\{p_0, \dots, p_{N_m-1}\}$ in $[0; 1]$ independently with p.d.f. $\mathrm{f}(p)$:

$$\mathrm{f}(p) = \left(\pi\sqrt{p(1-p)}\right)^{-1}. \tag{1}$$

Next we construct the fingerprint matrix $\mathbf{M} \in \mathcal{M}_{n,N_m}(\mathbb{F}_2)$: $\forall j \in [n]\ \forall i \in [N_m]\ \mathbf{M}(j,i) \sim \mathcal{B}(p_i)$, which is a Bernoulli distribution with parameter p_i. Each row of $\mathbf{M}$ is used as a fingerprint $\mathbf{m}_j$ which identify user $j \in [n]$ and is embedded into the content.

3.2. Collusion attacks

A coalition denotes a subset of users (called colluders) who wants to produce an unauthorised copy of the content. They forge a corrupted fingerprint $\mathbf{m} \in \mathbb{F}_2^{N_m}$ by mixing their sequences bit per bit in order to avoid to be traced by the distributor whenever the content will be found. This attack is called a collusion attack. Tardos' fingerprinting codes are collusion-secure against coalition $\mathcal{C} \subset [n]$ of size c. It is assumed the symbols are chosen according to a strategy beforehand defined by the colluders. The accusation process is only possible if the following conditions are satisfied:

- $N_m \sim O\left(c^2 \log\left(n/p_{fa}\right)\right)$ (the Peikert's theoretical lower bound [10]), p_{fa} denotes the false alarm probability: the probability of accusing an innocent (i.e. a user $j \notin \mathcal{C}$),
- the symbols of the pirated fingerprint are only selected into the symbols of the coalition, formally:

$$\mathbf{m} = \left(\mathbf{m}_{j'_0}(0) \dots \mathbf{m}_{j'_{N_m-1}}(N_m-1)\right), \quad \text{with } \left(j'_0 \dots j'_{N_m-1}\right) \in \mathcal{C}^{N_m}. \tag{2}$$

  This last condition is called the "marking assumption" [11].

The accusation process uses an accusation matrix in order to compute an accusation score S_j for each user j who is accused if the score is above a threshold T depending of the probability of accusing an innocent (see [9] for details).

3.3. Colluding strategies

We consider that the fingerprints codes are generated following Tardos' procedure and are embedded using a watermarking scheme. According to the security of the modulation, in the WOA (Watermarked Contents Only Attack [7]) framework, colluders would be able to compute an estimation of the secret key and to decode an estimation of the embedded sequences. In this context, we denote $\hat{\mathbf{m}}_j$, the fingerprint estimated by a colluder $j \in \mathcal{C}$ and ϵ, the estimation error, theoretically defined for each position $i \in [N_m]$ by:

$$\begin{aligned} \epsilon &= \Pr\left(\hat{\mathbf{m}}_j(i) = 1 | \mathbf{m}_j(i) = 0\right) \\ &= \Pr\left(\hat{\mathbf{m}}_j(i) = 0 | \mathbf{m}_j(i) = 1\right). \end{aligned} \tag{3}$$

As in [12], we call the strategy of the coalition the process used for forging a pirated sequence (estimated) $\hat{\mathbf{m}}$ from the estimated sequences $\hat{\mathbf{m}}_j$ of the coalition. A strategy is a function of the number of symbols "1" that a coalition estimates at position $i \in [N_m]$. For each position i, the strategy is completely defined by a $(c+1)$-vector θ:

$$\theta(k) = \Pr\left(\hat{\mathbf{m}}(i) = 1 \,\middle|\, \sum_{j \in \mathcal{C}} \hat{\mathbf{m}}_j(i) = k\right). \tag{4}$$

The pirated fingerprint $\mathbf{m}$ follows:

$$\forall i \in [N_m],\ \mathbf{m}(i) = \mathbf{m}_{j'}(i), \tag{5}$$

where j' is uniformly chosen in $\{j \in \mathcal{C} : \hat{\mathbf{m}}_j(i) = \hat{\mathbf{m}}(i)\}$. This condition respects the marking assumption Eq. (2): if, at one position i, they all have the symbol "1" (resp. "0"), necessarily $\mathbf{m}(i)$ equals "1" (resp. "0").

3.4. Worst Case Attacks

From the estimated symbols and a given accusation function, the coalition is able to build an attack which minimises the achievable rate R_s (in bits/sample) of the fingerprinting scheme defined by [6, 3][1]:

$$\mathrm{R}_s(\theta, \epsilon) = \mathbb{E}_P[I(M; M_{j_0}) | P = p], \tag{6}$$

where P is the random variable with p.d.f $\mathrm{f}(p)$ defined in Eq. (1), I denotes the mutual information, M the random variable associated to a bit of the pirated sequence and M_{j_0} , the random variable for a bit of the sequence of a colluder $j_0 \in \mathcal{C}$. According to the marking assumption defined in Eq. (2), M is chosen among the bits of the coalition $\mathcal{C}$. In [6], authors emerge the "Worst Case Attack" (WCA), the strategy θ which minimises $\mathrm{R}_s(\theta, \epsilon)$ for $\epsilon = 0$. In recent works [13], we have generalised this attack for any $\epsilon \in [0; 0.5]$. For $\epsilon \neq 0$ this attack, called ϵ-WCA, is more efficient than the classical WCA and allows an increase of the p_{fa}. The computation of the achievable rate is obtained by using conditional probabilities and combinatorial analysis (see [13] for details). We find the attack $\theta_{\epsilon\text{-WCA}}$ by minimising R_s using Simplex algorithm. As an example, Table 1 shows strategies $\theta_{\epsilon\text{-WCA}}$ for different values of ϵ with 3 or 4 colluders.

	$c = 3$	$c = 4$
$\epsilon = 0.$	(0. 0.651 0.349 1.)	(0. 0.487 0.5 0.513 1.)
$\epsilon = 0.05$	(0. 0.726 0.274 1.)	(0. 0.543 0.5 0.457 1.)
$\epsilon = 0.1$	(0. 0.830 0.170 1.)	(0. 0.620 0.5 0.379 1.)
$\epsilon = 0.15$	(0. 0.982 0.018 1.)	(0. 0.734 0.5 0.266 1.)
$\epsilon = 0.2$	(0. 1. 0. 1.)	(0. 0.908 0.5 0.091 1.)
$\epsilon > 0.2$	(0. 1. 0. 1.)	(0. 1. 0.5 0. 1.)

Table 1. Values of $\theta_{\epsilon\text{-WCA}}$ functions of ϵ for $c = 3, 4$. For $c = 2$, for all ϵ, $\theta_{\epsilon\text{-WCA}}$ = (0. 0.5 1.).

4. MODULATIONS FOR FINGERPRINT EMBEDDING

We present in this section different spread-spectrum modulations offering different security levels.

In the following, indexes j denote users $j \in [n]$. We consider a message $\mathbf{m}_j \in \mathbb{F}_2^{N_c}$ (N_c bits) we want to embed into a

[1] Note that in our framework, the rate is now also function of ϵ.

host signal $\mathbf{x} \in \mathbb{R}^{N_v}$. The secret key used for embedding are N_c carriers $\mathbf{u}_i \in \mathbb{R}^{N_v}$. These carriers are generated thanks to a Pseudo Random Number Generator initialised with a seed $K \in \mathbb{N}$. They come as Gaussian vectors and are further orthogonalised (using Gram-Schmidt procedure) with unit variance in order to provide a basis of the private subspace i.e. $\forall i \neq j\ ,\ \langle \mathbf{u}_i | \mathbf{u}_j \rangle = 0$. For each bit of the message we want to hide, we use a modulation s: $\mathbb{F}_2 \times \mathbb{R}^{N_v} \to \mathbb{R}$.

$$\mathrm{s}(\mathbf{m}_j(i), \mathbf{x}) = \alpha(i, \mathbf{x})(-1)^{\mathbf{m}_j(i)} - \lambda(\mathbf{x})\langle \mathbf{x}, \mathbf{u}_i \rangle, \quad (7)$$

where:

- $\alpha(i, \mathbf{x})$ allows to adjust the distortion of each carrier,
- $\lambda(\mathbf{x})$ allows to adjust informed embedding.

The watermark signal $\mathbf{y}_j$ is constructed as an addition of the host signal and the modulated carriers:

$$\mathbf{y}_j = \mathbf{x} + \mathbf{w}_j = \mathbf{x} + \sum_{i=0}^{N_c - 1} \mathrm{s}\left(\mathbf{m}_j(i), \mathbf{x}\right) \mathbf{u}_i, \quad (8)$$

where $\mathbf{w}_j$ denotes the watermark signal for user $j \in [n]$. We measure the embedding distortion by the Watermark-to-Content Ratio (WCR) and the distortion induced by the attack by the Noise-to Content Ratio (NCR) or the Watermark-to Noise Ratio (WNR).

We now consider the attacked vector $\mathbf{r}_j = \mathbf{y}_j + \mathbf{n}$, where $\mathbf{n}$ vector is composed with Gaussian iid coefficients. Decoding is assured by using correlations z between the (probably) attacked vector $\mathbf{r}_j$ and the carriers $\mathbf{u}_i$:

$$\mathrm{z}_{\mathbf{r}_j, \mathbf{u}_i} = \langle \mathbf{r}_j, \mathbf{u}_i \rangle. \quad (9)$$

If $\mathbf{m}'_j$ denotes the decoded message for user $j \in [n]$, for each bit we have:

$$\mathbf{m}'_j(i) = \begin{cases} 0 \text{ if } \mathrm{z}_{\mathbf{r}_j, \mathbf{u}_i} > 0, \\ 1 \text{ if } \mathrm{z}_{\mathbf{r}_j, \mathbf{u}_i} < 0. \end{cases} \quad (10)$$

4.1. Secure and insecure modulations

4.1.1. Security Attacks

In the WOA framework, an adversary owns several marked contents and tries to estimate the secret carriers. If the used modulation is not secure, he is able to estimate the carriers. The quality of the estimation depends on the number of marked signals this adversary has access to, of the embedding method and on the WCR. The estimation is a BSS (Blind Source Separation) problem. Source separation techniques as ICA (Independent Component Analysis [14]) are powerful to perform BSS problem when original sources are independently drawn (the modulations $\mathrm{s}\left(\mathbf{m}_j(i), \mathbf{x}\right)$ in our case).

4.1.2. Improved Spread-Spectrum (ISS)

The ISS modulation [15] is able to cancel the interference due to the host signal in order to improve the robustness of the watermarking scheme. $\alpha_{\text{ISS}}(i, \mathbf{x})$ and $\lambda_{\text{ISS}}(\mathbf{x})$ are computed to achieve host-interference rejection and error probability minimization given a target WCR and NCR: In the WOA framework, following the classification of [8], this modulation is *insecure*. Adversaries are able to estimate the carriers for example by ICA [16], provided that they can gather enough contents.

4.1.3. Circular Watermarking (CW)

The CW modulation [17] is based on ISS modulation but it uses a parameter $\mathbf{d}$, which is generated at each embedding from a standardized Gaussian signal $\mathbf{g}$. This perturbation is used to randomly spread the correlations of the mixed signals on the whole decoding regions:

$$\alpha_{\text{CW}}(i, \mathbf{x}) = \mathbf{d}(i)\alpha_{\text{ISS}}(i, \mathbf{x}), \quad (11)$$

with:

$$\mathbf{d}(i) = |\mathbf{g}(i)| / \|\mathbf{g}\|. \quad (12)$$

This modulation is *key-secure* [17], adversaries are able to estimate the private subspace generated by the carriers: $\mathrm{span}(\mathbf{u}_i)$ but they have no more information about the decoding regions.

4.1.4. ρ-Circular Watermarking (ρ-CW)

ρ-CW is a new modulation derived from the CW modulation: we spread the correlations on the decoding regions but these correlations are translated following a scalar parameter ρ on each dimension and then rescaled. Formally we construct:

$$\alpha_{\rho\text{-CW}}(i, \mathbf{x}) = \mathbf{d}_\rho(i)\alpha_{\text{ISS}}(i, \mathbf{x}), \quad (13)$$

with $\mathbf{d}_\rho \in \mathbb{R}^{N_c}$, $\mathbf{d}_\rho = (\mathbf{d} + \rho.\mathbf{1}) / \|\mathbf{d} + \rho.\mathbf{1}\|$ ($\mathbf{1}$ denotes the constant vector whose each component equals 1).

When $\rho >> 0$, ρ-CW is an insecure modulation, when $\rho = 0$, ρ-CW is the classical CW modulation. With this modulation, we are able to tune the security and the robustness because the quality of the estimation of the secret key relies both on the WCR and on the number of observations own by the coalition (see sec. 4.1.1). Fig. 1 shows ISS, CW and ρ-CW correlations with $N_c = 2,\ N_v = 512,\ WCR = -10dB,\ NCR = -10dB$.

5. SECURY VS ROBUSTNESS FOR SPREAD-SPECTRUM MODULATIONS FACING WORST CASE ATTACKS

In the fingerprinting scenario, a distributor creates different copies of a host content for users. Each copy is marked with a

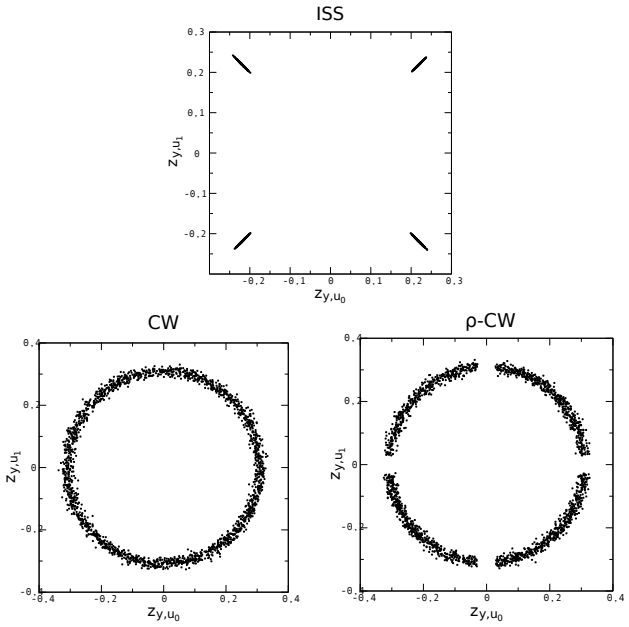

Fig. 1. Correlations of ISS, CW and ρ-CW ($\rho = 0.1$) signals over the secret carriers using 2000 observations: $N_c = 2,\ N_v = 512,\ WCR = -10dB,\ NCR = -10dB$.

message which identifies the concerned user. He first extracts a feature signal from an host content (image, video, audio or movie file) divided into N_t chunks of N_v components. Each fingerprint is a Tardos' code of N_m bits which will be insert into the host feature vector.

In this section, we use the new modulation ρ-CW defined in Eq. (13) for spread-spectrum watermarking. We hide $N_c = 16$ bits on each chunk of size $N_v = 512$. For our experimentations, host signals are standardized Gaussian distributed, $WCR = -10dB,\ NCR = -10dB$. First, we look for the estimation error ϵ that a colluder can compute based on the number N_t of chunks of his fingerprinting content and on the parameter ρ of the ρ-CW modulation. Next we measure the performances of ρ-CW considering the robustness point of view. Finally we quantify the compromise between security and robustness of the fingerprinting model by using on one hand, an insecure modulation (ISS) and on the other hand a secure one (ρ-CW).

5.1. Spread-spectrum embedding respecting the marking assumption

One important point in Tardos' fingerprint construction is to respect the marking assumption given Eq. (2). We have seen before that the strategy defined by the colluders on the estimated bits at one position will respect the marking assumption if the candidate symbol used to forge the pirated sequence is chosen among the correct symbol of a colluder. In our fingerprinting model using spread-spectrum watermarking, the strategy defined in Eq. (4) cannot be applied if $N_c > 1$ (if $N_c = 1$, colluders forge the pirated signal by mixing their chunks).

To solve this problem, we use the same embedding technique that the one proposed by Hartung and Kutter [18] with the difference that secure embedding is used: considering one chunk of size N_v, the secret key used to hide the N_c bits are N_c carriers $\mathbf{u}_i \in \mathbb{R}^{N_v}$ constructed as follows:

$$\forall\ j \in [N_v],\ \mathbf{u}_i(j) \begin{cases} \sim \mathcal{N}(0,1) \text{ if } j \in [\![i\frac{N_v}{N_c}; (i+1)\frac{N_v}{N_c} - 1]\!], \\ = 0 \text{ else.} \end{cases} \tag{14}$$

These carriers are further standardized. It means that, for a chunk $\mathbf{y}_j$ constructed as in Eq. (8), the first bit is only embedded into the $\frac{N_v}{N_c}$ first components of the chunk, the second bit in the $\frac{N_v}{N_c}$ following components, etc. Consequently the colluders can forge a pirated signal by mixing their fragments of $\frac{N_v}{N_c}$ components on each chunk according to a strategy hence respecting the marking assumption.

5.2. Computation of ϵ

Fig. 2 determines the value of the estimation error functions of the number of chunks N_t used by the colluder $j \in \mathcal{C}$ for the security attack. This estimation error is computed on average on 10 secret keys by measuring the bit error rate between the original messages and the messages decoded using estimated carriers. Note that ICA have limitations because it can only estimate the carriers up to their sign and it cannot estimate the order of the carriers (in the WOA framework – adversaries need the knowledge of a number of messages to do so, in the order of $\log_2(N_c)$). For these experimentation, we consider that the colluders are able to solve these limitations by using side information.

As can be seen, when $\rho = 0$, we use the classical key-secure CW, and it is not possible to correctly estimate the embedded symbols even with a large number of observations: ϵ stay stationary to 0.34. However, when ρ grows up, the estimation of the messages is more and more accurate. ϵ depends on values on ρ and N_t.

5.3. Security vs Robustness

First we consider the effects of robustness attacks on watermarked signals by the mean of the bit error rate η. η is theoretically defined for each position $i \in N_m$ by:

$$\begin{aligned} \eta &= \Pr\left(\mathbf{m}'(i) = 1 | \mathbf{m}(i) = 0\right), \\ &= \Pr\left(\mathbf{m}'(i) = 0 | \mathbf{m}(i) = 1\right), \end{aligned} \tag{15}$$

where $\mathbf{m}'$ is the attacked pirated sequence decoded by the distributor to identify members of the coalition.

Fig. 3 shows the bit error rate η of ρ-CW modulation functions of WNR for some values of ρ. As can be seen,

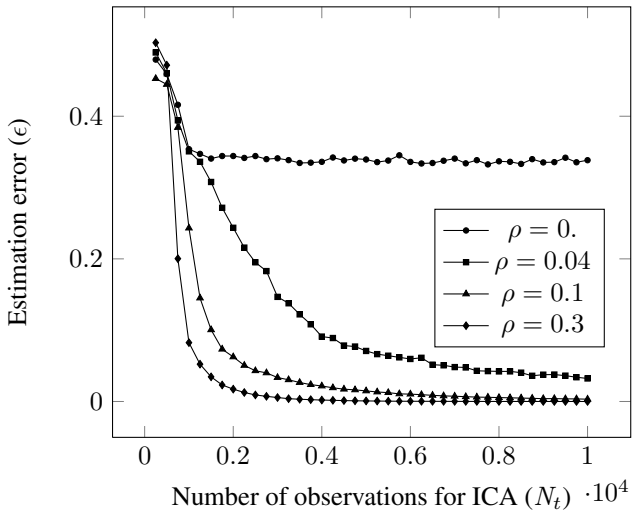

Fig. 2. Estimation error for ρ-CW w.r.t. number of chunks (N_t) used for the ICA security attack by one colluder $j \in \mathcal{C}$, $N_c = 16$ $N_v = 512$, $WCR = -10dB$, $NCR = -10dB$.

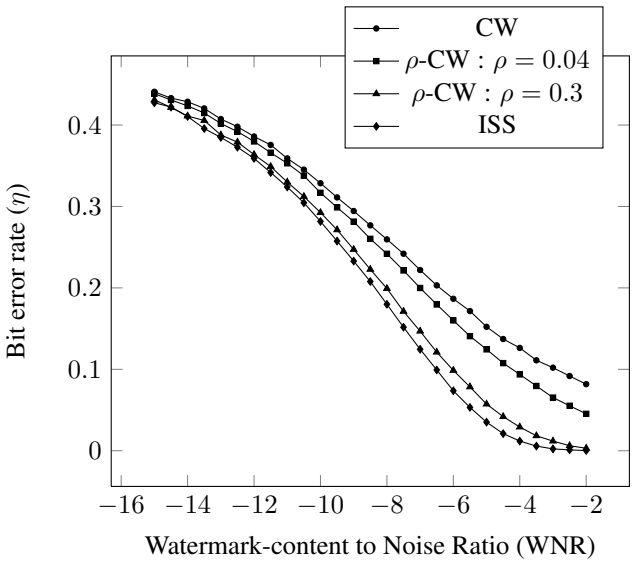

Fig. 3. Bit error rate (η) w.r.t. WNR for ISS and ρ-CW, $N_c = 16$ $N_v = 512$, $WCR = -10dB$, $NCR = -10dB$. As can be seen, robustness increases when ρ grows up, it highlights the compromise between security and robustness that watermark schemes have to deal with when comparing to Fig. 2.

robustness increases when ρ grows up, it highlights the compromise between security and robustness that watermark schemes have to deal with when comparing to Fig. 2.

In order to merge the $\epsilon - WCA$ with robustness attack, we define the new achievable fingerprinting rate (in bits/sample) by:

$$\mathrm{R}'_s(\theta, \epsilon, \eta) = \mathbb{E}_P[I(M'; M_{j_0})|P = p]. \quad (16)$$

where M' denotes the bit of the attacked pirated sequence and M_{j_0} the bit of the sequence of a colluder. The achievable rate R'_s is computed as the same way as Eq. (6) [13].

In Fig. 4, we compare the amount of noise that secure and insecure embedding schemes can undergo in order to achieve the same fingerprinting capacity. Formally this is done by solving η_1 (or its associated WNR) satisfying the following equation:

$$\mathrm{R}'_s(\theta_{\epsilon WCA}, \epsilon, \eta_1) = \mathrm{R}'_s(\theta_{WCA}, 0., \eta_2). \quad (17)$$

with $\epsilon = 0.3$ estimated by ICA with $N_t = 1300, \rho = 0.04$.

Practically, for the ISS modulation suffering a given WNR, we compute the corresponding η_2 on 10^6 signals. Next, we find η_1 using Eq. (17) (this computation is made using root-finding algorithm as in [13]) and its corresponding WNR for ρ-CW .

As can be seen in Fig. 4, using robust and secure scheme (ρ-CW) offers more robustness for the same achievable rate than using only a robust scheme like ISS when the robustness grows up[2]. Note also that the difference between ISS and ρ-CW (from the fingerprinting capacity point of view) is only significant for highly secure schemes (when $\epsilon < 0.25$, the curve on the figure given by Eq. (17) is close to identity).

For small WNRs the gain between secure embedding and insecure scheme is around $5dB$.

With theses experiments, we highlight the importance of using highly secure and robust watermarking compared to only highly robust ones.

6. CONCLUSION

In this paper, we have introduced a new spread-spectrum modulation, called ρ-CW, which is flexible for a given level of security and robustness. We have highlighted the impact of watermarking constraints when spread-spectrum is used for embedding Tardos' traitor tracing codes thanks to computations of achievable rates of fingerprinting models. In this case, we have shown the advantage of using secure and robust modulation (ρ-CW) instead of insecure but more robust one (ISS). Future works will be devoted to the study of optimal strategies on q-ary fingerprinting schemes when $q > 2$.

7. REFERENCES

[1] F. Xie, T. Furon, and C. Fontaine, "On-off keying modulation and Tardos fingerprinting," in *MM&Sec '08: Proceedings of the 10th ACM workshop on Multimedia and*

[2]Results for higher WNR are not available because η_1 and η_2 becomes rare events ($\sim 10^{-5}$) and are not estimated correctly with classical Monte-Carlo procedure.

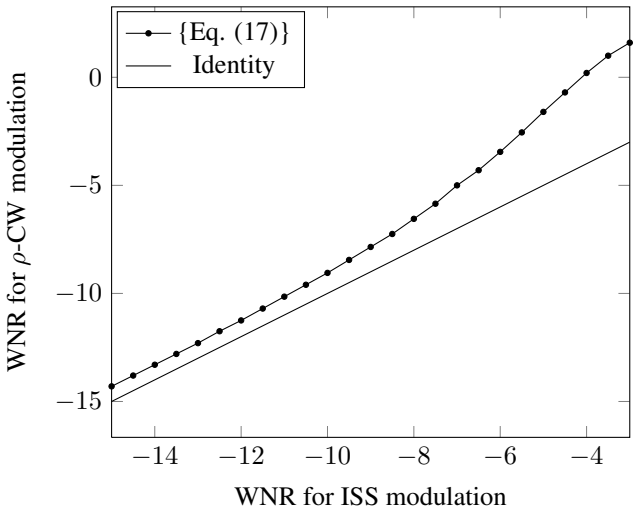

Fig. 4. WNR for ρ-CW (with ber η_1) w.r.t WNR for ISS (with ber η_2), $N_c = 16$ $N_v = 512$, $WCR = -10dB$, $NCR = -10dB$. Relation between η_1 and η_2 is given by: $\mathrm{R}'_s(\theta_{\epsilon WCA}, \epsilon, \eta_1) = \mathrm{R}'_s(\theta_{WCA}, 0., \eta_2)$. $\epsilon = 0.3$ obtained by ICA with $N_t = 1300, \rho = 0.04$.

security. New York, NY, USA: ACM, 2008, pp. 101–106.

[2] W. Trappe, M. Wu, Z. Wang, and K. Liu, "Anti-collusion fingerprinting for multimedia," *Signal Processing, IEEE Transactions on*, vol. 51, no. 4, pp. 1069–1087, Apr. 2003.

[3] P. Moulin, "Universal fingerprinting: Capacity and random-coding exponents," in *Information Theory, 2008. ISIT 2008. IEEE International Symposium on*, July 2008, pp. 220–224.

[4] B. Skoric, T. Vladimirova, M. Celik, and J. Talstra, "Tardos Fingerprinting is Better Than We Thought," *Information Theory, IEEE Transactions on*, vol. 54, no. 8, pp. 3663–3676, Aug. 2008.

[5] T. Furon, A. Guyader, and F. Cérou, "On the design and optimisation of Tardos probabilistic fingerprinting codes," in *Information Hiding*, ser. Lecture Notes in Computer Science, vol. 5284. Springer Berlin / Heidelberg, 2008, pp. 341–356.

[6] T. Furon, L. Pérez-Freire, A. Guyader, and F. Cérou, "Estimating the minimal length of Tardos code," in *Information Hiding*, ser. Lecture Notes in Computer Science, vol. 5806. Springer Berlin / Heidelberg, 2009, pp. 176–190.

[7] F. Cayre, C. Fontaine, and T. Furon, "Watermarking Security: Theory and Practice," *Signal Processing, IEEE Transactions on*, vol. 53, no. 10, pp. 3976–3987, Oct. 2005.

[8] F. Cayre and P. Bas, "Kerckhoffs-Based Embedding Security Classes for WOA Data Hiding," *Information Forensics and Security, IEEE Transactions on*, vol. 3, no. 1, pp. 1–15, Mar. 2008.

[9] G. Tardos, "Optimal probabilistic fingerprint codes," *J. ACM*, vol. 55, no. 2, pp. 1–24, 2008.

[10] C. Peikert, A. Shelat, and A. Smith, "Lower bounds for collusion-secure fingerprinting," in *SODA '03: Proceedings of the fourteenth annual ACM-SIAM symposium on Discrete algorithms*. Philadelphia, PA, USA: Society for Industrial and Applied Mathematics, 2003, pp. 472–479.

[11] B. Chor, A. Fiat, M. Naor, and B. Pinkas, "Tracing traitors," in *Proc. of Crypto'94*, ser. Lecture Notes in Computer Science, vol. 839. Springer Berlin / Heidelberg, 1994, pp. 257–270.

[12] T. Furon and L. Pérez-Freire, "Worst case attacks against binary probabilistic traitor tracing codes," *IEEE Transactions on Information Forensics and Security*, 2009.

[13] B. Mathon, P. Bas, F. Cayre, and B. Macq, "Considering security and robustness constraints for watermark-based tardos fingerprinting," in *IEEE SPS Multimedia Signal Processing Conference (MMSP 2010)*, Saint Malo, France, Oct. 2010.

[14] A. Hyvärinen, J. Karhunen, and E. Oja, *Independent Component Analysis*. John Wiley & Sons, 2001.

[15] H. Malvar and D. Florencio, "Improved spread spectrum: a new modulation technique for robust watermarking," *Signal Processing, IEEE Transactions on*, vol. 51, no. 4, pp. 898–905, Apr. 2003.

[16] B. Mathon, P. Bas, and F. Cayre, "Practical Performance Analysis of Secure Modulations for WOA Spread-Spectrum based Image Watermarking," in *MM&Sec '07: Proceedings of the 9th workshop on Multimedia & security*. New York, NY, USA: ACM, 2007, pp. 237–244.

[17] P. Bas and F. Cayre, "Achieving Subspace or Key Security for WOA using Natural or Circular Watermarking," in *MM&Sec '06: Proceedings of the 8th workshop on Multimedia and security*. New York, NY, USA: ACM, 2006, pp. 80–88.

[18] F. Hartung and M. Kutter, "Multimedia watermarking techniques," *Proceedings of the IEEE*, vol. 87, no. 7, pp. 1079–1107, July 1999.