Nothing But a Kiss: A Novel and Accurate Approach to Assessing the Performance of Multidimensional Distortion-Compensated Dither Modulation[⋆]

Fernando Pérez-González and Félix Balado

Dept. Tecnologías de las Comunicaciones, ETSI Telecom.
University of Vigo, E-36200 Vigo, Spain
{fperez,fiz}@tsc.uvigo.es

Abstract. A considerable amount of attention has been lately paid to a number of data hiding methods based on quantization, seeking to achieve in practice the results predicted by Costa for a channel with side information at the encoder. In this paper we analyze a multidimensional extension of the implementation of Costa's result known as DC-QIM. The presented analysis is based on measuring the probabilities of decoding error in the presence of two important kinds of additive channel distortions. DC-QIM is able to achieve a probability of decoding error of less than 1e-02 for 0 dB of watermark to noise ratio and only 20 samples per hidden bit. Empirical results supporting our analysis are given both for synthetic data and real images.

1 Introduction

Quantization methods have been used used for information hiding purposes since the early times of research in watermarking. However, it was only very recently when the idea was reviewed from a solid theoretical perspective in the form of a data hiding scheme known as Quantization Index Modulation (QIM) that was proposed by Chen and Wornell in [1]. QIM hides information by constructing a set of vector quantizers $\mathbf{Q}_{\mathbf{b}}(\cdot)$, each representing a different codeword $\mathbf{b}$. Then, given a host vector $\mathbf{x}$ and an information codeword $\mathbf{b}$, the embedder constructs the watermarked vector $\mathbf{y}$ by simply quantizing $\mathbf{x}$ with $\mathbf{Q}_{\mathbf{b}}(\cdot)$, i.e. $\mathbf{y} = \mathbf{Q}_{\mathbf{b}}(\mathbf{x})$.

Later, a crucial connection with an old paper by Costa [2] was made by some researchers [3,4] realizing that it was possible to implement a particular capacity-achieving scheme by following Costa's guidelines. The main improvement over QIM was made by adding back a fraction of the quantization error to the quantized signal. This compensation scheme gave rise to what Chen and Wornell called Distortion Compensated QIM (DC-QIM).

The original proposal of DC-QIM in [3] provides a general framework that can be readily adapted in order to use many existing structured vector quantizers. In particular, close attention was paid to a special case of QIM called Dithered Modulation (DM), which has the advantage of a very simple implementation when uniform quantizers are used.

[⋆] Work partially supported by the *Xunta de Galicia* under projects PGIDT01 PX132204PM and PGIDT02 PXIC32205PN, the European project Certimark (Certification Watermarking Technologies), IST-1999-10987, and the CYCIT project AMULET, reference TIC2001-3697-C03-01.

It was soon realized that QIM and its improvement DC-QIM could benefit from multidimensional extensions. Previously, there had been some efforts in the use of multidimensional quantization lattices like those of Chae et al. [5,6], who used D_4 and other quantization lattices [7] to implement several image and video data hiding algorithms. Actually, this scheme was a sort of multidimensional QIM, but nevertheless no performance analysis was given. Later on, and departing from their Scalar Costa Scheme (SCS) method —formally equivalent to DC-QIM—, Eggers and Girod [4] extended the unidimensional approach to a bidimensional one by using the optimal sphere-packing hexagonal lattice and ternary information symbols. They empirically showed that there is an implicit performance gain in the use of a new dimension. Finally, Brunk [8] considered the problem of estimating the capacity of very high dimensional QIM and DC-QIM schemes, trying to find the limiting capabilities when adding further dimensions. Only small attacking distortions were allowed for the used approximations to hold.

Chen and Wornell only sketched very roughly the performance of QIM as measured by the probability of decoding error in front of additive Gaussian noise [3]. Unfortunately, as showed in [9], approaches to assessing performance based on the so-called union bound fail for high distortion levels. The number of "kissing spheres" in the quantizer increases exponentially with dimensionality, thus largely overestimating the actual probability of decoding error. Recently, we have been able to produce very tight approximations for the probability of decoding error in the case of Dither Modulation (DM) data hiding and theoretically support its good performance.

To the authors' knowledge there does not exist a formal analysis for Multidimensional Distortion-Compensated QIM methods in terms of their probability of decoding error. In this paper, we take an important step in this direction by producing very accurate bounds for this probability in the particular case of DC-DM method. We will show how the multidimensional problem can be transformed into a one-dimensional one by adapting a technique recently discovered in the digital communications area that will allow us to reduce the number of kissing spheres that need be considered, down to a single one.

2 Problem formulation

Let $\mathbf{x}$ be a vector containing the samples of the host signal (henceforth an image) that will convey the hidden information. Before the encoding stage, a perceptual mask vector $\boldsymbol{\alpha}$ is computed from $\mathbf{x}$ in the domain where the samples are defined (e.g., spatial, DCT, etc), taking into account the characteristics of the human visual system (HVS). At each sample k, the maximum variance of the modification in the host signal sample that renders it perceptually unchanged is proportional to $\alpha^2[k]$.

In order to keep the discussion simple, we will assume that we want to hide only one binary digit of information that we consider to be mapped to an antipodal symbol, $b \in \{\pm 1\}$. In order to hide this symbol, a set $\mathcal{S} = \{k_1, \ldots, k_L\}$ of L key-dependent pseudorandomly chosen indices for the samples of $\mathbf{x}$ is employed. This way of choosing the indices allows us to assume statistical independence between the host signal samples. Also note that in the subsequent statistical analysis, if L is large, it is enough to consider the embedding of one bit: thanks to the pseudorandom selection of the host signal samples, if additional bits were to be encoded they would face, in average, exactly the same channel. A host signal of size N would allow us to hide $\lfloor N/L \rfloor$ bits in the same way.

The watermark $\mathbf{w}$ is produced from the desired binary information symbol b by using a certain function, $\mathbf{w} = g_K(\mathbf{x}, \boldsymbol{\alpha}, b)$, that we will detail in Sect. 3. Without loss of generality we will write the watermarked signal as the addition

$$\mathbf{y} = \mathbf{x} + \mathbf{w}\,. \tag{1}$$

An important issue when performing a rigorous analysis lies in the election of a proper measure for the so-called *embedding distortion*. A quite often used possibility is the global Mean-Squared Error (MSE) distortion that in the single-bit case is defined as

$$D_w = \frac{1}{L} \sum_{k \in \mathcal{S}} \mathrm{Var}\{w[k]\} \tag{2}$$

where $w[k]$ is a random process representing the watermark [1]. The MSE, appropriate as it is for measuring the total power devoted to the watermark, should be handled with extreme care if one wants to relate it to visibility constraints. In fact, bounding just the MSE seems to be inadequate for the data hiding problem, since it is not well-matched to the characteristics of the HVS. It is widely recognized that the masking phenomena affecting the HVS and exploited for the invisible embedment of information respond to local effects. All existing approaches to modeling distortions which are unnoticeable to the HVS take into account this fact, be it the Just Noticeable Distortion function (JND) [10] in the spatial domain, the Noise Visibility Function (NVF) [11] applicable to different domains or equivalent models in other domains of interest like the DCT [12]. Then, the main drawback encountered when posing an MSE constraint is that unacceptably high local distortions (from the perceptual perspective) could be globally compensated with very small local distortions in order to meet the established restriction. An alternative consisting in a weighted MSE is discussed in [13], but it should be clear that it suffers from the same global compensation effect.

In view of the previous discussion it seems reasonable to restrict the local variance of the watermark so that global compensations are not possible and at the same time perceptual weighting is taken into account. This is achieved by means of the following set of constraints:

$$\mathrm{Var}\{w[k]\} \leq c^2 \cdot \alpha^2[k], \text{ for all } k \in \mathcal{S} \tag{3}$$

where c is a *visibility* constant.[2] If the samples $w[k]$ are such that their variances take their extremal values in (3), it is immediate to write

$$D_w = \frac{c^2}{L} \sum_{k \in \mathcal{S}} \alpha^2[k] \tag{4}$$

[1] This randomness may be due to the way in which the watermark depends on the host image, as happens for instance in the DC-QIM method, or on a pseudorandom sequence, as occurs in spread spectrum based schemes.

[2] Note that our definition of $\alpha[k]$ differs from others in the literature (e.g., that in [14]) in that the visibility constant has been taken away. Of course, this is just a matter of convenience and does not alter the final results.

Therefore, it is important to note that simultaneously meeting the constraints in (3) automatically leads to a bound in the MSE, but the converse is not true, unless an extremely large value of the visibility parameter c is considered. This cannot occur if some structure is imposed on the watermark, such as $w[k] = c \cdot \alpha[k]$, for all $k \in \mathcal{S}$.

2.1 Channel characterization

Before arriving to the receiver we assume that the watermarked signal undergoes an additive probabilistic channel independent of $\mathbf{x}$, yielding a received signal $\mathbf{z} = \mathbf{y} + \mathbf{n}$. This channel distortion models certain attacking operations. By virtue of the pseudorandom choice of the indices in $\mathcal{S}$ we may assume for the multidimensional case that the samples in $\mathbf{n}$ are also mutually independent, with diagonal covariance matrix $\Gamma = \text{diag}(\sigma_n^2[k_1], \ldots, \sigma_n^2[k_L])$.

The *channel distortion* D_c is defined in a similar fashion as the embedding distortion, i.e.,

$$D_c = \frac{1}{L} \sum_{k \in \mathcal{S}} \sigma_n^2[k] \tag{5}$$

Then, it will be useful to introduce the square-root ratio

$$\xi \triangleq \sqrt{D_w / D_c} \tag{6}$$

that relates the power of embedding and channel distortions. In addition, we define the *watermark-to-noise ratio* as $\text{WNR} = 20 \log_{10} \xi$.

As before, if perceptual shaping is to be introduced in the noise distribution[3], a simple constraint on D_c in (5) will not be enough and rather a set of inequalities like that in (3) will become necessary. For the purposes of this paper, we will assume that the channel distortion will be limited by that resulting during the embedding procedure, i.e. $D_c \leq D_w$. This choice is justified whenever the attacker wants to produce an image with at least an equivalent quality as the watermarked image. In fact, this is the basis for attacking algorithms such as perceptual remodulation [15]. Needless to say, the previous argument cannot be repeatedly invoked over a sequence of attacks because, eventually, the invisibility conditions would be violated. A much less conservative approach is taken for instance in [16], where $D_c \geq D_w$ under the reasoning that restoration of the original signal must be included in the range of possible modifications caused by the attacking distortion. Last, in [4] both ranges are considered. In any case, it is worth saying that the novel methodology here presented can be adapted with little modification for any range of WNR.

Finally, regarding the probability distribution function (pdf) of the distortion, we will consider here two simple but illustrative cases: Gaussian and uniform. Gaussian channels have been extensively used in the previous work on this topic, see [3,4,16]. On the one hand, the Gaussian channel gives upper bounds to capacity and the optimal attack under global MSE distortion constraints. On the other hand, Gaussian distributions often appear with regard to unintentional attacks. As for the uniform pdf, we have chosen it for its simplicity that leads

[3] Note that an attacker can compute the perceptual mask himself and shape the noise accordingly by making $\boldsymbol{\sigma}_n$ proportional to $\boldsymbol{\alpha}$, but this will unlikely occur for unintentional attacks, so the two cases are possible.

to tractable analytical expressions of performance. Note that by no means we intend to consider here the whole gamut of possible attacks and that the question of which noise channel distribution is the most harmful remains open.

3 Multidimensional Distortion-Compensated Dither Modulation

In this section we discuss the structure of multidimensional QIM, in which each dimension is separately quantized. As mentioned in the Introduction, we solely consider the QIM implementation by means of uniform dithered quantizers, commonly called dithered modulation (DM). Also, we investigate the gain that the use of distortion compensation may produce, as proposed by [3]. For these reasons we call the presented scheme "multidimensional distortion-compensated dither modulation" (DC-DM). In order to simplify the subsequent discussion, we will assume that the samples of the perceptual mask are constant, i.e., $\alpha[k] = \alpha$, for all $k \in \mathcal{S}$. Wherever it is appropriate, we will give indications on how to adapt the method and its analysis to the more realistic case in which the samples of perceptual mask are variant.

In DC-DM each binary information symbol b is hidden by using a L-dimensional uniform quantizer $\mathbf{Q}_b(\cdot)$ on the host image, obtaining the watermark as

$$\mathbf{w} = \nu\,\mathbf{e}\,, \tag{7}$$

i.e. the L-dimensional quantization error $\mathbf{e} \triangleq \mathbf{Q}_b(\mathbf{x}) - \mathbf{x}$ weighted by an optimizable constant ν, $0 \leq \nu \leq 1$. Considering (1) this is equivalent to choosing $\mathbf{y}$ as

$$\mathbf{y} = \mathbf{Q}_b(\mathbf{x}) - (1-\nu)\mathbf{e} \tag{8}$$

Observe that when $\nu = 1$ we have multidimensional uncompensated DM as a particular case. The uniform quantizers $\mathbf{Q}_{-1}(\cdot)$ and $\mathbf{Q}_1(\cdot)$ are such that the corresponding centroids are the points in the lattices

$$\Lambda_{-1} = 2\Delta\mathbb{Z}^L + \mathbf{d} \tag{9}$$

$$\Lambda_1 = 2\Delta\mathbb{Z}^L + \mathbf{d} + \Delta(1,\ldots,1)^T \tag{10}$$

with $\mathbf{d}$ an arbitrary vector that may be key-dependent so as to introduce an additional degree of uncertainty. Since the presence of a known offset $\mathbf{d}$ in the above description of the lattices does not alter the final results, we will assume from now on that $\mathbf{d} = \mathbf{0} \triangleq (0,\ldots,0)^T$. Notice that $\Lambda_{-1} \cup \Lambda_1 \cong D_L^*$, i.e. the dual of the checkered or face centered cubic root lattice [7].

As for the statistical modeling of the watermark, note that if the quantization step is small enough we can consider each dimension of the quantization error $\mathbf{e}$ to be uniformly distributed over an L-dimensional cube with edge size 2Δ, centered at the corresponding quantization centroid. From (7) this in turn implies that the watermark is also uniformly distributed in a cube with edge size $2\nu\Delta$. Thus, it is immediate to write the embedding distortion in (2) as $D_w = \nu^2\Delta^2/3$.

3.1 Decoding and decision regions

Now let $\mathbf{z} = \mathbf{y} + \mathbf{n}$ be the watermarked image that has been corrupted by a noise vector $\mathbf{n}$. Given $\mathbf{z}$, decoding is accomplished by using a minimum Euclidean distance decoder

$$\hat{b} = \arg \min_{-1,1} \|\mathbf{z} - \mathbf{Q}_b(\mathbf{z})\|^2 \tag{11}$$

The decision regions associated to $\hat{b} = -1$ and $\hat{b} = 1$ are denoted by respectively $\mathcal{R}_{-1}$ and $\mathcal{R}_1$. In the sequel we will also find useful to identify the decision regions associated to each of the centroids in the lattices Λ_{-1} and Λ_1. To that end, let $\mathbf{c} \in \Lambda_{-1} \bigcup \Lambda_1$ be any centroid, then we will denote by $\mathcal{S}_{\mathbf{c}}$ the Voronoi cell associated to $\mathbf{c}$, i.e.,

$$\mathcal{S}_{\mathbf{c}} \triangleq \left\{ \mathbf{z} \in \mathbb{R}^L | \quad \|\mathbf{z} - \mathbf{c}\|^2 \leq \|\mathbf{z} - \mathbf{c}'\|^2, \, \forall \mathbf{c}' \in \Lambda_{-1} \bigcup \Lambda_1 \right\} \tag{12}$$

It follows immediately from definition (12) that

$$\mathcal{R}_b \; = \bigcup_{\mathbf{c} \in \Lambda_b} \mathcal{S}_{\mathbf{c}}, \quad b = \{-1, 1\} \tag{13}$$

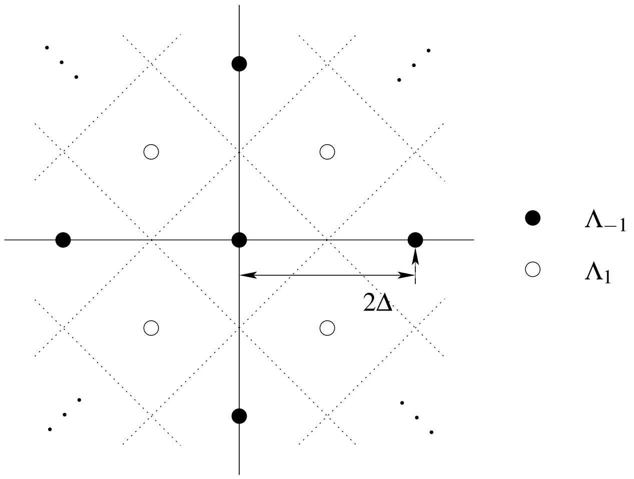

Fig. 1. Centroids and decision regions $(L = 2)$.

The centroids and decision regions $\mathcal{R}_{-1}$ and $\mathcal{R}_1$ for the case $L = 2$ are depicted in Fig. 1. The L-dimensional Voronoi cells $\mathcal{S}_{\mathbf{c}}$ are generalized truncated octahedra [7]. We will find useful to denote by $\mathcal{T}_{\mathbf{0}}$ the generalized octahedron that contains the origin and is limited by all the hyperplanes having the form:

$$\mathbf{v}^T \left(\mathbf{z} - \frac{\Delta}{2} \mathbf{v} \right) = 0 \tag{14}$$

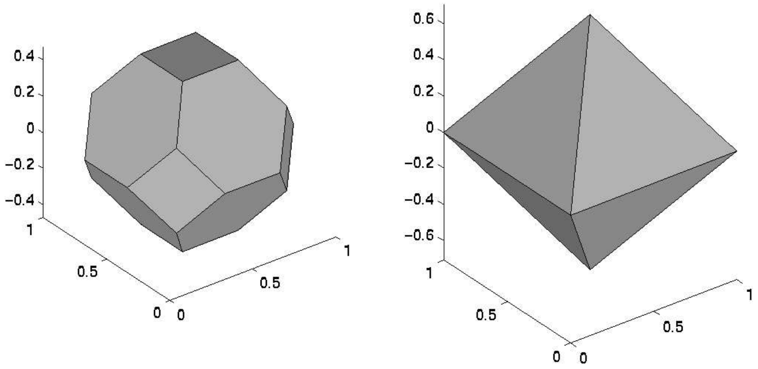

Fig. 2. Regions $\mathcal{S}_{\mathbf{0}}$ and $\mathcal{T}_{\mathbf{0}}$ ($L = 3$)

where $\mathbf{v}$ is any vector such that $v[k] \in \{\pm 1\}$, $k \in \mathcal{S}$ and $\mathbf{z} \in \mathbb{R}^L$. These hyperplanes simply bisect the segments that connect the origin $\mathbf{0}$ and its nearest neighbors in Λ_1. Obviously, $\mathcal{S}_{\mathbf{0}} \subseteq \mathcal{T}_{\mathbf{0}}$, with equality only when $L = 1, 2$. Both regions are depicted in Fig. 2 for $L = 3$.

Several geometrical properties will later allow us to obtain upper bounds to the decoding error probability:

Property 1. [7] Let $\mathcal{S}_{\mathbf{0}}$ be the Voronoi cell associated to the centroid at $\mathbf{0}$. Then, for any other codeword $\mathbf{c} \in \Lambda_{-1} \bigcup \Lambda_1$, its decision region is such that

$$\mathcal{S}_{\mathbf{c}} = \mathcal{S}_{\mathbf{0}} + \mathbf{c} \tag{15}$$

Property 2. By construction, it follows that the set $\mathcal{T}_{\mathbf{0}}$ is symmetric with respect to the coordinate planes.

Property 3. [9] $\mathcal{T}_{\mathbf{0}} \subset \mathcal{R}_{-1}$.

3.2 Non-constant perceptual mask

In the case where the perceptual mask is not constant and the embedding distortions in (2) are taken to their extremal values it is easy to see that the quantization step in each dimension $\Delta[k]$ should be now proportional to $\alpha[k]$. This has the effect of stretching the regions $\mathcal{R}_{\pm 1}$ and consequently the octahedron $\mathcal{T}_{\mathbf{0}}$. If the noise variance at each sample is proportional to $\alpha^2[k]$ (perceptually shaped noise), then it is possible to rescale both the octahedron and the noise by dividing by $\alpha[k]$ so that the original problem with a constant perceptual mask and independent indentically distributed (i.i.d.) noise is recovered.

If the noise variance is not perceptually shaped (for instance, constant variance noise), then it is still possible to divide both the octahedron and the noise by $\alpha[k]$, so that the regular octahedron is recovered; however, now the rescaled noise samples will have different variance. Fortunately, the procedure sketched in App. A for the computation of an upper bound to the probability of decoding error can be easily adapted by noting that scaling a random variable has a well-known effect on its characteristic function.

4 Performance analysis

The discussion on the decision regions made in the previous section allows us to undertake the performance analysis of DC-DM. In order to obtain the decoding error probability we may assume without loss of generality (Property 1) that a symbol $b=-1$ is sent, and that $\mathbf{x}$ is such that $\mathbf{Q}_{-1}(\mathbf{x})=\mathbf{0}$. Considering the detector (11) we have that

$$P_e = P\{\mathbf{z}\in\mathcal{R}_1\} \tag{16}$$

For the determination of P_e, one might be tempted to resort to the well-known union bound with the 2^L nearest neighbors of $\mathbf{0}$ belonging to Λ_1; unfortunately, for moderate values of D_c/D_w and L the results become impractical due to the overlap between the different decision regions that result when only two centroids (i.e., $\mathbf{c}=\mathbf{0}$ and its nearest neighbors $\mathbf{c}\in\Lambda_1$) are taken into account. On the other hand, consideration of a single nearest neighbor, as done in [3], produces overly optimistic results, as we will confirm later.

For obtaining a useful upper bound we will follow a different and novel strategy. Making use of properties P. 1 and P. 3 from the previous section, it is possible to conclude that

$$P_e \le P_s = P\left\{\mathbf{z}\in\overline{\mathcal{T}_{\mathbf{0}}}\right\} \tag{17}$$

where $\overline{\mathcal{T}}$ denotes the complement of $\mathcal{T}$ in $\mathbb{R}^L$.

Let $\mathbf{u}\triangleq\mathbf{n}-(1-\nu)\mathbf{e}$, then from the assumption $\mathbf{Q}_{-1}(\mathbf{x})=\mathbf{0}$ and making use of (8), we can write $\mathbf{z}=\mathbf{u}$. Recalling that $\mathbf{e}$ has i.i.d. components, $e[k]\sim U(-\Delta,\Delta)$, it follows that $\mathbf{u}$ will also have i.i.d. components, each having a pdf

$$f_u(u[k]) = \begin{cases} f_n(n[k]) * \frac{1}{(1-\nu)} f_e(e[k]/(1-\nu)), & 0<\nu<1 \\ f_n(n[k]), & \nu=1 \end{cases} \tag{18}$$

where $f_n(\cdot)$ and $f_e(\cdot)$ denote the marginal pdf's of respectively the noise and the quantization error components.

If the noise pdf is symmetric with respect to the coordinate planes, then $\mathbf{u}$ will inherit this symmetry; thus, from Property 2 the above evaluation of the bound P_s in (17) can be reduced to

$$\begin{aligned} P_s &= P\left\{\|\mathbf{u}\|^2 > \|\mathbf{u}-\Delta\cdot(1,\ldots,1)^T\|^2 \mid \mathbf{u}\in\mathcal{O}\right\} \\ &= P\left\{\|\mathbf{u}'\|^2 > \|\mathbf{u}'-\Delta\cdot(1,\ldots,1)^T\|^2\right\} = P\left\{\sum_{k\in\mathcal{S}} u'[k] > L\Delta/2\right\} \end{aligned} \tag{19}$$

where O is the positive orthant and $\mathbf{u}'$ is an auxiliary random vector with i.i.d. components such that $u'[k] \triangleq |u[k]|$ whose pdf is

$$f_{u'}(u'[k]) \triangleq \begin{cases} 2f_u(u'[k]), & u'[k] > 0 \\ 0, & \text{otherwise} \end{cases}, \quad k \in \mathcal{S}. \tag{20}$$

Now, let

$$r = \sum_{k\in\mathcal{S}} u'[k] \tag{21}$$

Then the pdf of the random variable r is the convolution of L independent random variables with pdf $f_{u'}(u')$ and from (17) P_s is the integral of its tail from $L\Delta/2$ to infinity. This arrangement allows to transform the L-dimensional problem into a unidimensional one. In fact, we have transformed a problem with 2^L kissing spheres into another with just one neighbor, which then becomes tractable as we will see next. By the central limit theorem (CLT), as $L \longrightarrow \infty$, $f_r(r)$ tends to a normal curve. Then, for L very large, $f_r(r)$ can be approximated by a Gaussian pdf whose mean and variance would suffice to compute the desired probability as[4]

$$P_s \approx Q\left(\frac{L\Delta/2 - E\{r\}}{\sqrt{\mathrm{Var}\{r\}}}\right) \tag{22}$$

Moreover, since the components of $\mathbf{u}'$ are i.i.d., we can write

$$\begin{aligned} E\{r\} &= L \cdot E\{u'\} \qquad (23) \\ \mathrm{Var}\{r\} &= L \cdot \mathrm{Var}\{u'\} \qquad (24) \end{aligned}$$

It is important to remark that the approximation (22) should be taken with a grain of salt because the process of building the one-sided distribution $u'[k]$ may produce highly skewed pdf's whose sum converges very slowly to a Gaussian distribution as L increases [17]. If this is the case, the Gaussian approximation to P_s may underestimate the importance of the tails of $f_r(r)$ and give results that are not even an upper bound to the true P_e. In App. A we show a novel technique that can be used to overcome this problem, which is extremely useful for the Gaussian noise case.

4.1 Approximation with uniform noise

With i.i.d. noise components we have $n[k] \sim U(-\eta,\eta)$, $k \in \mathcal{S}$, and for $\eta \geq (1-\nu)\Delta$, the pdf of $u'[k]$ becomes

$$f_{u'[k]}(u'[k]) = \begin{cases} \frac{1}{\eta}, & 0 < u'[k] \leq \eta - (1-\nu)\Delta \\ \frac{\eta+(1-\nu)\Delta-u'[k]}{2(1-\nu)\Delta\eta}, & \eta-(1-\nu)\Delta < u'[k] \leq \eta + (1-\nu)\Delta \\ 0, & \text{otherwise} \end{cases} \tag{25}$$

[4] $Q(x) \triangleq \frac{1}{\sqrt{2\pi}} \int_x^{\infty} e^{-x^2/2} dx$.

while for $\eta < (1-\nu)\Delta$ this pdf is

$$f_{u'[k]}(u'[k]) = \begin{cases} \frac{1}{(1-\nu)\Delta}, & 0 < u'[k] \leq (1-\nu)\Delta - \eta \\ \frac{\eta+(1-\nu)\Delta-u'[k]}{2(1-\nu)\Delta\eta}, & (1-\nu)\Delta - \eta < u'[k] \leq \eta + (1-\nu)\Delta \\ 0, & \text{otherwise} \end{cases} \tag{26}$$

which is obviously the same expression as (25) after swapping η and $(1-\nu)\Delta$.

Even though it is possible to derive an analytical expression for P_e in this case, the exact result becomes quite involved and has little interest. For this reason we will analyze it using the Gaussian approximation described above.

We also remark that under the assumption $\eta \leq \nu\Delta$ (or, equivalently, $D_c < D_w$) for accomplishing the perceptual restrictions, the upper bound (22) becomes in this case a good approximation of P_e. This is due to the fact that the one-sided distribution (25) is highly localized [17].

Considering (23-24), we can compute the mean and variance of r for the case (25) as

$$E\{r\} = L\,\frac{3\eta^2 + (1-\nu)^2\Delta^2}{6\eta} \tag{27}$$

$$\mathrm{Var}\{r\} = L\,\frac{3\eta^4 + 6\eta^2(1-\nu)^2\Delta^2 - (1-\nu)^4\Delta^4}{36\eta^2} \tag{28}$$

The same statistics in the case (26) become straightforwardly

$$E\{r\} = L\,\frac{3(1-\nu)^2\Delta^2 + \eta^2}{6(1-\nu)\Delta} \tag{29}$$

$$\mathrm{Var}\{r\} = L\,\frac{3(1-\nu)^4\Delta^4 + 6\eta^2(1-\nu)^2\Delta^2 - \eta^4}{36(1-\nu)^2\Delta^2} \tag{30}$$

Then P_e can be approximated by

$$P_e \approx \begin{cases} Q\left(\sqrt{L}\frac{3\nu\xi - 3\nu^2 - (1-\nu)^2\xi^2}{\sqrt{3\nu^4 + 6\nu^2(1-\nu)^2\xi^2 - (1-\nu)^4\xi^4}}\right), & \xi < \nu/(1-\nu) \\ Q\left(\sqrt{L}\frac{3\nu(1-\nu)\xi^2 - \nu^2}{\sqrt{3(1-\nu)^4\xi^4 + 6\nu^2(1-\nu)^2\xi^2 - \nu^4}}\right), & \xi \geq \nu/(1-\nu) \end{cases} \tag{31}$$

where in this case, it is possible to show that $\xi = \nu\Delta/\eta$.

It is possible to improve on the approximation (31) by considering that, if $\nu \geq 1/2$ then $P_e = 0$ when $\xi \geq \nu/(\nu - 1/2)$ because the pdf (25-26) has finite length. Also, as ξ is constrained to be greater than zero, errorless decoding can never happen for $\nu < 1/2$ regardless the value of ξ.

Evidently, there is a performance variability associated to the parameter ν controlling the distortion compensation. It is interesting to see that for $\xi = 1$ the probability of error can take the worst possible value, i.e. $P_e = 0.5$, when either $\nu = 1$ or $\nu = 0$.

Also, note that the appearance of a factor $\sqrt{L}$ governing the asymptotic performance means that usage of many dimensions can be assimilated to a form of repetition coding, where this factor would be the coding gain.

4.2 Approximation with Gaussian noise

Assume that $\mathbf{n}$ is a random vector with L i.i.d. components with zero mean and variance σ_g^2, that is

$$f_n(\mathbf{n}) = \left(2\pi\sigma_g^2\right)^{-L/2} \exp\left(-\frac{1}{2}\mathbf{n}\Gamma^{-1}\mathbf{n}^T\right) \tag{32}$$

where Γ is the noise covariance matrix that takes the form $\Gamma = \sigma_g^2\mathbf{I}$, with $\mathbf{I}$ the $L \times L$ identity matrix. Hence, the pdf of the random variable $u[k]$ in (18) becomes the convolution of a zero-mean Gaussian random variable with variance σ_g with a random variable uniformly distributed in the interval $(-(1-\nu)\Delta, (1-\nu)\Delta)$. Then, the pdf of $u'[k]$, $k \in \mathcal{S}$ in (20) becomes

$$f_{u'[k]}(u'[k]) = \begin{cases} \frac{1}{(1-\nu)\Delta}\left(Q\left(\frac{u'[k]-(1-\nu)\Delta}{\sigma_g}\right) - Q\left(\frac{u'[k]+(1-\nu)\Delta}{\sigma_g}\right)\right), & u'[k] > 0 \\ 0, & \text{otherwise} \end{cases} \tag{33}$$

Since the one-sided Gaussian distribution is highly skewed, CLT approximation to the bound P_s is only valid for very large L, as we have already discussed. Even for high values of L, say $L = 100$, the CLT approximation holds only for values of ξ close to one. On the other hand, this approximation becomes very simple to compute; for this reason and for the purpose of comparison, in Section 5 we give the results, noting that its practical utility is limited by the actual values of L and ξ. In this case, for using (22) it would be necessary to compute numerically (23) and (24). We must mention that the bound obtained by following this procedure is asymptotically tight as $\xi \to \infty$, with $\xi = \nu\Delta/(\sqrt{3}\sigma_g)$. In the limit, the probability that $\mathbf{n}$ falls in $\mathcal{R}_{-1}$ but not in $\mathcal{T}_{\mathbf{0}}$, becomes negligible.

The procedure for computing a true upper bound to P_s in the Gaussian case for moderate values of L is given in App. A. This procedure has been adapted from a technique originally proposed in the digital communications area, by Beaulieu in [18] and [19], which is suitable for analyzing the performance of equal gain diversity communications receivers in fading channels. The main advantage of reduction in the kissing numbers that we have shown in this Section is that from this point on we are able to exploit Beaulieu's technique and provide very accurate bounds.

5 Experimental results

Next, the theoretical probabilities of error obtained in previous sections are compared and validated against empirical data. Plots are presented for values of WNR starting at WNR $= 0$ dB ($\xi = 1$). First we can see in Figs. 3 and 4 the comparison of empirical data generated through Monte Carlo simulations with the theoretical predictions for the two considered types of noise distortion and for the case of perceptually shaped noise.

Apart from the theoretical predictions given in this paper, we also depict in these figures the single-neighbor approximation that would follow from adapting the results in [3] as well as the corresponding union bound using the 2^L nearest neighbors. We can see in Figs. 3 and 4 that the single-neighbor approach clearly underestimates the true P_e because it does not take

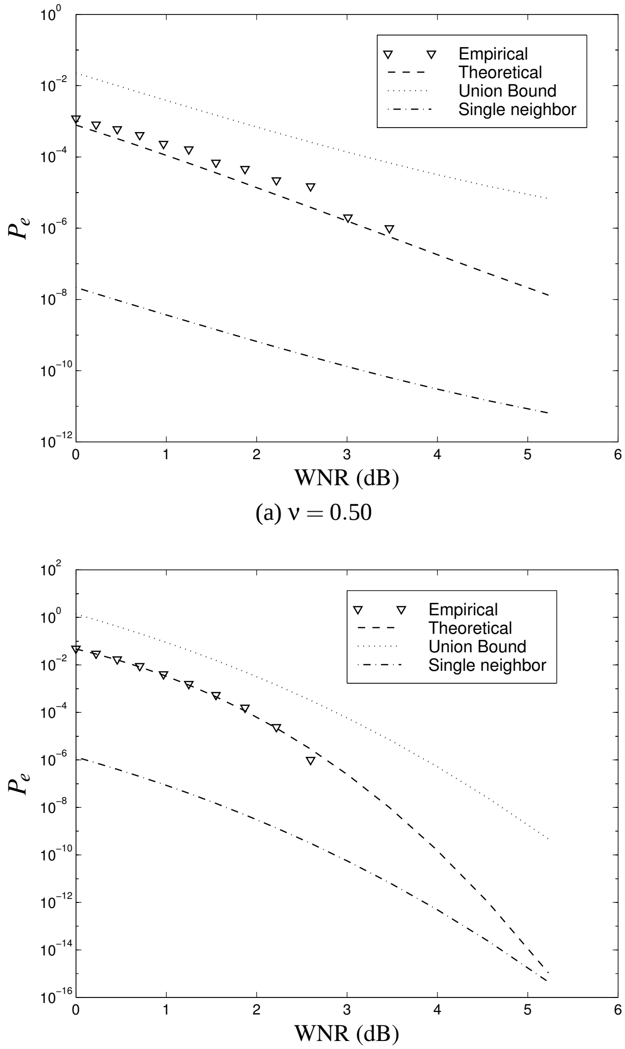

(a) $\nu = 0.50$

(b) $\nu = 0.80$

Fig. 3. Monte Carlo simulation, uniform noise, $L = 20$

into account the actual number of neighbors at such distance, that grows exponentially with dimensionality. On the other hand, it becomes evident that the union bound, that considers all aforementioned neighbors but disregards any overlapping between their respective decision regions, is much too loose to be useful, especially for low values of WNR.

Regarding the theoretical predictions given in Sect. 4 we can see that the approximation using the CLT is acceptable for the uniform distortion, due to the finite length of this pdf. As we see in the case for $\nu = 0.5$, in the mid-range values of ν this approximation slightly underestimates the real P_e. This discrepancy is explained due to the fact that the Gaussian approximation of (21) is above the true pdf at the error event values. Concerning the Gaussian

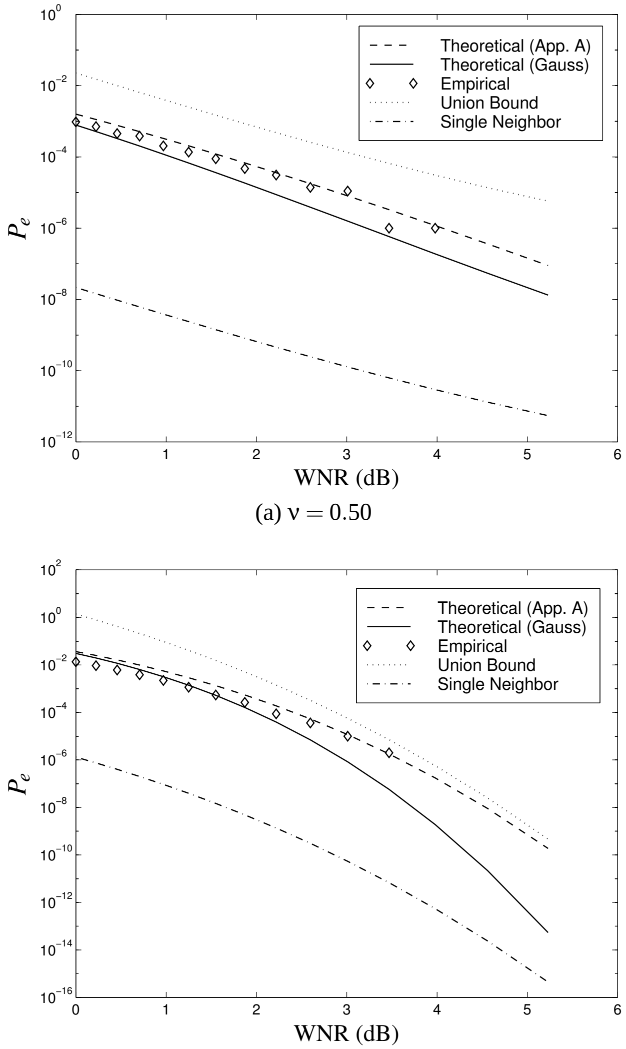

(a) $\nu = 0.50$

(b) $\nu = 0.80$

Fig. 4. Monte Carlo simulation, Gaussian noise, $L = 20$

channel distortion, it is clear that the CLT approximation is not good anymore, as it had been previously warned. even though, it improves to become progressively tight for decreasing ν as predicted. However the prediction using our novel approach succeeds to tightly upperbound the probability of decoding error for any WNR. It is important to highlight that Beaulieu's approach could be also applied to the uniform case to get even more reliable bounds.

Is is remarkable that the lower the value of ν, the less important the pdf of noise. If we compare the plots for $\nu = 0.5$ we see that they are practically identical, while for the $\nu = 0.8$ case there is a notable difference between uniform and Gaussian noise. Last, we have numerically computed what value of ν yields the optimum performance for $\xi = 1$, that is in both cases $\nu = 1/2$.

Next, in Fig. 5 we confirm the previous Monte Carlo results using a real implementation of DC-DM. The host signal $\mathbf{x}$ chosen is the well-known *Lena* image (256×256), and the embedding takes place in the spatial domain. The perceptual mask $\boldsymbol{\alpha}$ is computed using the method proposed in [14] with a visibility factor $c = 1$, which happens to be very conservative. In Fig. 5 Gaussian noise is used as the channel distortion and two cases are considered for the noise variance: (a) locally proportional to the perceptual mask energy; (b) equal to its average value. The theoretical bounds are depicted using the respective approximations derived in App. A.

Some practical aspects are taken into account in the tests shown here. First, after embedding the watermark, $\mathbf{y}$ is quantized to integer values and limited between $[0, 255]$ in order to produce a real image. Second, the decoder obtains the perceptual mask from $\mathbf{z}$ without knowing $\mathbf{x}$. This provokes that the $\boldsymbol{\alpha}$ used by the decoder be just an approximation to the one used by the encoder.

Finally, in Fig. 6 we consider the effect of applying Wiener filtering to the watermarked image. We can see that the result of this linear filter is more injurious than just adding Gaussian noise with the same power. This leaves the door open to theoretically analyzing the performance of DC-QIM to other types of channel distortion which may be more harmful.

6 Conclusions

In this paper we have given a theoretical analysis together with empirical verifications on the attainable performance of multidimensional DC-DM using uniform quantizers. Gaussian and uniform random additive attacks were considered.

Improvements could be expected if other lattices than the one employed were used. In general, these lattices would lead to a problem non-separable in each one of the dimensions, and therefore harder to solve. Anyway better options for multidimensional methods using quantization are actually known to be available. Specifically, further improvements are possible using quantized projections like Spread-Transform Dither Modulation (STDM) [3] or Quantized Projection (QP) [9]. These projections reduce the number of “kissing points” for the same embedding energy and thus make it possible to diminish the probability of error.

For instance, preliminary tests show that QP would be able to overcome the performance of the optimum sphere-packing scheme for $L = 2$, i.e. the A_2 hexagonal lattice, using host signal and attacking distortion with Gaussian statistics.

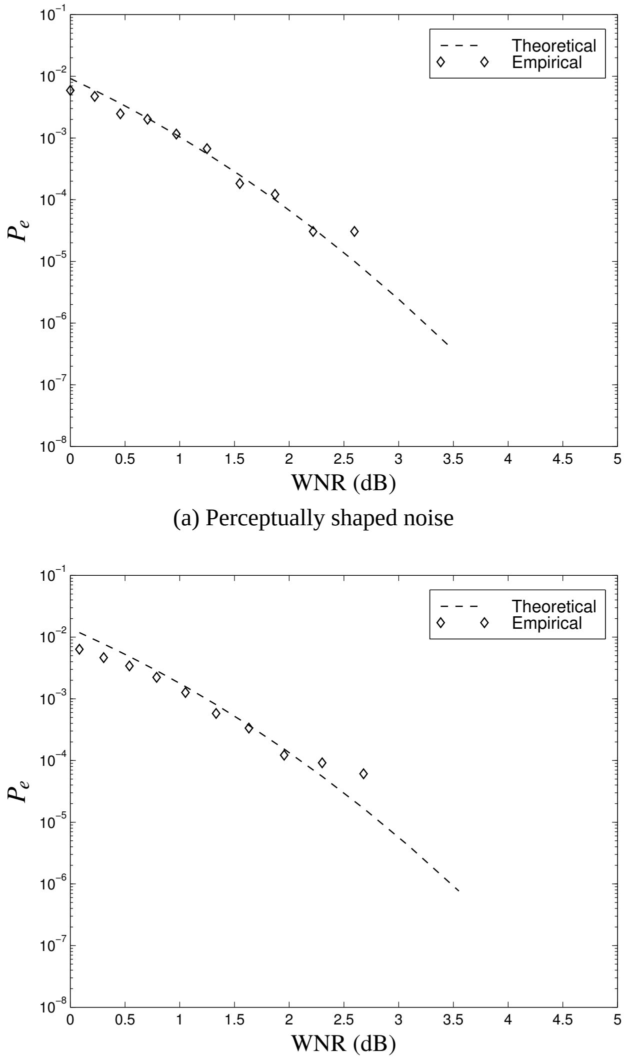

(a) Perceptually shaped noise

(b) Non-perceptually shaped noise

Fig. 5. Real DC-DM implementation using *Lena* image, Gaussian noise, $\nu = 0.70$, $L = 20$

A True upper bound with Gaussian noise

We start by defining the normalized random variable $m[k] \triangleq u'[k]/\sigma_g$, that allows us to rewrite (19) as

$$P_s = P\left\{\sum_{k\in\mathcal{S}} m[k] > \frac{L\xi\sqrt{3}}{2\nu}\right\}, \tag{34}$$

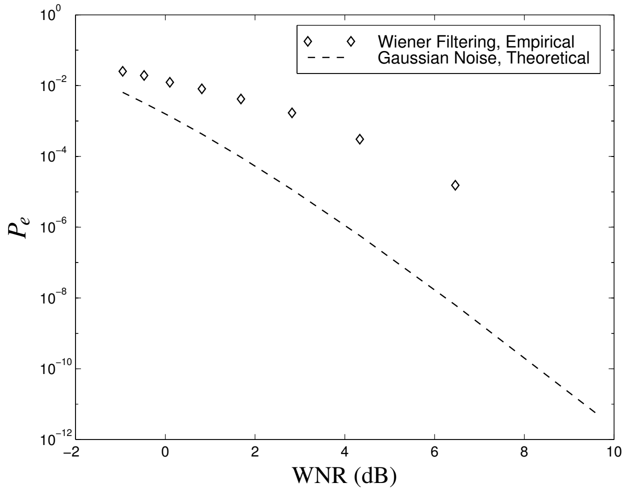

Fig. 6. Real DC-DM implementation using *Lena* image, Wiener filtering, $\nu = 0.50$, $L = 20$

where ξ takes the same form as in the previous section. Considering (33), the one-sided normalized random variable $m[k]$ has the following pdf

$$f_{m[k]}(m[k]) = \begin{cases} \frac{1}{\mu}\Big(Q\left(m[k]-\mu\right) - Q\left(m[k]+\mu\right)\Big), & m[k] > 0 \\ 0, & \text{otherwise} \end{cases} \tag{35}$$

where, for convenience, we have defined

$$\mu \triangleq \frac{(1-\nu)\Delta}{\sigma_g} \tag{36}$$

In order to compute P_s from its definition in (34) the characteristic function $M(\omega)$ of $m[k]$ has to be first obtained. Let $\omega_l = 2\pi l/T$ for any positive integer l and with T a sufficiently large real number. Then, following [18], P_s may be calculated as

$$P_s \approx \frac{1}{2} + \frac{2}{\pi} \sum_{\substack{l=1 \\ l \text{ odd}}}^{\infty} \frac{|M(\omega_l)|^L \sin(L\theta(\omega_l))}{l} \tag{37}$$

where $\theta(\omega)$ is defined as

$$\theta(\omega) \triangleq \arg\{M(\omega)\} - \omega\frac{\xi\sqrt{3}}{2\nu} \tag{38}$$

and $\arg(x)$ denotes the four-quadrant phase of the complex number x. The series in (37) is pointwise convergent with an accuracy that depends on the value of T. A greater accuracy is obtained for larger values of T but this requires truncation to more terms in the series for a practical implementation.

When the noise variance is not perceptually shaped, (19) remains valid after the rescaling procedure explained in Sect. 3.2. As already noted, in this case the variance of $u'[k]$ is different for each k; it is straightforward to see that the characteristic functions of these random variables can be written as $U'_k(\omega) = M(\omega\,\sigma_n[k]/\alpha[k])$. Now the calculation of P_s is made, following [18], as

$$P_s \approx \frac{1}{2} + \frac{2}{\pi} \sum_{\substack{l=1 \\ l \text{ odd}}}^{\infty} \frac{\prod_{k\in\mathcal{S}} |U'_k(\omega_l)| \sin(\sum_{k\in\mathcal{S}} \phi_k(\omega_l))}{l} \tag{39}$$

where $\phi_k(\omega)$ is defined as

$$\phi_k(\omega) \triangleq \arg\{U'_k(\omega)\} - \omega\,\Delta/2 \tag{40}$$

It only remains the calculation of $M(\omega)$, which can be shown to be [9]

$$M(\omega) = e^{-\omega^2/2}\frac{\sin(\mu\omega)}{\mu\omega} + j\frac{1}{\mu\omega}\left[\Phi\left(\frac{\mu}{\sqrt{2}}\right) + e^{-\omega^2/2}\,\mathrm{Re}\left\{e^{-j\mu\omega}\Phi\left(\frac{-\mu+j\omega}{\sqrt{2}}\right)\right\}\right] \tag{41}$$

with $\Phi(\cdot)$ the error function with complex argument.

References

1. Chen, B., Wornell, G.W.: Provably robust digital watermarking. In: Proc. of SPIE. Volume 3845 of Multimedia Systems and Applications II., San José, USA (1999) 43–54
2. Costa, M.H.: Writing on dirty paper. IEEE Trans. on Information Theory **29** (1983) 439–441
3. Chen, B., Wornell, G.W.: Quantization index modulation: A class of provably good methods for digital watermarking and information embedding. IEEE Trans. on Information Theory **47** (2001) 1423–1443
4. Eggers, J.J., Su, J.K., Girod, B.: A blind watermarking scheme based on structured codebooks. In: Proc. of IEE Conf. on Secure Images and Image Authentication, London, UK (2000)
5. Chae, J.J., Mukherjee, D., Manjunath, B.S.: A robust data hiding technique using multidimensional lattices. In: Procs. of the IEEE Forum on Research and Technology Advances in Image Processing, Santa Bárbara, USA (1998) 319–326
6. Chae, J., Mukherjee, D., Manjunath, B.S.: Color image embedding using multidimensional lattice structures. In: Procs. of the IEEE Intnal. Conference on Image Processing (ICIP'98). Volume 1., Chicago, USA (1998) 460–464
7. Conway, J., Sloane, N.: Sphere Packings, Lattices and Groups. 3rd edn. Volume 290 of Comprehensive Studies in Mathematics. Springer (1999)
8. Brunk, H.: Quantizer characteristics important for quantization index modulation. In Wong, P.W., Delp, E.J., eds.: Security and Watermarking of Multimedia Contents III. Volume 4314 of Proc. of SPIE., San José, USA (2001) 686–694
9. Pérez-González, F., Balado, F., Hernández, J.R.: Performance analysis of existing and new methods for data hiding with known host information in additive channels. (2001) Submitted to IEEE Transactions on Signal Processing.
10. Wolfgang, R.B., Podilchuk, C.I., Delp, E.J.: Perceptual watermarks for digital images and video. Proceedings of the IEEE **87** (1999) 1108–1125
11. Voloshynovskiy, S., Herrigel, A., Baumgärtner, N., Pun, T.: A stochastic approach to content adaptive digital image watermarking. In: 3rd International Workshop on Information Hiding, Desden, Germany, Springer-Verlag (1999)
12. Hernández, J.R., Amado, M., Pérez-González, F.: DCT-domain watermarking techniques for still images: Detector performance analysis and a new structure. IEEE Trans. on Image Processing **9** (2000) 55–68 Special Issue on Image and Video Processing for Digital Libraries.

13. Moulin, P., Mıhçak, M.K.: The data hiding capacity of image sources. (2001)
14. Hernández, J.R., Pérez-González, F., Rodríguez, J.M.: Coding and synchronization: A boost and a bottleneck for the development of image watermarking. In: Proc. of the COST #254 Int. Workshop on Intelligent Communications, L'Aquila, Italy, SSGRR (1998) 77–82
15. Voloshynovskiy, S., Pereira, S., Herrigel, A., Baumgärtner, N., Pun, T.: Generalized watermark attack based on watermark estimation and perceptual remodulation. In Wong, P.W., Delp, E.J., eds.: Electronic Imaging 2000: Security and Watermarking of Multimedia Content II. Volume 3971 of SPIE Proceedings., San José, USA (2000)
16. Moulin, P., O'Sullivan, J.A.: Information-theoretic analysis of information hiding. (2001)
17. Bury, K.: Statistical Models in Applied Science. Robert E. Krieger Publishing Company, Malabar, Florida (1975)
18. Beaulieu, N.C.: An infinite series for the computation of the complementary probability distribution function of a sum of independent random variables and its application to the sum of Rayleigh random variables. IEEE Trans. Commun. **38** (1990) 1463–1474
19. Beaulieu, N.C., Abu-Dayya, A.A.: Analysis of equal gain diversity on Nakagami fading channels. IEEE Trans. Commun. **39** (1991) 225–234