Improving the performance of RDM watermarking by means of trellis coded quantisation

A. Abrardo, M. Barni, F. Pérez-González and C. Mosquera

Abstract: We propose a method to improve the performance of the recently introduced rational dither modulation (RDM) watermarking scheme. The improvement is obtained by modifying the essentially scalar nature of RDM with the introduction of a vector quantiser. The vector quantiser is based on a properly modified version of classical trellis coded quantisation, thus leading to the new trellis-coded RDM scheme (TC-RDM). Due to the infinite memory of the trellis TC-RDM relies on, the standard Viterbi algorithm is no longer optimum, hence we introduce a new suboptimal encoding algorithm that keeps the computational complexity reasonably low while ensuring a minimal loss with respect to the optimal scheme. The simulations we carried out on synthetic sequences show that TC-RDM permits improvement of the performance of RDM by 2–4 dB in terms of robustness against Gaussian noise addition.

1 Introduction

QIM watermarking, introduced in [1], is known to greatly outperform the older algorithms based on the spread spectrum paradigms. At the same time, practical implementations of QIM, such as SCS and ST-DM [1, 2], are vulnerable against valuemetric scaling: it only needs that the features hosting the watermark are scaled by an unknown gain factor, say ρ, to imped correct decoding of the watermark [3]. This is a serious problem, since in most cases multiplication by a constant gain does not damage the quality of the host document at all. On the contrary, sometimes it has a beneficial effect on the the watermarked data, as in the image case where multiplication by a factor $\rho > 1$ is often used to increase the image contrast.

Possible remedies include estimation of the gain factor at the decoder, possibly with the aid of a pilot signal [4,5], use of spherical (equi-energetic) codewords [6,7], or use of an image-dependent quantisation step [8,9]. A simple and effective solution to the problem has been proposed recently [10,11]. Such a solution, named rational dither modulation (RDM), achieves invariance against the gain attack by quantising a rational function of consecutive features instead of the features themselves. In this way gain invariance is obtained at the expense of a minor modification of the classical DM scheme, leading to a very simple and effective algorithm. As it is shown in [10,11], by properly designing the RDM algorithm the same performance of conventional DM can be achieved asymptotically. An approach which somewhat resembles RDM, though with significant differences and in a different context, is described in [12].

The basic RDM scheme can be improved in many ways, for instance by applying distortion compensation, or by introducing some form of channel coding to increase the robustness of the watermark. Here we focus on another aspect of RDM that needs improvement. In the form described in [10,11,13], RDM is an essentially scalar algorithm, since each feature ratio is quantised by itself, by means of a scalar quantiser. Yet it is known that better results can be obtained by means of vector quantisation. This is indeed the purpose of this paper: to replace the scalar quantiser of RDM with a trellis-based vector quantiser, whereby a set of consecutive ratios are quantised all together. As it is shown in Section 5, simulations carried out on i.i.d. Gaussian features demonstrate that the new trellis-coded RDM (TC-RDM) scheme permits a gain in the range of 2–4 dB with respect to plain RDM in standard working conditions ($P_e \simeq 10^{-3}$).

The rest of this paper is organised as follows. In Section 2 the basic notation used throughout the paper is introduced. In Section 3 the classical RDM algorithm is revised. Section 4 describes the new TC-RDM algorithm, with particular attention to the description of the trellis-based quantiser. Simulation results are shown in Section 5. Section 6 presents our conclusions.

2 Notation

This section introduces the basic notation used throughout the paper. We assume that the features hosting the watermark are arranged into a one-dimensional vector $\mathbf{X} = (x_1 \dots x_{N_t})$ of length N_t. The to-be-hidden message is indicated by a binary vector $\mathbf{W}$ whose length is equal to N_w and the vector with the watermarked features by $\mathbf{Y}$. We assume that the host feature samples are i.i.d. random variables having zero mean and variance σ_x^2 (actually we use these assumptions only for the experimental validation of our method, while the theoretical part of the paper does not depend on

IEE Proceedings online no. 20055152

doi:10.1049/ip-ifs:20055152

Paper first received 20th December 2005 and in revised form 8th June 2006

A. Abrardo and M. Barni are with the Department of Information Engineering, University of Siena, Siena 53100, Italy

F. Pérez-González and C. Mosquera are with the Department Teoría de la Señal y Comunicaciones, ETSI Telecom, Universidad de Vigo, Spain

them). Note that we will use the same symbol both to indicate a random variable and the specific values assumed by the variable, the exact meaning of each symbol being clearly identified by the context wherein the symbol is used.

We assume that the decoder works on a manipulated version of the watermarked feature vector. We indicate such a manipulated vector by $\mathbf{Z}$. In particular the manipulations addressed here include the multiplication of $\mathbf{Y}$ by a constant factor ρ unknown to the decoder (gain attack) and the addition of a white Gaussian noise vector $\mathbf{N}$ (AWGN attack).

The evaluation of the performance of the system requires that some other definitions are given. The embedding distortion D_w is defined as the average value of the watermarking signal, that is $D_w = E[\| \mathbf{Y} - \mathbf{X}\|^2]/N_t$. In a similar way, the attack distortion is defined as $D_a = E[\|\mathbf{N}\|^2]/N_t$. Note that the presence of the gain factor ρ does not have any influence on D_a whose only goal is to measure the strength of the AWGN part of the attack. It is also useful to introduce the document-to-watermark ratio, $\mathrm{DWR} = \sigma_x^2/D_w$, the watermark-to-noise ratio, $\mathrm{WNR} = D_w/D_a$ and the document-to-noise ratio, $\mathrm{DNR} = \sigma_x^2/D_a$. These quantities are often given in decibels.

3 Rational dither modulation

In this section we briefly review the RDM algorithm. For a more detailed description readers are referred to [10, 11].

Let $\mathbf{y}_k^h$ denote the vector with samples from y_k to y_h. For example, in the following we will often use the symbol $\mathbf{y}_{k-L}^{k-1} = (y_{k-L}, y_{k-L+1}, \ldots, y_{k-1})$. We consider the set $\mathcal{G}$ of functions $g : \mathcal{Y}^L \to \mathbb{R}$ having the property that for any $\rho > 0$, and any vector $\mathbf{y}$, $g(\rho\mathbf{y}) = \rho g(\mathbf{y})$. Given the kth bit to be hidden w_k, embedding goes through the application of standard DM watermarking to the ratio $x_k/g(\mathbf{y}_{k-L}^{k-1})$, i.e.

$$y_k = g(\mathbf{y}_{k-L}^{k-1}) Q_{w_k}\left(\frac{x_k}{g(\mathbf{y}_{k-L}^{k-1})}\right) \tag{1}$$

where the quantiser $Q()$ is chosen according to the particular bit to be embedded, hence justifying the subscript appearing in the equation. Initialisation is made to an arbitrary state agreed by the embedder and the decoder. The decoder receives z_k and based on its previous knowledge about $\mathbf{z}_{k-L}^{k-1}$, decodes the hidden bit by applying the standard DM decoding procedure to the ratio between z_k and $g(\mathbf{z}_{k-L}^{k-1})$:

$$\hat{w}_k = \arg \min_{w_k \in \{-1,1\}} \left| \frac{z_k}{g(\mathbf{z}_{k-L}^{k-1})} - Q_{w_k}\left(\frac{z_k}{g(\mathbf{z}_{k-L}^{k-1})}\right) \right| \tag{2}$$

From the above equation and the properties of the function g, it is immediate to see that RDM is intrinsically immune against the gain attack. Note also that RDM is very close to standard DM watermarking; in embedding, the only difference regards the division of x_k by $g(\mathbf{y}_{k-L}^{k-1})$ prior to quantisation, while in decoding, due to the unavailability of $\mathbf{y}_{k-L}^{k-1}$ the divisor becomes $g(\mathbf{z}_{k-L}^{k-1})$. As to the choice of $g()$, the set $\mathcal{G}$ includes, but is not limited to, the l_p vector-norms, given by:

$$g(\mathbf{y}_{k-L}^{k-1}) = \left(\frac{1}{L}\sum_{i=1}^{L} |y_{k-i}|^p\right)^{1/p} \tag{3}$$

For instance, in [10, 11] the squared Euclidean norm is adopted, by letting $p = 2$ in (3). Previous works on RDM demonstrated the good performance of the system that approach that of conventional DM for $L \to \infty$. At the same time with respect to other gain-invariant schemes such as SS and ISS [14, 15], RDM has a great advantage in terms of capacity.

The main weakness of the basic RDM algorithm is its essentially scalar nature, in that the overall quantisation codebook is nothing but a rectangular lattice. As we demonstrate in the following sections, better performance can be obtained by using a vector quantisation approach.

4 Trellis-coded RDM

In this section we give a detailed description of TC-RDM watermarking. We first outline the basic ideas behind it, then we pass to a formal description of the algorithm.

4.1 The basic idea

The main idea behind TC-RDM is to replace the scalar quantisers used by RDM with a set of vector quantisers built by relying on a dirty trellis mechanism similar to that described in [6, 16]. The overall scheme of the TC-RDM embedder is depicted in Fig. 1. Let $\mathcal{R}$ be the codebook that will be used to quantise the feature ratios, in the following we will assume that each codeword $\mathbf{r} \in \mathcal{R}$ is a P-long vector of reals. By starting from the watermark sequence $\mathbf{W}$, we form a bit sequence $\mathbf{B}$ obtained by interleaving the watermark bits in $\mathbf{W}$ with a number of so called *free bits* $\mathbf{V}$ that will be used to guide the choice of the codeword in $\mathcal{R}$ that minimises the watermark embedding distortion. The bits in $\mathbf{B}$ are feeded in blocks of size M to a shift register structure whose content represents the status of the TC-RDM embedder. Let the number of shift register blocks each of size M be v. At each time step t, the TC-RDM embedder uses the content of the shift registers (comprising both the actual input of M bits and the previous $v-1$ blocks representing the memory of the embedder) to index a codeword of $\mathcal{R}$, let us call it $\mathbf{r}_t = (r_{t,1} \ \ldots \ r_{t,P})$. It is this vector of P values at a time, that is interpreted as the sequence of ratios to be used in (1) instead of $Q_{w_k}(x_k/g(\mathbf{y}_{k-L}^{k-1}))$. It is clear that the sequence of quantised ratios $\mathbf{R}$ will depend both on the informative bits $\mathbf{W}$ and the free bits $\mathbf{V}$. Being the informative bits fixed, the output sequence of ratios will then be a function of $\mathbf{V}$, whose actual value will be chosen in such a way to minimise the embedding distortion (see (6) below). The role of the informative and free bits is easily understood if we adopt the conventional random binning approach to watermark embedding. According to it, the informative bits determine the bin (or coset) within which the codeword with the quantised ratios must be search for, and the free

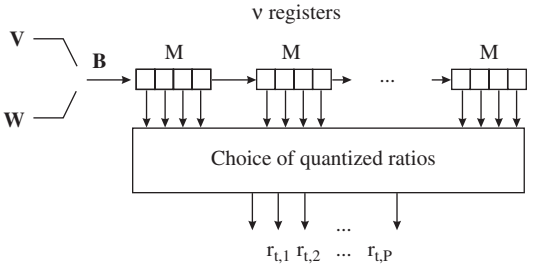

Fig. 1 *Overall scheme of the watermark embedder*

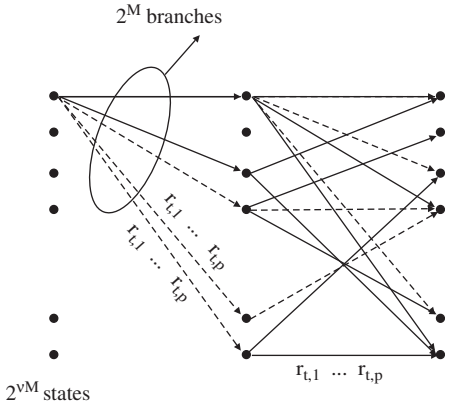

Fig. 2 *Trellis representation of the embedder*

bits determine the choice of a particular codeword within the bin. As usual this choice (corresponding to the choice of the free bits) is made in such a way that the embedding distortion is minimised. As will be detailed in the following section, the choice of the quantised values and the free bits is made by means of a trellis like structure, hence justifying the appellative TC-RDM.

4.2 The algorithm

We now need to describe how the quantised ratios and the free bits are chosen. These two steps are carried out jointly by means of a trellis structure as that depicted in Fig. 2. For sake of simplicity we will consider the case of binary inputs, i.e. the input vector at generic time instant t is given by $\mathbf{b}_t = (b_{t,1}, b_{t,2}, \ldots, b_{t,M})$, $b_{t,j}$ being the jth input bit at tth transition. The content of the v shift registers define a trellis structure with 2^{vM} states. At the output of each state we find 2^M branches each corresponding to a different input block. Each branch is associated to a vector of quantised ratios $\tilde{\mathbf{r}} = (\tilde{r}_1 \ldots \tilde{r}_P)$, that in turn univocally defines the watermarked signal through the RDM equation

$$y_k = g(\mathbf{y}_{k-L}^{k-1}) \cdot r_k \tag{4}$$

At each transition we can identify the branches for which the informative bits correspond to the actual informative bits at the input of the embedder (indicated as solid branches in the figure). The union of all the branches with the right informative bits define a set of paths on the trellis. The goal of the embedder is to choose the path that results in the minimum distortion [Note 1].

To be specific, let us assume that k out of M bits of $\mathbf{b}_t$ are the actual information bits to be hidden in the host features at time instant t, while the remaining $M-k$ are free bits. That is, the first k bits are taken (in a sequential fashion) from the message sequence $\mathbf{W}$, whereas the others are left free. The message sequence $\mathbf{W}$ is split into $N_w/k = N$ chunks [Note 2] $(\mathbf{w}_1 \ldots \mathbf{w}_N)$, with $\mathbf{w}_i = \mathbf{w}_{(i-1)k+1}^{(i-1)k+k}$, then the trellis input is composed as follows

$$\mathbf{b}_i = (\mathbf{w}_i, \mathbf{v}_i) \tag{5}$$

where $\mathbf{v}_i$ is the vector with the $M-k$ free bits. Hence, for each information sequence of kN bits, a set (bin) of $2^{(M-k)N}$ codewords is obtained.

As we said, the output of the trellis, formed by blocks of P values at a time, is interpreted as the sequence of ratios to be used in (1) instead of $Q_{w_k}(x_k/g(\mathbf{y}_{k-L}^{k-1}))$. Note that since for each block of k input bits the trellis produces P output values, and the dimensionality of $\mathbf{Y}$ must be equal to that of $\mathbf{R}$, it is necessary that the length of $\mathbf{W}$ is equal to $k \cdot N_t/P$.

We must now define a strategy to choose the free bits the output sequence of quantised ratios $\mathbf{R}$ depends on. To do so, we consider an MSE informed embedding approach, whereby the free bits are chosen so that the mean squared error between the watermarked sequence (that is a function of $\mathbf{R}$ via (4)) and the host features vector $\mathbf{X}$ is minimised.

To proceed, let us remember that $\mathbf{w}_t$ and $\mathbf{v}_t$ represent the vectors of informative and free bits at time instant t, respectively. We have $\mathbf{V} = (\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_N)$. Note that, according to the above notations, we have $\mathbf{B} = (\mathbf{W}, \mathbf{V})$ and $\mathbf{R} = \mathbf{R}(\mathbf{W}, \mathbf{V})$. The goal of the embedder is to find $\mathbf{V}$, and hence $\mathbf{R}(\mathbf{W}, \mathbf{V})$ so that the transmitted signal $\mathbf{Y}$ defined by means of (4) is as *close* as possible to the side information $\mathbf{X}$.

By summarising, for a given information vector $\mathbf{W}$, the MSE criterion for information embedding consists in evaluating the free bits vector $\mathbf{V} = \tilde{\mathbf{V}}$ as:

$$\tilde{\mathrm{V}} = \underset{\mathrm{V}}{\arg\min} \sum_{n=1}^{N_t} (x_n - y_n(\mathrm{W}, \mathrm{V}))^2 \tag{6}$$

In order to actually perform the above minimisation, it is essential to note that, due to the feedback introduced in the computation of $\mathbf{Y}$ (see (4)), the term y_n depends on all the previous inputs, i.e. the memory is infinity. Thus, the only possibility to perform the minimisation in (6) is by means of exhaustive search. This means that the embedder should consider the whole encoding tree, with a complexity which increases exponentially with N (and hence with N_t). Of course, this approach is not feasible and alternative suboptimum strategies must be envisaged. In the next section we describe a trellis-based suboptimum Viterbi approach with a fixed number of survivors.

As to decoding, this is a straightforward operation, since the decoder only needs to calculate the sequence of ratios $z_k/g(\mathbf{z}_{k-L}^{k-1})$ and use it to feed the trellis. By running a standard Viterbi decoding algorithm the path on the trellis that is closest to the sequence of received ratios is selected, and the corresponding bit sequence is given as the output of the decoder [Note 3]. As it can be readily seen, due to the properties of $g()$, the output of the decoder is invariant with respect to the gain factor ρ, hence ensuring immunity against the gain attack.

4.3 Suboptimum embedding strategy

To start with, let us observe that every path on the trellis results in a transmitted sequence $\mathbf{Y}(\mathbf{W}, \mathbf{V})$ defined by (4). Moreover, for each partial path $\mathbf{y}_1^t$ it is possible to evaluate the partial path distance:

$$D^2(\mathbf{x}_1^t, \mathbf{y}_1^t) = \sum_{n=1}^{tP} (x_n - y_n)^2 \tag{7}$$

where the sum goes from 1 to tP since at each step t, the trellis outputs P ratios. The proposed suboptimum

Note 1: As to trellis initialisation, we assume that the initial state is known both to the embedder and the detector, since they may have agreed about it beforehand.
Note 2: We neglect border effects for simplicity.

Note 3: It is worth remembering that in this case the Viterbi algorithm is run on the complete trellis.

approach works by evaluating the partial distance between the host feature sequence $\mathbf{x}_1^t$ and all the sequences $\mathbf{y}_1^t$ deriving from the partial trellis paths $\mathbf{r}_1^t$, where the various paths differ because of a different choice of the free bits sequence $\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_t$. Of course the number of possible paths increases exponentially with t. To avoid such a complexity, each time two or more paths enter the same state the partial distances associated to them are evaluated and only the N_s paths with the smallest distances are retained. In this way, the complexity of the algorithm is proportional to N_s and increases only linearly with the length of the host feature sequence.

As far as memory requirement is of concern, it is worth noting that the proposed embedding strategy requires each survivor to store the accumulated metric and the previous state (as in the standard Viterbi decoder). However, since the evaluation of the accumulated metric in the kth step requires the knowledge of $g(\mathbf{y}_{k-L}^{k-1})$, it is also necessary that each survivor at step $k-1$ store the L-dimensional vector $\mathbf{y}_{k-L}^{k-1}$, that are the candidate transmitted samples for that survived path. To sum up, comparing the complexity of the proposed embedding strategy with that of the classical Viterbi algorithm, we get a N_s times higher computation complexity and the necessity of storing $N_s \times L$ more real numbers for each state. Of course, this is true only for the embedding step since, as stated before, decoding can be performed by standard Viterbi algorithm.

It is worth noting again that, though this way of operating closely resembles the classical Viterbi's algorithm [17], it does not lead to an optimum decoding strategy due to the infinite memory of the trellis. In fact, for each transition the distance between $\mathbf{y}_t^{t+1}$ and $\mathbf{x}_t^{t+1}$ depends also on the previous transitions due to the intrinsic memory of RDM. Nevertheless, it is expected that the higher N_s the closer the solution will be to the optimum one. Of course, if $N_s \rightarrow \infty$ the proposed algorithm coincides with the exhaustive approach, thus yielding the optimum MSE solution (with an exponentially high computational complexity). On the contrary, for $N_s = 1$ the proposed scheme works as the classical Viterbi algorithm with only one survivor per node.

4.4 Choice of good codes

The choice of optimum codes for the TC-RDM scheme presented in this paper is out of the scope of the present work. We instead consider a classical coding design for multilevel signals, namely Ungerboeck codes [18]. These codes present very good coding properties (high minimum codewords distance) and can be easily applied to the TC-RDM scenario presented in the previous section, as it will be shown in the following.

The scheme of an Ungerboeck code is shown in Fig. 3. The M input bits $\mathbf{b}_t$ at time instant t enter a binary rate $M/(M+1)$ convolutional code which outputs $M+1$ bits $\mathbf{q}_t$. Such output bits are subsequently mapped into a P-dimensional vector $\mathbf{r}_t \in \mathbb{R}^P$. Classical Ungerboeck codes have been designed for communication of signals over noisy channels, thus considering both one-dimensional (base-band signals) and two-dimensional (pass-band signals) signal sets, i.e. $P = 1,2$. For the sake of implementation simplicity we will refer in the following to the $P = 1$ case, even if all the considerations that follow can be easily extended to the $P = 2$ case.

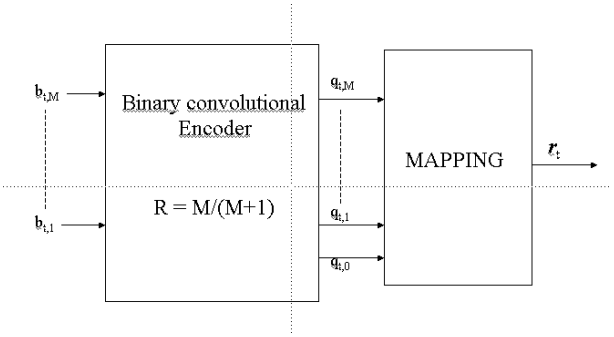

Fig. 3 *Scheme of the Ungerboeck code*

As for the binary convolutional code we refer to the systematic convolutional encoder structure with feedback shown in Fig. 4. The state of the trellis is given by the configuration of the v shift registers indicated by D in the figure and, hence, the number of possible states is 2^{Mv}. In this scheme the M most significative output bits are identical to the input bits (systematic structure). Instead, the least significative output bit q_0 depends on the state of the trellis only. The rationale for this choice derives directly from the Ungerboeck codes' design rules. We will discuss this issue in more details later on in this section. Note that the proposed codes are univocally determined by the generator polynomials $G^{(j)}(D) = 0 \times D^v + g_{v-1}^{(j)} D^{v-1} + g_{v-2}^{(j)} D^{v-2} + \cdots + g_1^{(j)} D^1 + 0 \times D^0$, and by the feedback polynomials $H^{(j)}(D) = 1 \times D^v + h_{v-1}^{(j)} D^{v-1} + h_{v-2}^{(j)} D^{v-2} + \cdots + h_1^{(j)} D^1 + 1 \times D^0$, with $j = 1, \ldots, M$.

As for the mapping rule, we consider the classical mapping by set partitioning shown in Fig. 5 for the one-dimensional case and for $M = 2$. This mapping follows from successive partitioning of a channel-signal set into subsets with increasing minimum distances (from Δ_0 to Δ_M, with $\Delta_i = 2\Delta_{i-1}$), where the choice of the actual signal to be transmitted at time instant t depends on the $M+1$ output bits. In the classical set up proposed by Ungerboeck, output bits $\mathbf{q}_t$ are used to partition the initial signal set, finally giving a single point, as shown in Fig. 5. In our case, since the goal is that of getting a good quantiser rather than a good channel code, we must associate to the resulting point a signal subset able to span all the space of real numbers. This goal is obtained by filling the space with points at distance $\Delta_{M+1} = 2\Delta_M$, taking the resulting point as reference. In this way, we get 2^{M+1} regular mono-dimensional lattices shifted one another by Δ_0 (see Fig. 6) which span all the space of real number with good quantisation properties. Seen from a different perspective, the creation of an infinite periodical lattice from a single point amounts to creating an infinite number of parallel transitions among states (without having to introduce an infinite number of free bits). At the same time, the selection of the optimum point in the lattice does not require that an exhaustive search is performed, since due to the regularity of the lattice it corresponds to a uniform scalar quantiser [Note 4].

Turning back to the code structure shown in Fig. 4, it can be easily observed that the least significant output

Note 4: Note that the scalar quantisation described here refers to the selection of the quantised ratio of a given branch of the trellis, it goes without saying that due to the memory of the Ungerboeck structure depicted in Fig. 3, an overall vector quantiser is obtained.

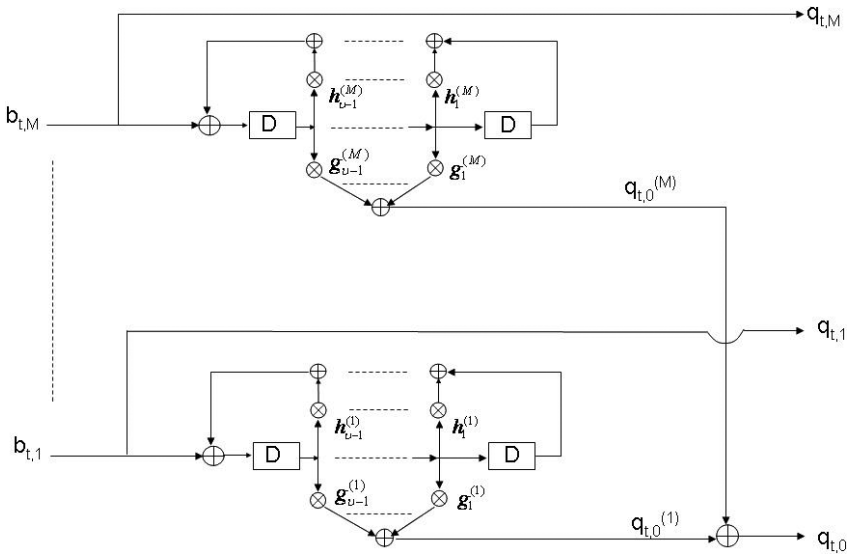

Fig. 4 *Scheme of the binary systematic convolutional encoder structure with feedback*

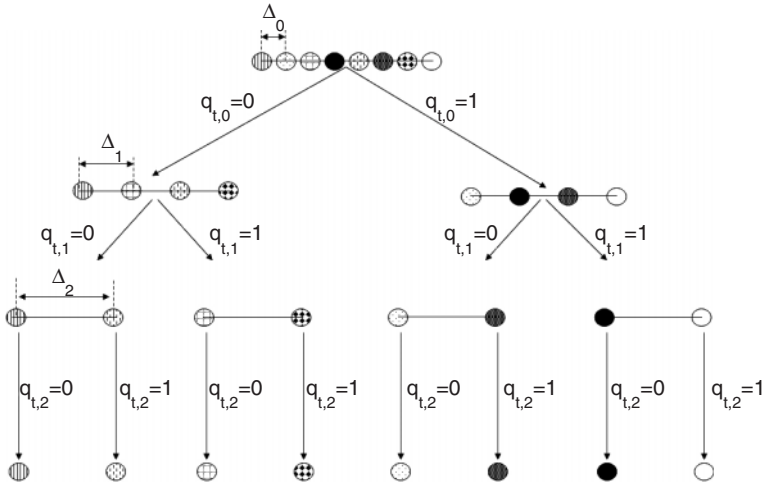

Fig. 5 *Mapping by set partitioning*

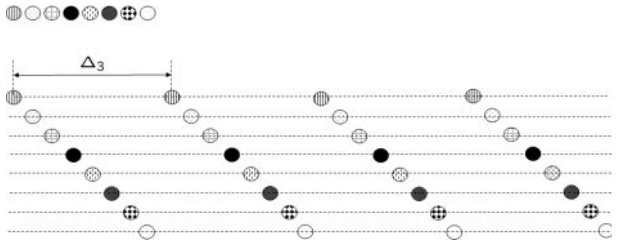

Fig. 6 *Regular mono-dimensional lattices built from the constellation points of the Ungerboeck code*

bit $q_{t,0}$ does not depend on the input bits. This means that two paths of the trellis which diverge from the same state are characterised by the same output bit $q_{t,0}$, i.e. from the mapping rule shown in Fig. 5, they are formed by points which have a distance of at least Δ_1 (e.g. paths which diverge from the same state cannot be at minimum distance Δ_0). On the other hand, it is also straightforward to observe that paths which converge to the same state, are also characterised by the same output bit $q_{t,0}$, and hence have the same distance property. Hence, the code configuration shown in Fig. 4 allows to get the important property that two paths of the trellis that diverge at a given time instant and converge after some steps have a minimum distance which is at least equal to $2\Delta_1 = \Delta_2$. This is the main property of Ungerboeck codes which allows to get channel codewords with minimum distance that is at least two times the minimum distance of the equivalent (e.g. with the same transmission rate and the same average transmission power) uncoded system.

Let us focus again to the TC-RDM problem. Given the code structure described above, it is now necessary to find the association between the informative and free bits $\mathbf{w}_t$ and $\mathbf{v}_t$ and the input bits $\mathbf{b}_t$ of the binary encoder. The natural association is to put the informative bits as the k least significant input bits, and the free bits as the $M-k$ most significant ones. In this way, the informative bits are responsible for selecting the signal subsets, i.e. they determine a bin which spans all the original space thus allowing good quantisation properties while maintaining good coding properties on account of the high free distance that can be achieved by a proper setting of the generator and feedback polynomials. A systematic procedure to find codes with good free distance properties is shown in [18]. Such a procedure has been used in this paper to get the optimum polynomials configurations which will be reported in the Results section for different v values.

4.5 TC-RDM, TCQ and signal precoding

Before ending our description of TC-RDM it is worth pointing out the relationship between TC-RDM and some techniques recently proposed in the different fields of source and channel coding. As a matter of fact, the proposed structure is akin to trellis coded quantisation (TCQ) proposed in [19] for source coding, in the sense that in TCQ a trellis path is sought that minimises a certain distortion measure with respect to the source (here the host features). In fact, the original TCQ proposal also relies on an Ungerboeck code. The main difference is the presence of a more involved trellis diagram due to the recursivity inherent to RDM. This same idea has been used in digital communications to achieve the so-called *shaping gain* [20], which in the case of precoded signals [21] holds a striking resemblance with our scheme, as the very large host power to watermark power guarantees that border effects are negligible. Furthermore, for side-informed data-hiding problems, source coding is a convenient way of closing the gap to Costa's capacity, as the lack of border effects due to the periodic replication of codewords in a certain bin helps in reducing the embedding distortion. As an example, without periodic replication, a source coder of rate $R = 1$ bit/sample using four output codewords gives for a Gaussian source a gain with respect to the Lloyd-Max quantiser of $\sim$0.14 dB, while with the periodic replication this gain is $\sim$1.4 dB. Of course, these gains afforded by the periodic structure tend to diminish as the number of output codewords increases, but the previous comparison hints at the advantages of using TCQ structures for low-size alphabets in an informed embedding context.

5 Simulation results

In the following we will show the results of the proposed TC-RDM scheme obtained through computer simulations. As for the coding scheme we consider the Ungerboeck codes described in the previous section. In particular, we consider two different codes, namely C_1 and C_2, which are both characterised by $M = 2$, M being the input dimension of the code (see Fig. 4). As for the number v of shift registers of the binary convolution code (see Fig. 4) we consider $v = 2$ for C_1 and $v = 3$ for C_2, which yields 16 and 64 states, respectively. The generator polynomials for the two codes are shown in Table 1.

The number k of informative bits per transmitted sample is taken equal to one during all simulations. This means that one of the two input bits of the code is the informative bit (the first one), while the other bit is a free bit (the second one). Moreover, we have generated i.i.d. Gaussian host features with variance σ_x^2 and we have set the minimum squared distance between constellation points of the Ungerboeck code to $\Delta_0^2 = \sigma_x^2/10^3$. As shown in simulation results later on in this section, this value of Δ_0 allows to get a distortion $DWR \cong 30$ dB which is requested in most of real-world watermarking applications. Moreover, we have considered one-dimensional codewords ($P = 1$). This means that the proposed embedding strategy allows to transmit one informative bit per each side information's sample, that is a transmission rate equal to one. For better readability of the following curves, remember that N_s is the number of survivors per each state of the suboptimum embedding strategy (see Section 4.3) and L is the number of previous samples used to evaluate the function g() (see Section 3) that has been set equal to the l_2 vector-norms, i.e.:

$$g(\mathbf{y}_{k-L}^{k-1}) = \left(\frac{1}{L}\sum_{i=1}^{L}|y_{k-i}|^2\right)^{1/2} \tag{8}$$

In order to verify the performance of the embedding scheme proposed in Section 4 we have firstly carried out computer simulations for different values of N_s. In Fig. 7 we plot the DWR of the watermarked host feature sequence versus the number of survivors N_s, obtained by using the code C_1 and for three different values of L, $L = 1$, $L = 10$ and $L = 50$. Note that for $L = 1$, increasing N_s allows to improve the embedder's performance, even if a floor is quickly achieved for $N_s \simeq 6$. For higher values of L, the floor is achieved even for lower N_s. This is due to the fact that, by increasing L, TC-RDM tends to behave as the classical TCM scheme (due to the lower fluctuations of $g(\cdot)$) where the optimum embedder is already achieved for $N_s = 1$, as in the case of $L = 50$.

In order to get an insight into the performance gain that can be obtained by TC-RDM when compared conventional RDM, in Fig. 8 the bit error probability P_e for the two codes C_1 and C_2 is given as a function of WNR in the case of $L = 10$, by considering the same rate-one Ungerboeck code described above (i.e. both TC-RDM and RDM carry one bit per host feature symbol) and by setting $N_s = 4$. We found that TC-RDM allows a performance gain of $\sim$3.5 dB for $P_e = 0.001$ and $v = 2$ (C_1 code) and a gain of $\sim$4 dB for $P_e = 0.001$ and $v = 3$ (C_2 code). The performance gain becomes even higher for lower values of P_e. In Fig. 9 the same results of Fig. 8 are shown for the case of $L = 50$. It is shown that in this case the gain obtained by TC-RDM is $\sim$2 dB for $v = 2$ and $\sim$2.5 dB for $v = 3$. Hence, the performance gain that can be obtained by TC-RDM is higher for lower L values [Note 5]. This is not surprising since for low values of L the function $g(\mathbf{z}_{k-L}^{k-1})$ computed by the receiver introduces a variability in the received signal level that is reminiscent of fading in communications. Hence, for low values of L the effect of channel coding is that of introducing both code gain and diversity gain as in the case of communication in presence of frequency-selective fading [17].

It is worth noting that for most real-world applications WNR would be less than 0, i.e. the noise may well be stronger than the watermark. In order to let the system working at lower WNR (and, of course, lower transmission rate) it is necessary to implement a form of channel precoding. As an example, the classical spread transform-dither modulation (ST-DM) approach, which allows to get a W dB gain with respect to DM (W is the spreading factor in dB) could be easily extended to

Table 1: Generator polynomials for the two codes used during simulations, named C_1 and C_2

	$C_1: M = 2, v = 2$	$C_2: M = 2, v = 3$
$G^{(1)}(D)$	D	$D^2 + D$
$G^{(2)}(D)$	D	D^2
$H^{(1)}(D)$	$D^2 + D + 1$	$D^3 + 1$
$H^{(2)}(D)$	$D^2 + 1$	$D^3 + D^2 + 1$

Note 5: Note that this is only a relative behaviour with respect to RDM, since in absolute terms the TC-RDM system is more robust for larger L, as is clear from a comparison of Figs 5 and 6.

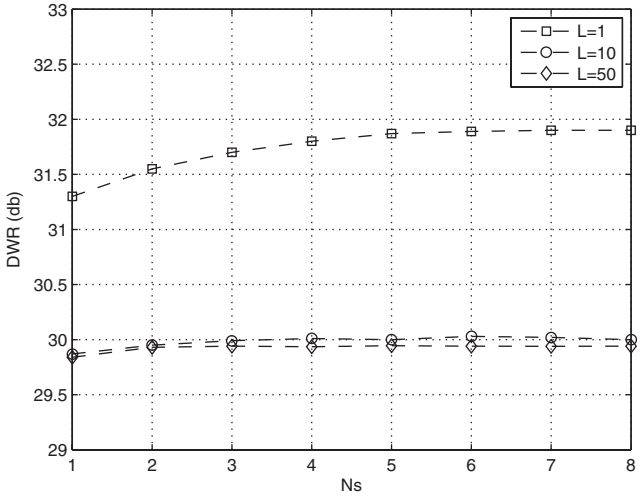

Fig. 7 *DWR against number of survivors for various L values*
As the number of survivors increases the suboptimum embedding algorithm tends to the optimum one, in that no further improvement of DWR is expected

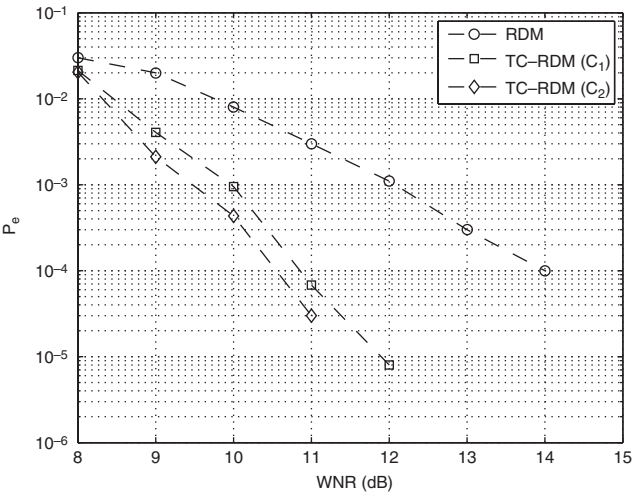

Fig. 9 *Comparison between TC-RDM and conventional RDM for $L = 50$*
The improvement brought by TC-RDM is less evident than in the $L = 10$ case on account of the reduced diversity gain

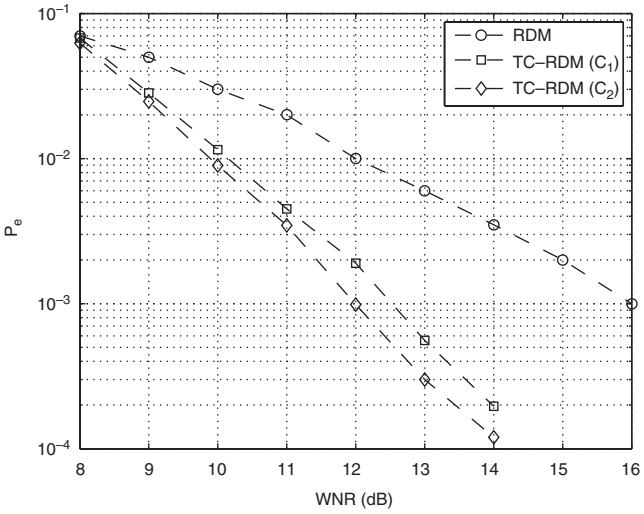

Fig. 8 *Comparison between TC-RDM and conventional RDM for $L = 10$*
For low values of P_e the improvement brought by TC-RDM is evident. The plot has been obtained by embedding one bit for each sample of the host feature sequence; lower bit error rates are expected for lower embedding rates (channel coding)

TC-RDM embedding strategy, thus leading to ST-TC-RDM. In this case, performance curves would be simply shifted back by W dB. Of course, the transmission rate would be reduced by W dB as well. Alternative techniques to spread transform for channel precoding could be investigated with the aim of performing better than ST-TC-RDM. Such an issue is out of the scope of this paper and will be investigated in future works.

6 Conclusions

A vector extension of the RDM algorithm for gain-invariant QIM watermarking has been presented. The vector extension, somewhat relying on the same rationale of TCQ, is obtained by introducing the concept of free bits, i.e. a set of bits at the input of a trellis code that can be adjusted in such a way to minimise the embedding distortion. From a different perspective, the same results can be obtained by means of a redundant trellis (like the dirty trellis described in [6]) whereby several paths can be associated to the same input message. According to the RDM paradigm, the redundant trellis is used to quantise a sequence of rational functions of the host feature sequence. Due to the memory introduced by RDM, optimum embedding would require an exhaustive search over all the possible paths on the trellis, hence we introduced a sub-optimum embedding scheme that permits to approach the performance of the optimum scheme at a reduced computational cost. The performance improvement allowed by TC-RDM has been evaluated by testing it on a sequence of i.i.d. normal features, obtaining a gain ranging from 2 through 4 dB with respect to RDM.

It is worth noting that the, dare we say, 'vectoriality' of the proposed quantiser could be augmented by increasing the output dimension P of the trellis, which is considered equal to one in this paper. Of course, increasing P may lead to performance benefits at the expenses, however, of implementation complexity increase. Indeed, in order to get a good multi-dimensional mapping it is necessary to increase the dimension of the code (i.e. to increase M). On the other hand, the increase in the convolutional encoder memory is an alternative way for getting an higher vectoriality and an higher complexity. An interesting question that arises is that, for a given implementation complexity, $P > 1$ allows to get better performance than $P = 1$. This issue will be investigated in future works.

Future works will also include the introduction of distortion compensation and channel coding to further improve the robustness of the algorithm. The possibility of increasing the security of the system by randomising the rational function RDM relies on will be also investigated. Finally, the gap between theory and practice will need to be covered, by applying TC-DRM to the watermarking of real data such as audio, still images or video sequences.

7 References

1 Chen, B., and Wornell, G.: 'Quantization index modulation: a class of provably good methods for digital watermarking and

information embedding'. *IEEE Trans. Inf. Theory*, 2001, **47**, pp. 1423–1443
2 Eggers, J.J., Bäuml, R., Tzschoppe, R., and Girod, B.: 'Scalar Costa scheme for information embedding'. *IEEE Trans. Signal Process.*, 2003, **4**, pp. 1003–1019
3 Bartolini, F., Barni, M., and Piva, A.: 'Performance analysis of st-dm watermarking in presence of non-additive attacks'. *IEEE Trans. Signal Process.*, 2002, **52**, pp. 2965–2974
4 Eggers, J.J., Bäuml, R., and Girod, B.: 'Estimation of amplitude modifications before SCS watermark detection'. In Wong, P.W., and Delp, E.J. (Eds): Proc. SPIE Security and Watermarking of Multimedia Contents IV, San Jose, CA, 19-25 Jan. 2002, (*Proc. SPIE*, **4675**), pp 387–398
5 Lee, K., Kim, D.S., Kim, T., and Moon, K.A.: 'Em estimation of scale factor for quantization-based audio watermarking'. In Kalker, T., and C.I.Y. (Eds): Proc. Second Int. Workshop on Digital Watermarking, IWDW, Seoul, Korea, 20-22 Oct. 2003, pp. 316–327
6 Miller, M.L., Doerr, G.J., and Cox, I.J.: 'Applying informed coding and embedding to design a robust, high capacity, watermark'. *IEEE Trans. Image Process.*, 2004, **13**, pp. 792–807
7 Abrardo, A., and Barni, M.: 'Informed watermarking by means of orthogonal and pseudo-random dirty paper coding'. *IEEE Trans. Signal Process.*, 2005, **53**, pp. 824–833
8 Oostven, J., Kalker, T., and Staring, M.: 'Adaptive quantization watermarking'. In Wong, P.W., and Delp, E.J. (Eds): Proc. SPIE Security, Steganography, and Watermarking of Multimedia Contents VI, San Jose, CA, Jan. 2004, (*Proc. SPIE*, **5306**), pp. 296–303
9 Li, Q., and Cox, I.: 'Using perceptual models to improve fidelity and provide invariance to valumetric scaling for quantization index modulation watermarking'. Proc. IEEE Int. Conf. on Acoustic Speech and Signal Processing, ICASSP'05, Philadelphia, PA, USA, March 2005, pp. 18–23
10 Perez-Gonzalez, F., Mosquera, C., Barni, M., and Abrardo, A.: 'Rational dither modulation: a novel data-hiding method robust to value-metric scaling attacks'. MMSP 2004, IEEE Workshop on Multimedia Signal Processing. Siena, Italy, Sept. 2004
11 Perez-Gonzalez, F., Mosquera, C., Barni, M., and Abrardo, A.: 'Ensuring gain invariance in high-rate data-hiding'. In Wong, P.W., and Delp, E.J. (Eds): Proc. SPIE Security, Steganography, and Watermarking of Multimedia Contents VII, San Jose, CA, Jan. 2005, (*Proc. SPIE*, **5681**), pp. 206–217
12 Lib, T., Venkatesan, R., and Mihcak, M.K.: 'Scale-invariant image watermarking via optimization algorithms for quantizing randomized statistics'. Proc. ACM Multimedia and Security Workshop, Magdeburg, Germany, Sept. 2004, pp. 20–21
13 Perez-Gonzalez, F., Mosquera, C., Barni, M., and Abrardo, A.: 'Rational dither modulation: a high-rate data-hiding method invariant to gain attacks'. *IEEE Trans. Signal Process.*, 2005, **53**, pp. 3960–3975
14 Cox, I.J., Kilian, J., Leighton, T., and Shamoon, T.: 'Secure spread spectrum watermarking for multimedia'. *IEEE Trans. Image Process.*, 1997, **6**, pp. 1673–1687
15 Malvar, H.S., and Florencio, D.A.F.: 'Improved spread spectrum: a new modulation technique for robust watermarking'. *IEEE Trans. Image Process.*, 2003, **51**, pp. 898–905
16 Miller, M.L., Doerr, G.J., and Cox, I.J.: 'Dirty-paper trellis codes for watermarking'. Proc. 9th IEEE Int. Conf. Image Processing. ICIP'02, Rochester, NY, USA, 2002, vol. II. pp. 129–132
17 Proakis, J.G.: 'Digital communications' (McGraw-Hill, 1989, 2nd edn.)
18 Ungerboeck, G.: 'Channel coding with multilevel/phase signals' *IEEE Trans. Inf. Theory*, 1982, **IT-28**, pp. 55–67
19 Marcellin, M., and Fischer, T.: 'Trellis coded quantization of memoryless and gaussmarkov sources'. *IEEE Trans. Commun.*, 1990, **36**, pp. 82–93
20 Forney, Jr, G.: 'Trellis shaping,' *IEEE Trans. Inf. Theory*, 1992, **38**, pp. 281–300
21 Fischer, R.: 'Precoding and signal shaping for digital transmission' (John Wiley and Sons, 2002)