Secret dither estimation in lattice-quantization data hiding: a set-membership approach*

Luis Pérez-Freire, Fernando Pérez-Gonzalez and Pedro Comesaña

Signal Theory and Communications Dept.
University of Vigo
36310 Vigo, Spain

ABSTRACT

In this paper, security of lattice-quantization data hiding is considered under a cryptanalytic point of view. Security in this family of methods is implemented by means of a pseudorandom dither signal which randomizes the codebook, preventing unauthorized embedding and/or decoding. However, the theoretical analysis shows that the observation of several watermarked signals can provide sufficient information for an attacker willing to estimate the dither signal, quantifying information leakages in different scenarios. The practical algorithms proposed in this paper show that such information leakage may be successfully exploited with manageable complexity, providing accurate estimates of the dither using a small number of observations. The aim of this work is to highlight the security weaknesses of lattice data hiding schemes whose security relies only on secret dithering.

Keywords: Lattice data hiding, DC-DM, security, residual entropy, set-membership estimation

1. INTRODUCTION

Recently, there have been some attempts at addressing the concept of watermarking security[1,2] from a cryptanalytic point of view, where all the information about watermarking schemes is public, and security relies only on the use of secret keys that parameterize the embedding and/or decoding processes. In additive spread-spectrum methods, for instance, security usually depends on the secrecy of the pseudorandom sequence which is added to the host signal. Theoretical studies[1,2] have shown that, for these methods, some information about the secret key leaks from the observation of watermarked signals, so an attacker can take advantage of this information leakage to obtain an estimate of that pseudorandom sequence. Thereafter, the attacker may use his knowledge about the secret key to devise smart attacks to defeat the watermarking scheme.

In this paper we consider the security of lattice-quantization data hiding schemes for which security relies only on the secret dither signal that randomizes the codebook.[3,4] In our set-up, the attacker is able to gather a collection of N_o watermarked signals along with additional information in some instances. Depending on the degree of additional information, three different scenarios are considered, following the nomenclature introduced in[1]: a) Known Message Attack (KMA), where the attacker knows the messages embedded in each watermarked signal; b) Constant Message Attack (CMA), where all watermarked signals are assumed to convey the same message; and c) Watermarked Only Attack (WOA), where no additional information is available. Although all these scenarios are addressed under theoretical and practical points of view, the main focus is put on the KMA scenario because of its central role in the analysis of the other cases, showing the relationship between them.

*This work was partially funded by *Xunta de Galicia* under projects PGIDT04 TIC322013PR and PGIDT04 PXIC32202PM; MEC project DIPSTICK, reference TEC2004-02551/TCM; FIS project IM3, reference G03/185 and European Comission through the IST Programme under Contract IST-2002-507932 ECRYPT. ECRYPT disclaimer: The information in this paper is provided as is, and no guarantee or warranty is given or implied that the information is fit for any particular purpose. The user thereof uses the information at its sole risk and liability.

Further author information: (Send correspondence to F.P.G.)
L.P.F.: E-mail: lpfreire@gts.tsc.uvigo.es
F.P.G.: E-mail: fperez@gts.tsc.uvigo.es
P.C.: E-mail: pcomesan@gts.tsc.uvigo.es

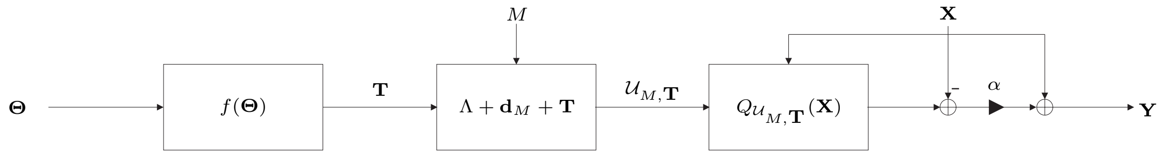
$\boldsymbol{\Theta}$
$f(\boldsymbol{\Theta})$
$\mathbf{T}$
M
$\Lambda + \mathbf{d}_M + \mathbf{T}$
$\mathcal{U}_{M,\mathbf{T}}$
$Q_{\mathcal{U}_{M,\mathbf{T}}}(\mathbf{X})$
$\mathbf{X}$
-
α
$\mathbf{Y}$

Figure 1. Block diagram of lattice-based DC-DM.

For the theoretical part, the measures of security used in this paper were introduced in[2]; following Shannon's approach to assess the security of cryptographic systems,[5] information leakages are measured here by the mutual information between the observations and the secret dither (a real-valued vector), and the remaining uncertainty about the latter is quantified by the residual (conditional) entropy.

In the practical part, two different algorithms for estimating the secret dither are proposed and tested, comparing their performance with that of the optimal estimator. The proposed algorithms are suitable for performing dither estimation with arbitrary lattices in any number of dimensions, although the computational complexity of such algorithms may be increased with the dimensionality (roughly speaking, the complexity will depend on the number of facets of the Voronoi regions of the considered lattice). To the best of the authors' knowledge no results in this direction were published before for quantization-based schemes, with the exception of,[6] but the problem was restricted there to scalar quantization with no distortion compensation, and as such it can be regarded as a subproblem of our general formulation.

The main notational conventions used in the rest of the paper are the following: scalar variables and vectors are respectively denoted by italicized letters and boldface letters; capital letters denote random variables, whereas lowercase letters are their occurrences. Calligraphic letters are reserved for sets, and all vectors are regarded as column vectors.

2. THEORETICAL SECURITY OF LATTICE DATA HIDING

Before proceeding with the theoretical analysis, we will briefly explain the basics of embedding in lattice data-hiding; for more details and other aspects such as decoding, the interested reader is referred to.[3] Consider a n-dimensional lattice Λ and the set $\mathcal{M} = \{0, \ldots, L_M - 1\}$ of possible messages. For each message $m \in \mathcal{M}$ let us define the associated coset of Λ as $\mathcal{U}_m = \Lambda + \mathbf{d}_m$, where Λ is a n-dimensional lattice, and $\mathbf{d}_m$ is the minimum-norm *coset representative* corresponding to message m. The codebook $\mathcal{U}$ is defined by the union of all cosets, $\mathcal{U} = \bigcup_{m=0}^{L_M-1} \mathcal{U}_m$. Given a certain host signal $\mathbf{x}$ and a to-be-transmitted message $m \in \mathcal{M}$, each watermarked signal is generated as

$$\mathbf{y} = \mathbf{x} + \alpha(Q_{\mathcal{U}_m,\mathbf{t}}(\mathbf{x}) - \mathbf{x}) = Q_{\mathcal{U}_m,\mathbf{t}}(\mathbf{x}) + (1-\alpha)(\mathbf{x} - Q_{\mathcal{U}_m,\mathbf{t}}(\mathbf{x})), \tag{1}$$

where α is the distortion compensation parameter, and $Q_{\mathcal{U}_{m,\mathbf{t}}}(\cdot)$ is an Euclidean quantizer whose centroids are defined by the coset $\mathcal{U}_m + \mathbf{t}$. The term $\mathbf{t}$ in (1) is the secret dither vector, which is known only by embedder and decoder, and whose aim is to improve the security of the scheme by randomly translating the codebook.[3, 4] The second term of (1) is the so-called *self-noise* term. The data hiding scheme just described is commonly known as Distortion Compensated - Dither Modulation (DC-DM),[3] and is summarized in the block diagram of Fig. 1. As it is usual in the analysis of quantization-based methods for data hiding,[3, 4] we will assume a low embedding distortion regime, thus we can consider that the host pdf is approximately uniform inside each quantization cell and all centroids in $\mathcal{U}$ occur with similar probabilities. This assumption (which we will refer to in the sequel as the *flat-host assumption*) implies that the statistical distribution of the self-noise is approximately uniform over $(1-\alpha)\mathcal{V}(\Lambda)$, where $\mathcal{V}(\Lambda)$ denotes the Voronoi region of Λ, defined as

$$\mathcal{V}(\Lambda) \triangleq \{\mathbf{x} \in \mathbb{R}^n : Q_\Lambda(\mathbf{x}) = \mathbf{0}\}. \tag{2}$$

The flat-host assumption permits us to simplify the theoretical analysis, restricting our attention to the modulo-reduced random variable $\tilde{\mathbf{Y}} \triangleq \mathbf{Y} \mod \Lambda = \mathbf{Y} - Q_\Lambda(\mathbf{Y})$. Hence, the pdf of $\tilde{\mathbf{Y}}$ conditioned on the embedded message and the secret dither is

$$f(\tilde{\mathbf{y}}|m,\mathbf{t}) = \begin{cases} \mathrm{vol}(\mathcal{Z}(\Lambda))^{-1}, & \tilde{\mathbf{y}} \in (\mathbf{d}_m + \mathbf{t} + \mathcal{Z}(\Lambda)) \bmod \Lambda \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

where $\mathcal{Z}(\Lambda) \triangleq (1-\alpha)\mathcal{V}(\Lambda)$. In our model, as is customary in theoretical analyses of watermarking methods, the host samples are considered independent and identically distributed (i.i.d.). In case that the dimensionality of the host is $N_v = N \cdot n$, embedding is made in disjoint groups of n samples using a different secret dither for each of them. Clearly, for maximizing the residual entropy about $\mathbf{t}$, its samples should also be i.i.d. In the following, we only consider n-dimensional hosts, but all the results on mutual informations and entropies are easily generalizable to hosts of dimension $N \cdot n$ by multiplying them by a factor N, due to the independence assumed.

Under these premises, a theoretical security analysis will be developed in the following for the three scenarios introduced in Section 1. Obviously, the security level of the system will depend on the statistical distribution of the secret dither, or better to say, of its modulo-Λ reduced version, $\tilde{\mathbf{T}}$, due to the Λ-periodicity of the watermark generation (see Eq. (1)). In other words, the watermarked signal depends in last instance on $\tilde{\mathbf{T}}$, and hence the secrecy of the codebook depends only on the statistics of $\tilde{\mathbf{T}}$.[†] Therefore, we will assume hereinafter that the support of $\mathbf{T}$ is bounded by $\mathcal{V}(\Lambda)$.

2.1. Known Message Attack

When a sequence of watermarked signals $\{\tilde{\mathbf{Y}}_1, \ldots, \tilde{\mathbf{Y}}_{N_o}\}$ and its associated messages $\{M_1, \ldots, M_{N_o}\}$ are observed, the information leakage about $\mathbf{T}$ can be calculated as

$$I(\tilde{\mathbf{Y}}_1, \ldots, \tilde{\mathbf{Y}}_{N_o}; \mathbf{T}|M_1, \ldots, M_{N_o}) = h(\mathbf{T}) - h(\mathbf{T}|\tilde{\mathbf{Y}}_1, \ldots, \tilde{\mathbf{Y}}_{N_o}, M_1, \ldots, M_{N_o}), \quad (4)$$

since $\mathbf{T}$ is independent of the embedded messages. The second term of (4) is the residual entropy of the dither after N_o observations, which depends on its statistical distribution conditioned to the observations. Through the application of Bayes' rule, it is easy to see that the conditional pdf of the dither is given by

$$f(\mathbf{t}|\tilde{\mathbf{y}}_1, \ldots, \tilde{\mathbf{y}}_{N_o}, m_1, \ldots, m_{N_o}) = \frac{f(\mathbf{t}) \cdot \prod_{i=1}^{N_o} f((\tilde{\mathbf{y}}_i - \mathbf{d}_{m_i} - \mathbf{t}) \bmod \Lambda | M_i = 0, \mathbf{T} = \mathbf{0})}{f(\tilde{\mathbf{y}}_1, \ldots, \tilde{\mathbf{y}}_{N_o}|m_1, \ldots, m_{N_o})}, \quad (5)$$

where we have made use of the flat-host assumption. By recalling Eq. (3), it is clear that each term in the numerator of (5) is nonzero iff $(\tilde{\mathbf{y}}_i - \mathbf{d}_{m_i} - \mathbf{t}) \bmod \Lambda \in \mathcal{Z}(\Lambda)$, or equivalently, iff $\mathbf{t} \in \mathcal{D}_i$, with $\mathcal{D}_i$ given by

$$\mathcal{D}_i \triangleq (\tilde{\mathbf{y}}_i - \mathbf{d}_{m_i} - \mathcal{Z}(\Lambda)) \bmod \Lambda, \; i = 1, \ldots, N_o, \quad (6)$$

Hence, it is clear that the support of $\mathbf{t}$ is contained in $\mathcal{S}_{N_o} \triangleq \bigcap_{i=1}^{N_o} \mathcal{D}_i$, independently of the distribution of $\mathbf{T}$. We well refer to $\mathcal{S}_{N_o}$ in the following as the *feasible region*. Obviously, the watermarker is interested in maximizing the residual entropy, for which he should choose $\mathbf{T} \sim U(\mathcal{V}(\Lambda))$. The reason is that in this case the conditional pdf of the dither results to be uniform in $\mathcal{S}_{N_o}$, this way maximizing the residual entropy. Furthermore, we have that

$$h(\mathbf{T}|\tilde{\mathbf{Y}}_1, \ldots, \tilde{\mathbf{Y}}_{N_o}, M_1, \ldots, M_{N_o}) = E[\log(\mathrm{vol}(\mathcal{S}_{N_o}))], \quad (7)$$

so for the case of a single observation ($N_o = 1$) it is straightforward to see that

$$h(\mathbf{T}|\tilde{\mathbf{Y}}_1, M_1) = \log(\mathrm{vol}(\mathcal{Z}(\Lambda))) = \log((1-\alpha)^n \mathrm{vol}(\mathcal{V}(\Lambda))), \quad (8)$$

and it is immediate to notice that the information leakage is given by

$$I(\tilde{\mathbf{Y}}_1; \mathbf{T}|M_1) = h(\mathbf{T}) - h(\mathbf{T}|\tilde{\mathbf{Y}}_1, M_1) = \log(\mathrm{vol}(\mathcal{V}(\Lambda))) - \log((1-\alpha)^n \mathrm{vol}(\mathcal{V}(\Lambda))) = -n \log(1-\alpha), \quad (9)$$

independently of the specific lattice chosen for embedding. This result clearly shows a trade-off between security and achievable rate: theoretical analyses[4,7] show that, in AWGN channels, the value of α must approach

[†] However, attacks at a cryptographic level might be indeed interested in knowing the exact value of $\mathbf{t}$.

1 for maximizing the achievable rate in the high-SNR region; however, bear in mind that for $\alpha \approx 1$, one observation suffices to get an accurate estimate of the centroids in $\mathcal{U}_m$, and consequently of the secret dither, due to the structure imposed to the codebook. This is reflected in the residual entropy of the dither (8), for which $\lim_{\alpha\to 1} h(\mathbf{T}|\tilde{\mathbf{Y}}_1, M_1) = -\infty$.

For the general case of $N_o > 1$, one may be tempted to upper bound the mutual information by assuming that all observations will provide the same amount of information, but the resulting bound will be too loose in most cases. In fact, the mutual information can be shown to be a strictly increasing, concave function of the number of observations N_o. This result will be checked in Section 2.4.

2.2. Constant Message Attack

The easiest way of addressing this scenario is to regard it as a collection of several KMA problems. When the message embedded is unknown but unchanged for the whole sequence of observations, the conditional pdf of the dither after N_o observations can be expressed as

$$f(\mathbf{t}|\tilde{\mathbf{y}}_1,\ldots,\tilde{\mathbf{y}}_{N_o},\mathrm{CM}) = \frac{1}{|\mathcal{M}|}\sum_{m=1}^{|\mathcal{M}|} f(\mathbf{t}|\tilde{\mathbf{y}}_1,\ldots,\tilde{\mathbf{y}}_{N_o}, m,\ldots,m), \tag{10}$$

where CM stands for *constant message*. This means that the feasible region $\mathcal{S}_{N_o}^{CMA}$ for the dither in the CMA case is simply the union of the feasible regions of $|\mathcal{M}|$ different KMA problems. Formally,

$$\mathcal{S}_{N_o}^{CMA} = \bigcup_{m=1}^{|\mathcal{M}|} (\mathcal{S}_{N_o} + \mathbf{d}_m), \tag{11}$$

where $\mathcal{S}_{N_o}$ is the feasible region defined for the KMA case in the previous section. Since $\mathrm{vol}(\mathcal{S}_{N_o})$ is reduced with the number of observations, the different regions that constitute $\mathcal{S}_{N_o}^{CMA}$ will be disjoint for sufficiently large N_o; in such case, the residual entropy is again maximized if $\mathbf{T} \sim U(\mathcal{V}(\Lambda))$ is chosen. Due to (10), the residual entropy can be upper bounded as

$$h(\mathbf{T}|\tilde{\mathbf{Y}}_1,\ldots,\tilde{\mathbf{Y}}_{N_o},\mathrm{CM}) \le h(\mathbf{T}|\tilde{\mathbf{Y}}_1,\ldots,\tilde{\mathbf{Y}}_{N_o}, M_1,\ldots,M_{N_o}) + \log(|\mathcal{M}|), \tag{12}$$

resulting in a lower bound to the information leakage. Equality in (12) is achieved when the regions $\mathcal{S}_{N_o} + \mathbf{d}_m$ are disjoint, thus the bound will be asymptotically tight for increasing N_o. Anyway, if the value of α is above a certain threshold (which is increasing with the nesting ratio) the above upper bound is an equality for every N_o. This value of α must be such that the conditional pdf's of the watermarked signal conditioned to each message do not overlap; it is easy to see that such value is given by $\alpha_t = 1 - 1/|\mathcal{M}|$ for any lattice, as long as the lattice partition is self-similar.

2.3. Watermarked Only Attack

In this case we have to compute the residual entropy $h(\mathbf{t}|\tilde{\mathbf{Y}}_1,\ldots,\tilde{\mathbf{Y}}_{N_o})$. The pdf of the secret dither conditioned on the past observations is given by

$$f(\mathbf{t}|\tilde{\mathbf{y}}_1,\cdots,\tilde{\mathbf{y}}_{N_o}) = \frac{f(\mathbf{t})\cdot\prod_{i=1}^{N_o} f(\tilde{\mathbf{y}}_i|\mathbf{t})}{f(\tilde{\mathbf{y}}_1,\cdots,\tilde{\mathbf{y}}_{N_o})} = \sum_{i=1}^{|\mathcal{M}|^{N_o}} f(\mathbf{t}|\tilde{\mathbf{y}}_1,\ldots,\tilde{\mathbf{y}}_{N_o},\mathbf{m}_i^{N_o})\cdot \Pr\{\mathbf{m}_i^{N_o}|\tilde{\mathbf{y}}_1,\ldots,\tilde{\mathbf{y}}_{N_o}\}, \tag{13}$$

where $\mathbf{m}_i^{N_o}$, $i = 1,\ldots,|\mathcal{M}|^{N_o}$, covers the whole range of message sequences of length N_o. Notice that the pdf (13) is not necessarily uniform, since it is, in fact, a weighted sum of uniform pdf's.

Clearly, this is the most pessimistic scenario for the attacker, due to the additional uncertainty introduced by the unknown message sequence. However, if the value of α is above a given threshold (that again depends on the nesting ratio) which guarantees that there are only $|\mathcal{M}|$ message sequences with non-null probability, then it is easy to see that the residual entropy in this case is given by

$$h(\mathbf{T}|\tilde{\mathbf{Y}}_1,\ldots,\tilde{\mathbf{Y}}_{N_o}) = h(\mathbf{T}|\tilde{\mathbf{Y}}_1,\ldots,\tilde{\mathbf{Y}}_{N_o}, M_1,\ldots,M_{N_o}) + \log(|\mathcal{M}|). \tag{14}$$

The value of α that fulfills the condition necessary for (14) to hold can be shown to be $\alpha_t' = 1 - \frac{1}{2|\mathcal{M}|}$ for self-similar lattice partitions. For values of α below α_t', the residual entropy may be larger, in general.

Finally, one interesting observation is that, in an asymptotical regime of low embedding distortion it is possible to achieve *perfect secrecy* in Shannon's sense, i.e., the information leakage about the secret dither is null, no matter how many observations are available to the attacker. It is easy to realize that this is possible if and only if $I(\mathbf{T};\tilde{\mathbf{Y}}) = 0$ for $N_o = 1$. In the low embedding distortion regime that we are considering, it is possible to find values of α that allow to achieve perfect secrecy. For instance, using self-similar lattice partitions, such value is given by $\alpha_0 = 1 - 1/|\mathcal{M}|$.

2.4. Lattice comparison

The comparison given in this section will be focused on the KMA problem, although it can be extended to the CMA scenario taking into account the comments given in sections 2.2 and 2.3. Recall that, from Eq. (7), the residual entropy is given by the expectation of the log-volume of the feasible region. In order to simplify the analysis, hereinafter we will assume that $\alpha \geq 0.5$, since in this case the feasible region $\mathcal{S}_{N_o}$ can be shown to be always a (modulo-Λ) convex set. Thus, if the problem is properly shifted, the modulo operation can be dropped out from the expressions of the feasible regions; bear in mind that the entropy is invariant to translations, so this simplification does not change the results. Moreover, notice that the conditional pdf of $\mathbf{T}$ does not depend on the specific sequence of messages embedded, as long as the latter is known; this implies that, for the expectation, the message sequence can be assumed to be deterministic. Despite these simplifications, exact analytical evaluation of (7) for arbitrary lattices may be extremely involved; this is why one must resort to Monte-Carlo simulations and bounding techniques. A possible upper bound to the residual entropy can be derived as follows:[‡]

$$\begin{aligned} E[\log(\mathrm{vol}(\mathcal{S}_{N_o}))] &= \int_{\mathcal{V}(\Lambda)^{N_o}} \log(\mathrm{vol}(\mathcal{S}_{N_o}))f(\tilde{\mathbf{y}}_1,\ldots,\tilde{\mathbf{y}}_{N_o})d\tilde{\mathbf{y}}_1,\ldots,d\tilde{\mathbf{y}}_{N_o} \\ &\leq \int_{\mathcal{Z}(\Lambda)^{N_o}} \log(\mathrm{vol}(\mathcal{S}^u_{N_o}))f(\tilde{\mathbf{v}}_1,\ldots,\tilde{\mathbf{v}}_{N_o})d\tilde{\mathbf{v}}_1,\ldots,d\tilde{\mathbf{v}}_{N_o} \\ &\leq \sum_{k=1}^{n}\int_{\Omega_k^{N_o}} \log(w_k)f(\tilde{v}_1^k,\ldots,\tilde{v}_{N_o}^k)d\tilde{v}_1^k,\ldots,d\tilde{v}_{N_o}^k, \end{aligned} \tag{15}$$

where 1) $\mathcal{S}^u_{N_o}$ is an outer bound to $\mathcal{S}_{N_o}$, and is defined as $\mathcal{S}^u_{N_o} \triangleq \bigcap_{i=1}^{N_o}\mathcal{D}_i''$, with $\mathcal{D}_i'' \triangleq \tilde{\mathbf{v}}_i - (1-\alpha)\mathcal{B}(\Lambda)$; 2) $\mathcal{B}(\Lambda)$ is an orthotope that outer bounds $\mathcal{V}(\Lambda)$; 3) $f(\tilde{v}_1^k,\ldots,\tilde{v}_{N_o}^k)$ is the joint pdf of the vector formed by the k-th components of $\tilde{\mathbf{v}}_1,\ldots,\tilde{\mathbf{v}}_{N_o}$; 4) $\Omega_k = [-\mu_k,\mu_k)$ is the support interval for $\tilde{v}_i^k$, $i = 1,\ldots,N_o$; 5) w_k is the volume (length) of the bounding feasible region (interval) in the k-th dimension, given by $w_k = \mathrm{vol}(\bigcap_{i=1}^{N_o}(\tilde{v}_i^k - \Omega_k))$. The last inequality in (15) follows from the fact that, in the case of independent observations with independent components, the residual entropy can be written as the sum of entropies per dimension. The above upper bound is an equality only for the scaled cubic lattice[§] $\Delta\mathbb{Z}^n = (x_1,\ldots,x_n)$, $x_i \in \Delta\mathbb{Z}$, and in particular for the well-known Scalar Costa Scheme (SCS) proposed by Eggers,[4] which uses the scalar lattice $\Delta\mathbb{Z}$ for embedding. Furthermore, for the cubic lattice it is possible to arrive at a closed-form expression: from (15), the residual entropy in the k-th dimension is given by $E[\log(w_k)]$; it can be shown, under the assumption of $\alpha \geq 0.5$, that the residual entropy per dimension is

$$\frac{1}{n}h(\mathbf{T}|\tilde{\mathbf{Y}}_1,\ldots,\tilde{\mathbf{Y}}_{N_o},M_1,\ldots,M_{N_o}) = \log((1-\alpha)\Delta) - H_{N_o} + 1, \tag{16}$$

where $H_{N_o} = \sum_{i=1}^{N_o}\frac{1}{i}$ is the N_o-th harmonic number.

In order to compare the security level offered by different standard lattices, numerical experiments through Monte Carlo simulations were carried out. To provide a fair comparison, the lattices were conveniently scaled so as to present the same embedding distortion per dimension as the cubic lattice $\Delta\mathbb{Z}^n$ with $\Delta = 1$, that is, 1/12. For the low-embedding-distortion regime that we are considering, the embedding distortion in DC-DM is given by $\sigma_e^2 = \alpha^2 P(\Lambda)$, where $P(\Lambda)$ denotes the second order moment per dimension of $\mathcal{V}(\Lambda)$.

[‡] Note that the message sequence is dropped out from the pdf's, for the sake of compactness.

[§] Although the scaling factor need not be the same for each dimension.

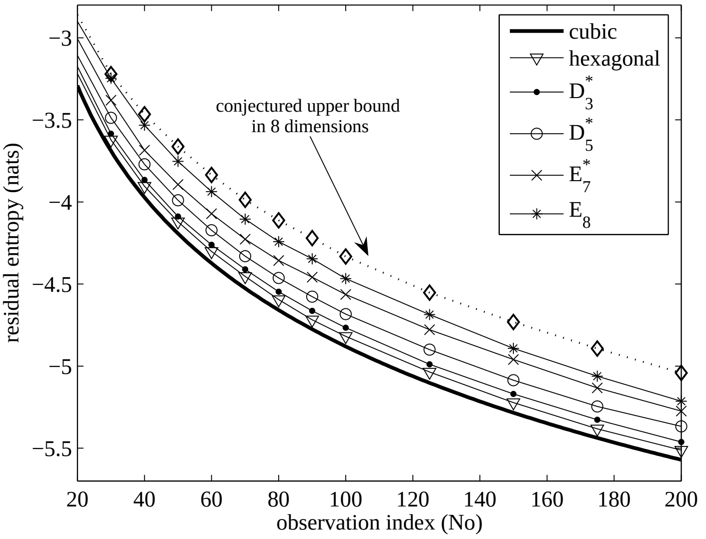

Figure 2. Residual entropies per dimension for different lattices (KMA, $alpha = 0.5$)

For the comparison, we considered the root lattices and their duals[8] (the best known quantizers for $n \leq 8$). Monte-Carlo simulations showed that the lattices that maximize the residual entropy for each n are those with the best source coding properties, i.e., those with the smallest normalized second order moment, $G(\Lambda) \triangleq \frac{P(\Lambda)}{\mathrm{vol}(\mathcal{V}(\Lambda))^{2/n}}$. This could be explained by the fact that, for fixed $P(\Lambda)$, the lattice with the smallest $G(\Lambda)$ provides the maximum volume of $\mathcal{V}(\Lambda)$, and consequently the highest a priori entropy, due to the uniformity of $\mathbf{T}$. Fig. 2 gives a comparison between the residual entropy per dimension using the cubic lattice and that using the best quantizer for each dimensionality. It is a well-known result that the region of $\mathbb{R}^n$ with the smallest normalized second order moment is the sphere, thus we conjecture that it provides an upper bound to the residual entropy, and is indeed so for the cases considered (this is illustrated in Fig. 2 with the result for sphere in 8 dimensions). Unfortunately, the space can not be tessellated with spherical regions (except for $n = 1$), although it was shown in[9] that as n increases there exist lattices whose normalized second order moment tend to that of a sphere. It can be shown that the upper bound (15) to the entropy per dimension with a hypothetical spherical Voronoi cell, for large n, may be approximated as

$$\begin{aligned}\frac{1}{n}h(\mathbf{T}|\tilde{\mathbf{Y}}_1,\ldots,\tilde{\mathbf{Y}}_{N_o},M_1,\ldots,M_{N_o}) &\leq \int_{\Omega_k^{N_o}} \log(w_k) f(\tilde{v}_1^k,\ldots,\tilde{v}_{N_o}^k) d\tilde{v}_1^k,\ldots,d\tilde{v}_{N_o}^k \\ &\approx \log\left(2r - \frac{2\sqrt{2}\cdot r}{\sqrt{2+n}}\cdot\left[\log\left(\frac{N_o}{2}\right)\right]^{\frac{1}{2}}\right) + \log(1-\alpha),\ \text{for } N_o \geq 2, \quad (17)\end{aligned}$$

where $r = \frac{\sqrt{2+n}}{2\sqrt{3}}$ is the radius of the n-dimensional sphere necessary to yield $P(\Lambda) = 1/12$. This result implies that the residual entropy, for a given N_o, is always increasing with n. Finally, we want to note the results given in (16) and (17), as well as those obtained by Monte Carlo integration, are in agreement with the concavity of the mutual information mentioned in Section 2.1.

2.5. Security of DC-DM vs. Costa

Lattice DC-DM schemes are deeply rooted in the theoretical construction developed by Costa.[10] However, the generation of the codebook is totally different in both schemes: whilst $\mathcal{U}$ is structured due to the lattice quantizer, the codebook in Costa's construction is completely random. The purpose of the brief comparison given in the following is to show how much can be gained in terms of security by using a codebook with these characteristics. Due to the lack of space, the security analysis for Costa's scheme is not included in this paper; some of the results that we use here were already published in,[11] whereas others will be published elsewhere.

For the KMA case with $N_o = 1$, it can be shown that

$$I(\mathbf{Y};\mathcal{U}|M) = \frac{n}{2}\log\left[\frac{P+\sigma_X^2}{(1-\alpha)^2\sigma_X^2}\right], \quad h(\mathcal{U}|\mathbf{Y},M) = h(\mathcal{U}) - \frac{n}{2}\log\left[\frac{P+\sigma_X^2}{(1-\alpha)^2\sigma_X^2}\right], \tag{18}$$

where σ_X^2 and P stand for host and watermark power, respectively, and $h(\mathcal{U})$ denotes the entropy of the codebook, given by $h(\mathcal{U}) = \frac{n\cdot|\mathcal{U}|}{2}\log\left[2\pi e(P+\alpha^2\sigma_X^2)\right]$. Although (18) depend on the ratio σ_X^2/P it is easy to see that, for $\sigma_X^2/P \to \infty$, the information leakage for Costa tends to $-n\log(1-\alpha)$, exactly as for DC-DM (see Eq. (9)). Actually, the information leakage in DC-DM also depends on the ratio σ_X^2/P, but the results in sections 2.1 to 2.3 were derived for the asymptotical regime of low embedding distortion (equivalently, *high-rate* quantization), which is that of practical interest in real applications. Nevertheless, the results are radically different if the comparison is made in terms of residual entropy: whereas for DC-DM the entropy of the codebook is bounded by $\log(\text{vol}(\mathcal{V}(\Lambda)))$, the residual entropy in Costa's scheme is unbounded when $\sigma_X^2/P \to \infty$, because of the dependence of $|\mathcal{U}|$ with this ratio. This shows the clear advantage, in terms of security, of the scheme based on random codebooks over the scheme based solely on the dither.

For $N_o = 1$, the CMA and WOA cases are totally equivalent. Assuming that the watermarker is transmitting information at the reliable rate allowed by the channel, we have that

$$I(\mathbf{Y};\mathcal{U}|\text{CM}) = I(\mathbf{Y};\mathcal{U}) = \frac{n}{2}\log\left[\frac{(P+\sigma_X^2)\left(P\sigma_X^2(1-\alpha)^2+\sigma_N^2(P+\alpha^2\sigma_X^2)\right)}{P(P+\sigma_X^2+\sigma_N^2)(1-\alpha)^2\sigma_X^2}\right] = I(\mathbf{Y};\mathcal{U}|M) - I(\mathbf{Y};M|\mathcal{U}), \tag{19}$$

where σ_N^2 is the noise power introduced by the channel. This result is clearly related to those given in (12) and (14) for DC-DM. Here, we can see that the uncertainty about the codebook increases exactly in the same quantity as the reliable transmission rate. Contrary to what happened in DC-DM, there exists no value of α that yield null information leakage; this is only possible when $\sigma_N^2 = 0$, since $I(\mathbf{Y};M|\mathcal{U}) = \infty$.

Finally, for large values of n, if $|\mathcal{U}_m| >> N_o$, which is reasonable since $|\mathcal{U}_j|$ increases also exponentially with n, then the information leakages for $N_o > 1$ in Costa's scheme can be accurately approximated as

$$\begin{aligned}
I(\mathbf{Y}_1,\ldots,\mathbf{Y}_{N_o};\mathcal{U}|\mathbf{M}) &\approx N_o\cdot I(\mathbf{Y}_1;\mathcal{U}|M),\\
I(\mathbf{Y}_1,\ldots,\mathbf{Y}_{N_o};\mathcal{U}|\text{CM}) &\approx N_o\cdot I(\mathbf{Y}_1;\mathcal{U}|\mathbf{M}) - I(\mathbf{Y}_1;M|\mathcal{U}),\\
I(\mathbf{Y}_1,\ldots,\mathbf{Y}_{N_o};\mathcal{U}) &\approx N_o\cdot\left[I(\mathbf{Y}_1;\mathcal{U}|\mathbf{M}) - I(\mathbf{Y}_1;M|\mathcal{U})\right].
\end{aligned}$$

Thus, the information leakages for Costa's scheme increases (approximately) linearly with the number of observations, contrarily to the concave increase in DC-DM. The above approximations are based on the fact that the probability of observing two or more watermarked signals related to the same codeword vanishes exponentially with n, as long as $N_o \leq |\mathcal{U}_m|$.

3. PRACTICAL ALGORITHMS FOR SECRET DITHER ESTIMATION

The algorithms proposed here are mainly intended for the KMA case; the modifications needed to cope with the CMA and WOA cases will be addressed in sections 3.3 and 3.4. These algorithms take advantage of the boundedness of the *feasible regions* $\mathcal{S}_{N_o}$ (which will be also termed *feasible solution sets*) for the dither. Recall that in the KMA scenario, the conditional pdf of $\mathbf{t}$ is uniformly distributed over $\mathcal{S}_{N_o}$ (because we assumed $\mathbf{T} \sim U(\mathcal{V}(\Lambda))$). Thus, the optimal estimator (in an ML sense) simply computes iteratively the series of solution sets $\{\mathcal{S}_k\}$, $k = 1,\ldots,N_o$, and chooses randomly (with equal probability) a vector $\hat{\mathbf{t}} \in \mathcal{S}_{N_o}$. Although intuitively simple, such estimator is not practical in general, since exact computation of the solution sets may be extremely involved for non-trivial lattices, because of the arbitrary shape of the feasible region for $N_o > 1$. Nevertheless, the attacker may not be interested in obtaining the exact $\{\mathcal{S}_k\}$, but may be satisfied with an accurate approximation of the solution sets. This is the approach that will be adopted.

The algorithms that will be considered here work well when the considered sets are convex. As it was said in Section 2.4, the feasible region is a (modulo-Λ) convex set whenever $\alpha \geq 0.5$. If we translate all observations by $-\tilde{\mathbf{y}}_1 + \mathbf{d}_{m_1}$, then the modulo operation is transparent, so the feasible regions for each observation (Eq. (6)) can be now simplified to[¶]

$$\mathcal{D}_i = \tilde{\mathbf{v}}_i + (1-\alpha)\mathcal{V}(\Lambda),\ i = 1, \ldots, N_o, \tag{20}$$

where $\tilde{\mathbf{v}}_i \triangleq (\tilde{\mathbf{y}}_i - \mathbf{d}_{m_i} - \tilde{\mathbf{y}}_1 + \mathbf{d}_{m_1}) \mod \Lambda$, rendering the problem convex.

3.1. Algorithm I: Inner polytope algorithm

Assuming that we know the equations of the hyperplanes that bound the Voronoi region, the feasible solution set can be described through a set of linear inequalities, which in turn describe a polytope in n-dimensional space. Hence, the feasible solution set can be expressed as

$$\mathcal{S}_{N_o} = \left\{ \mathbf{z} \in \mathbb{R}^n : \boldsymbol{\phi}_k^T \mathbf{z} \leq \boldsymbol{\phi}_k^T \tilde{\mathbf{v}}_i + \boldsymbol{\phi}_k^T \mathbf{z}_{0,k} \right\},\ k = 1, \ldots, n_f;\ i = 1, \ldots, N_o, \tag{21}$$

where $\boldsymbol{\phi}_k$ is a vector outward-pointing and normal to the k-th facet of $\mathcal{V}(\Lambda)$, and $\mathbf{z}_{0,k}$ is a point in such facet. We are interested in computing an approximation of the feasible region. For such an approximation to be valid, it must outer bound $\mathcal{S}_{N_o}$ (as tightly as possible), since we do not want to discard any point in $\mathcal{S}_{N_o}$ a priori; furthermore, it would be desirable that the approximate region is easy to describe, thus an ellipsoid is a reasonable choice. Since the problem of finding the ellipsoid of minimum volume that contains $\mathcal{S}_{N_o}$ is ill-posed, we follow the alternative albeit suboptimal approach of finding the maximum volume ellipsoid $\mathcal{E}(\boldsymbol{\theta}, \mathbf{P})$ contained in $\mathcal{S}_{N_o}$, and scaling it by a factor of n around its center. This way, the resulting ellipsoid is guaranteed to bound the polytope $\mathcal{S}_{N_o}$.[12] The computation of the center $\boldsymbol{\theta}^*$ and the positive definite matrix $\mathbf{P}^*$ of such ellipsoid can be written as a convex minimization problem with second order cone constraints[12]:

$$\begin{aligned} (\hat{\boldsymbol{\theta}}, \hat{\mathbf{P}}) = \arg\min_{\boldsymbol{\theta}, \mathbf{P}} \quad & \log\det(\mathbf{P}^{-1/2}) \\ \text{subject to} \quad & ||\mathbf{P}^{1/2}\boldsymbol{\phi}_k|| \leq \boldsymbol{\phi}_k^T \tilde{\mathbf{v}}_i + \boldsymbol{\phi}_k^T \mathbf{z}_{0,i} - \boldsymbol{\phi}_k^T \boldsymbol{\theta}, \forall\ k = 1, \ldots, n_f;\ i = 1, \ldots, N_o. \end{aligned} \tag{22}$$

As it will be checked in Section 3.5, this approach yields tight approximations to $\mathcal{S}_{N_o}$, but it presents an obvious drawback: the potential complexity of the minimization problem arising from the huge number of constraints imposed by large n and N_o.

3.2. Algorithm II: Set membership estimator

Set-membership estimators (SME)[13] are commonly found in the Automatic Control literature. SME algorithms differ from classical estimation algorithms in two basic aspects: 1) the only assumption about the noise present in the observations is that it is additive and bounded, and 2) they construct a series of solution sets that enclose the region of $\mathbb{R}^n$ where the parameter to be estimated lies in.

SME algorithms can be straightforwardly applied to our problem by slightly modifying the description of the feasible region for each observation: in this case, we need to parameterize $\mathcal{D}_i$ as the intersection of a finite number of parallel hyperplanes. Assuming that the Voronoi cell of the considered lattice is composed of n_f pairwise parallel facets (see Fig. 3-(a)),[‖] the feasible solution set for the i-th observation can be specified by a matrix $\boldsymbol{\Phi}_{n \times n_f/2}$, and a vector $\boldsymbol{\gamma}_{n_f/2 \times 1}$ such that $\mathcal{D}_i = \bigcap_{j=1}^{n_f/2} \mathcal{F}_{i,j}$, where

$$\mathcal{F}_{i,j} = \{ \mathbf{z} \in \mathbb{R}^n : |\tilde{\mathbf{v}}_i^T \boldsymbol{\phi}_j - \mathbf{z}^T \boldsymbol{\phi}_j| \leq \gamma_j \}, \tag{23}$$

being $\boldsymbol{\phi}_j$ the j-th column of $\boldsymbol{\Phi}$, and γ_j the j-th element of $\boldsymbol{\gamma}$, computed as $\boldsymbol{\phi}_j^T \mathbf{z}_{0,k}$. Hence, the series of solution sets $\{\mathcal{S}_k\}$ is given by

$$\mathcal{S}_k = \bigcap_{i=1}^{k} \mathcal{D}_i = \bigcap_{i=1}^{k} \bigcap_{j=1}^{n_f/2} \mathcal{F}_{i,j},\ k = 1, \ldots, N_o. \tag{24}$$

[¶] Obviously, the offset $-\tilde{\mathbf{y}}_1 - \mathbf{d}_{m_1}$ must be removed from the final estimate.

[‖] Should this not be true, the problem can still be described in a similar manner.

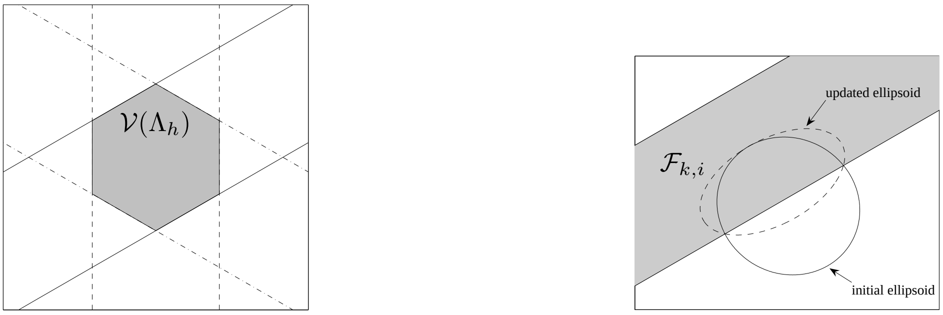

(a) (b)

Figure 3. (a) Voronoi region of the hexagonal lattice delimited by three pairs of parallel hyperplanes. (b) Intersection between an ellipsoid and a pair of hyperplanes.

The common approximation (although other approaches can be found in the literature) in SME algorithms for simplifying the computation of (24) is to choose an ellipsoid $\mathcal{E}(\hat{\boldsymbol{\theta}}_k, \hat{\mathbf{P}}_k)$ such that $\mathcal{S}_k \subseteq \mathcal{E}(\hat{\boldsymbol{\theta}}_k, \hat{\mathbf{P}}_k)$. Thus, computation of the $(k+1)$-th solution set amounts to obtaining an ellipsoid $\mathcal{E}(\hat{\boldsymbol{\theta}}_{k+1}, \hat{\mathbf{P}}_{k+1}) \supseteq \mathcal{E}(\hat{\boldsymbol{\theta}}_k, \hat{\mathbf{P}}_k) \cap \mathcal{D}_k$. A further simplification consists of computing such ellipsoid iteratively:

1. First, make $\mathcal{E}(\mathbf{c}_0, \mathbf{B}_0) = \mathcal{E}(\hat{\boldsymbol{\theta}}_k, \hat{\mathbf{P}}_k)$
2. Compute $\mathcal{E}(\mathbf{c}_{i+1}, \mathbf{B}_{i+1}) \supseteq \mathcal{E}(\mathbf{c}_i, \mathbf{B}_i) \cap \mathcal{F}_{k,i+1}, \ i = 0, \dots, n_f/2 - 1$
3. Finally, make $\mathcal{E}(\hat{\boldsymbol{\theta}}_{k+1}, \hat{\mathbf{P}}_{k+1}) = \mathcal{E}(\mathbf{c}_{n_f/2}, \mathbf{B}_{n_f/2})$

This way, in each step we are intersecting one ellipsoid with one set $\mathcal{F}_{k,i}$, as it is depicted in Figure 3-(b). Clearly, we are interested in finding the ellipsoid with minimum volume that contains such intersection. The SME algorithm that addresses such minimization problem is the so-called *Optimal Volume Ellipsoid* (OVE) algorithm.[14] The minimization problem has analytic solution and it reads as

$$\mathbf{c}_{i+1}^* = \mathbf{c}_i + \frac{\tau_i \mathbf{B}_i \boldsymbol{\phi}_i}{\left(\boldsymbol{\phi}_i^T \mathbf{B}_i \boldsymbol{\phi}_i\right)^{1/2}}, \quad \mathbf{B}_{i+1}^* = \delta_i \left(\mathbf{B}_i - \sigma_i \frac{\mathbf{B}_i \boldsymbol{\phi}_i \boldsymbol{\phi}_i^T \mathbf{B}_i}{\boldsymbol{\phi}_i^T \mathbf{B}_i \boldsymbol{\phi}_i}\right),$$

where τ_i, σ_i, δ_i are variables that depend on the observation $\tilde{\mathbf{v}}_k$, the current ellipsoid $\mathcal{E}(\mathbf{c}_i, \mathbf{B}_i)$ and $\mathcal{F}_{k,i+1}$ (details about their calculation can be found in[14]), and finally $\boldsymbol{\phi}_i$ is the i-th column of matrix $\boldsymbol{\Phi}$.

One interesting feature of the approach just described, and common to all SME algorithms, is that further refinements on the solution set are possible by recirculating the observed data, i.e., by feeding to the system the same set of observations repeatedly. This is possible since the resulting bounding ellipsoid in the i-th iteration depends both on the $(i-1)$-th bounding ellipsoid and the i-th observation. This feature, as will be checked in Section 3.5, allows to achieve performance similar to that of the above *inner polytope* algorithm.

3.3. Estimation in the CMA scenario

The CMA scenario implies minor changes to the estimation algorithms devised above for the KMA case. Actually, estimation in the CMA case could be performed as

1. Assume that the sequence of observation is watermarked with message $\mathbf{m} \in \mathcal{M}$,
2. Perform estimation as in the KMA case,
3. Once $\hat{\mathcal{S}}_{N_o}$ has been obtained, compute the approximate feasible region $\hat{\mathcal{S}}_{N_o}^{CMA}$ as in Eq. (11).
4. If the resulting feasible regions $(\hat{\mathcal{S}}_{N_o} + \mathbf{d}_m)$ overlap, then the dither is contained in their intersection with higher probability (provided that $\mathbf{T}$ is uniformly distributed over $\mathcal{V}(\Lambda)$); otherwise, the centroid is equally likely in any of the feasible regions.

3.4. Estimation in the WOA scenario

The WOA scenario introduces an additional source of uncertainty that invalidates the straightforward application of the algorithms proposed in sections 3.1 and 3.2, which make explicit use of knowledge of the embedded messages. A possible solution would be to consider all the possible sequences of embedded messages, so as to transform the WOA problem into $|\mathcal{M}|^{N_o}$ parallel KMA problems. Obviously, this *brute-force* approach is not practical due to the huge number of possible message sequences, which makes the problem unmanageable. Fortunately, the a priori search space can be reduced considerably if we consider the posterior probability of the message sequences, or in other words: given a sequence of observations, certain message sequences have null or negligible probability of occurrence. Assuming equiprobable and mutually independent messages, it can be readily shown that ML selection of the message is tantamount to

$$(\hat{m}_1, \ldots, \hat{m}_{N_o}) = \arg \max_{m_1, \ldots, m_{N_o}} f(\tilde{\mathbf{y}}_1, \ldots, \tilde{\mathbf{y}}_{N_o} | m_1, \ldots, m_{N_o}), \tag{25}$$

It can be shown that $f(\tilde{\mathbf{y}}_1, \ldots, \tilde{\mathbf{y}}_{N_o} | m_1, \ldots, m_{N_o})$ can be factorized as

$$f(\tilde{\mathbf{y}}_1, \ldots, \tilde{\mathbf{y}}_{N_o} | m_1, \ldots, m_{N_o}) = \prod_{k=1}^{N_o} \int f(\tilde{\mathbf{y}}_k | m_k, \mathbf{t}) \cdot f(\mathbf{t} | \tilde{\mathbf{y}}_1, \ldots, \tilde{\mathbf{y}}_{k-1}, m_1, \ldots, m_{k-1}) d\mathbf{t}, \tag{26}$$

where the conditional pdf of the secret dither is given by (5), and the term corresponding to $k = 1$ is simply $f(\tilde{\mathbf{y}}_1 | m_1) = f(\tilde{\mathbf{y}_1})$. Computation of the integral terms in (26) is straightforward since the conditional pdf of the secret dither is uniform; hence, such terms are proportional to the volume of the intersection between $\mathcal{D}_k$ and $\mathcal{S}_{k-1}$. Formally, each factor of (26) can be computed as

$$f_k = \text{vol}\left(\mathcal{D}_k \cap \mathcal{S}_{k-1}\right) \cdot \left((1-\alpha)^n \text{vol}(\mathcal{V}(\Lambda)) \cdot \text{vol}(\mathcal{S}_{k-1})\right)^{-1}. \tag{27}$$

The message selection rule just defined is optimal, but the use of SME estimators will introduce some loss of optimality due to the ellipsoidal approximation. With the ML approach, SME estimation in the WOA scenario can be thought of as a tree search where each node at level k is the hypothesized message for the k-th observation, and whose probability is given by the k-th factor of (26). The probability of each path in the tree is given by the product of the corresponding factors. Once the final node level is reached, the probability of each path is obtained, so we can either pick the path with the highest probability, or simply discard those paths with low probabilities (discarding can be made also in partial paths, to simplify the tree search). Bear in mind that the number of feasible paths with non-null probability, or equivalently, the number of branches that must be considered in the tree search, increases for decreasing values of α, making more difficult the estimation. Finally, an interesting byproduct of the proposed approach is that it allows to obtain an estimate of the transmitted message.

3.5. Experimental results

This section provides a comparison of the practical performance of the different estimators proposed in Section 3. The results are given only for the KMA scenario; nevertheless, since our approach decomposes both CMA and WOA cases in several KMA instances, it is clear that the performance of the estimators in these scenarios will be determined by that of the KMA case. The minimization problem (22) was solved using the optimization packages YALMIP[15] and SeDuMi[16] for Matlab.

Figures 4-(a) and 4-(b) show the performance of the different estimators when embedding lattices are the hexagonal and E_8, respectively. For the description of its Voronoi regions, the reader can refer to.[17] The value represented in the plots is the *empirical* residual entropy obtained by averaging the outputs of algorithms I and II, assuming that the conditional dither is uniformly distributed in the estimated feasible region. This value permits to quantify directly the loss in optimality incurred by the proposed estimators, by comparing to the theoretical value of the residual entropy. It can be observed that, although the inner polytope algorithm provides the best performance, the loss of optimality of the SME (OVE) algorithm can be compensated for by increasing the number of recirculations. This feature makes SME algorithms appealing for their application to high-dimensional lattices, because of their low complexity.

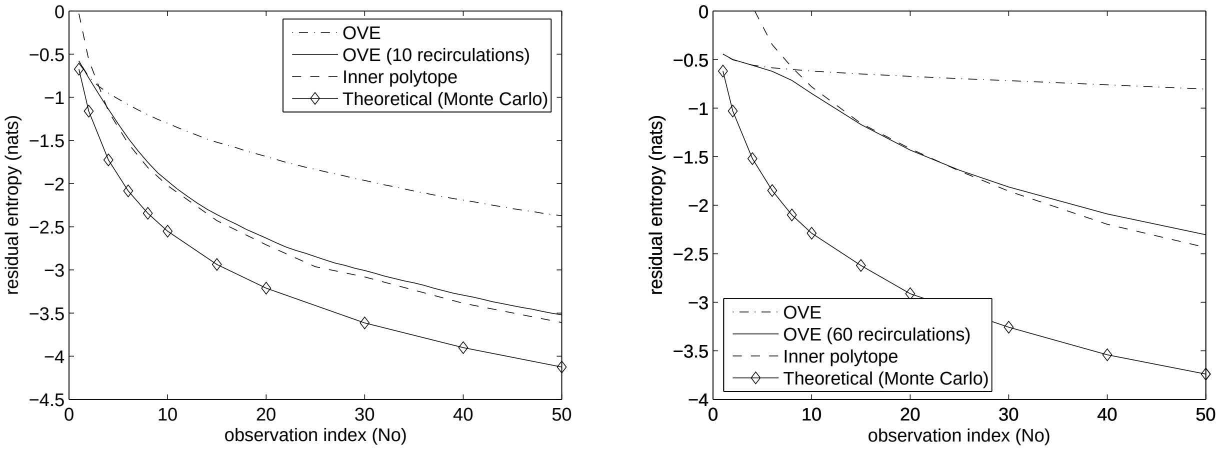

(a)

(b)

Figure 4. Performance comparison for the proposed estimators, with hexagonal lattice (a) and E_8 (b). (KMA, $\alpha = 0.5$).

4. CONCLUSIONS AND FINAL REMARKS

The comparison given in Section 2.5 shows that the security weaknesses of DC-DM are not inherent to Costa's schemes, but due to the structure imposed to the codebook. In fact, the information leakages about the secret dither may by important, depending on the considered scenario (KMA, CMA, WOA). However, perfect secrecy can be still (theoretically) achieved in the WOA scenario, although this is not strictly true in practice, since the pdf's of the watermarked signals are not exactly uniform; therefore, a small information leakage still may be exploited.

As it was seen in Section 2.4, the security level of DC-DM is dependent of the specific embedding lattice considered; in fact, security can be improved by increasing the dimensionality and choosing the appropriate lattice, but it comes at the price of higher computational complexity in embedding and decoding. Furthermore, the gain for small n is rather limited.

On the other hand, the performance of the proposed estimation algorithms show that attacks are feasible with manageable complexity. In particular, the SME estimator poses several advantages:

- Affordable complexity: the algorithm is iterative, and each iteration involves elementary operations. in each iteration $n(n+1)/2 + n$ (number of free parameters of the symmetric matrix + center) parameters must be estimated, where n is the dimensionality of the (lattice) problem.
- The asymptotic convergence of SME algorithms to the true parameters has been proved.
- Increasing the number of recirculations allows to compensate for the loss of optimality. Nevertheless, note that they are optimal in case of scalar embedding lattices.

Finally, we would like to say that, despite the results presented here, perspectives for DC-DM are not so hopeless: the security level of DC-DM can still be increased by maintaining the structured codebook, which is desirable from a practical point of view. The solution is to increase the degree of uncertainty about the codebook, introducing for instance rotations to the embedding lattice and/or permutations to the host prior embedding. Obviously, the price that must be paid for these approaches, apart from complexity, is the increase in the entropy of the key.

ACKNOWLEDGMENTS

The authors are indebted to Dr. Teddy Furon for his valuable comments and suggestions.

REFERENCES

1. F. Cayre, C. Fontaine, and T. Furon, "Watermarking security: Theory and practice," *IEEE Transactions on Signal Processing* **53**, pp. 3976–3987, October 2005.
2. P. Comesaña, L. Pérez-Freire, and F. Pérez-González, "Fundamentals of data hiding security and their application to spread-spectrum analysis," in *7th Information Hiding Workshop, IH05*, M. Barni, J. Herrera-Joancomartí, S. Katzenbeisser, and F. Pérez-González, eds., *Lecture Notes in Computer Science* **3727**, pp. 146–160, Springer Verlag, (Barcelona, Spain), June 2005.
3. 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**, pp. 1423–1443, May 2001.
4. J. J. Eggers, R. Bäuml, R. Tzschoppe, and B. Girod, "Scalar Costa Scheme for information embedding," *IEEE Transactions on Signal Processing* **51**, pp. 1003–1019, April 2003. Special Issue on Signal Processing for Data Hiding in Digital Media and Secure Content Delivery.
5. C. E. Shannon, "Communication theory of secrecy systems," *Bell system technical journal* **28**, pp. 656–715, October 1949.
6. P. Bas and J. Hurri, "Security of DM quantization watermarking schemes: a practical study for digital images," in *Fourth International Workshop on Digital Watermarking*, M. Barni, I. Cox, T. Kalker, and H. J. Kim, eds., **3710**, pp. 186–200, Springer, (Siena, Italy), September 2005.
7. U. Erez and R. Zamir, "Achieving $\frac{1}{2}$log(1+SNR) over the Additive White Gaussian Noise Channel with Lattice Encoding and Decoding," *IEEE Transactions on Information Theory* **50**, pp. 2293–2314, October 2004.
8. J. H. Conway and N. J. A. Sloane, *Sphere packings, lattices and groups*, vol. 290 of *a series of comprehensive studies in Mathematics*, Springer-Verlag, New York, third ed., 1999.
9. R. Zamir and M. Feder, "On lattice quantization noise," *IEEE Transactions on Information Theory* **42**, pp. 1152–1159, July 1996.
10. M. H. M. Costa, "Writing on dirty paper," *IEEE Transactions on Information Theory* **29**, pp. 439–441, May 1983.
11. L. Pérez-Freire, P. Comesaña, and F. Pérez-González, "Information-theoretic analysis of security in side-informed data hiding," in *7th Information Hiding Workshop, IH05*, M. Barni, J. Herrera-Joancomartí, S. Katzenbeisser, and F. Pérez-González, eds., *Lecture Notes in Computer Science* **3727**, pp. 131–145, Springer Verlag, (Barcelona, Spain), June 2005.
12. S. Boyd and L. Vandenberghe, *Convex optimization*, SIAM Studies in Applied Mathematics, Cambridge University Press, Cambridge, UK, 2004.
13. M. Milanese and A. Vicino, "Optimal estimation theory for dynamic systems with set membership uncertainty: an overview," *Automatica* **27**(6), pp. 997–1009, 1991.
14. M. F. Cheung, S. Yurkovich, and K. M. Passino, "An optimal volume ellipsoid algorithm for parameter set estimation," *IEEE Transactions on Automatic Control* **38**, pp. 1292–1296, August 1993.
15. J. Löfberg, "YALMIP: A toolbox for modeling and optimization in MATLAB," in *Proceedings of the CACSD Conference*, (Taipei, Taiwan), 2004. Available from `http://control.ee.ethz.ch/~joloef/yalmip.php`.
16. J. F. Sturm, "Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones.," *Optimization methods and software* **11-12**(1-4), pp. 625–653, 1999. Version 1.1 available from `http://sedumi.mcmaster.ca/`.
17. J. H. Conway and N. J. A. Sloane, "Voronoi regions of lattices, second moments of polytopes, and quantization," *IEEE Transactions on Information Theory* **28**, pp. 211–226, March 1982.