To appear in *IEEE Trans. Inform. Theory.*

Quantization Index Modulation: A Class of Provably Good Methods for Digital Watermarking and Information Embedding

Brian Chen and Gregory W. Wornell

Submitted June 1999
Revised September 2000

Abstract

We consider the problem of embedding one signal (e.g., a digital watermark), within another "host" signal to form a third, "composite" signal. The embedding is designed to achieve efficient trade-offs among the three conflicting goals of maximizing information-embedding rate, minimizing distortion between the host signal and composite signal, and maximizing the robustness of the embedding.

We introduce new classes of embedding methods, termed quantization index modulation (QIM) and distortion-compensated QIM (DC-QIM), and develop convenient realizations in the form of what we refer to as dither modulation. Using deterministic models to evaluate digital watermarking methods, we show that QIM is "provably good" against arbitrary bounded and fully-informed attacks, which arise in several copyright applications, and in particular it achieves provably better rate-distortion-robustness trade-offs than currently popular spread-spectrum and low-bit(s) modulation methods. Furthermore, we show that for some important classes of probabilistic models, DC-QIM is optimal (capacity-achieving) and regular QIM is near-optimal. These include both additive white Gaussian noise channels, which may be good models for hybrid transmission applications such as digital audio broadcasting, and mean-square-error constrained attack channels that model private-key watermarking applications.

Index Terms—dither modulation, quantization index modulation, information embedding, digital watermarking, steganography, data hiding, digital audio broadcasting, hybrid transmission

1 Introduction

A number of applications have emerged recently [1] that require the design of systems for embedding one signal, sometimes called an "embedded signal" or "watermark", within another signal, called a "host signal". The embedding must be done such that the embedded signal is "hidden," i.e., causes no serious degradation to its host. At the same time, the embedding must be robust to common

This work has been supported in part by MIT Lincoln Laboratory Advanced Concepts Committee, the National Science Foundation under Grant No. CCR-0073520, Microsoft Research, and a National Defense Science and Engineering Graduate Fellowship.

B. Chen and G. W. Wornell are with the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139. (e-mail: {bchen,gww}@allegro.mit.edu).

degradations of the watermarked signal—the watermark must survive whenever the host signal does. In some applications these degradations are the result of benign processing and transmission; in other cases they result from deliberate attacks.

Several of these applications relate to copyright notification and enforcement for audio, video, and images that are distributed in digital formats. In these cases the embedded signal either notifies a recipient of any copyright or licensing restrictions or inhibits or deters unauthorized copying. For example, this embedded signal could be a digital "fingerprint" that uniquely identifies the original purchaser of the copyrighted work. If illicit copies of the work were made, all copies would carry this fingerprint, thus identifying the owner of the copy from which all illicit copies were made. In another example, the embedded signal could either enable or disable copying by some duplication device that checks the embedded signal before proceeding with duplication. Such a system has been proposed for allowing a copy-once feature in digital video disc recorders [2]. Alternatively, a standards-compliant player could check the watermark before deciding whether or not to play the disc [3].

Other applications include automated monitoring of airplay of advertisements on commercial radio broadcasts. Advertisers can embed a digital watermark within their ads and count the number of times the watermark occurs during a given broadcast period, thus ensuring that their ads are played as often as promised. In other applications, the embedded signal may be used for authentication of—or detection of tampering with—the host signal. For example, a digital signature could be embedded in a military map. A number of other national security applications are described in [4] and include covert communication, sometimes called "steganography" or low probability of detection communication, and so-called traitor tracing, a version of the digital fingerprinting application described above used for tracing the source of leaked information.

One final application for which the digital watermarking methods developed in this paper are well-suited is the backwards-compatible upgrading of an existing communication system, an example of which is so-called hybrid in-band on-channel digital audio broadcasting [5, 6]. In this application one would like to simultaneously transmit a digital signal with existing analog (AM and/or FM) commercial broadcast radio without interfering with conventional analog reception. Thus, the analog signal is the host signal and the digital signal is the watermark. Since the embedding does not degrade the host signal too much, conventional analog receivers can demodulate the analog host signal. In addition, next-generation digital receivers can decode the digital signal embedded within the analog signal, which may be all or part of a digital audio signal, an enhancement signal used to refine the analog signal, or simply supplemental information such as station

identification or traffic information. More generally, the host signal in these hybrid transmission systems could be some other type of analog signal such as video [7], or even a digital waveform—for example, a digital pager signal could be embedded within a digital cellular telephone signal.

In general, designers of information embedding systems for these kinds of applications seek to achieve high embedding rates with high levels of robustness and low levels of embedding-induced distortion. However, in general these three goals are conflicting. Thus, in this paper we characterize methods in terms of the efficiency with which they trade off rate, distortion, and robustness. For instance, for any minimum embedding rate requirement and maximum acceptable level of embedding distortion, the more efficient an embedding method is, the higher the robustness that can be achieved.

A great many information-embedding algorithms have been proposed [1] in this still emerging field. Some of the earliest proposed methods [8, 9, 7] employ a quantize-and-replace strategy: after first quantizing the host signal, these systems change the quantization value to embed information. A simple example of such a system is so-called low-bit(s) modulation (LBM), where the least significant bit(s) in the quantization of the host signal are replaced by a binary representation of the embedded signal. More recently, additive spread-spectrum based methods, which embed information by linearly combining the host signal with a small pseudo-noise signal that is modulated by the embedded signal, have received considerable attention in the literature as an alternative to LBM-type methods [10, 11, 12, 13].

In this paper we show that both LBM-type strategies and additive spread-spectrum are in general *not* good choices for most information embedding and digital watermarking applications. As an alternative, this paper introduces a new class of information embedding strategies we refer to as "quantization index modulation (QIM)" that is in general preferable and in many specific scenarios optimal. We further develop computationally efficient implementations of QIM in the form of what we refer to as "dither modulation." We evaluate both specific realizations of uncoded and coded QIM, and the asymptotic performance limits of coded QIM using information-theoretic analysis. Other emerging information theoretic results on the digital watermarking problem are developed in, e.g., [14, 15, 16, 17, 18, 19, 20].

The specific organization of the paper is as follows. In Sec. 2 we develop two useful equivalent models for the information-embedding problem. In Sec. 3, we classify traditional approaches to this problem, and in the process identify some of their shortcomings. Sec. 4 introduces the QIM class of embedding methods, and Sec. 5 develops practical realizations that are compared to corresponding implementations of traditional approaches. Next, Sec. 6 establishes conditions under which different

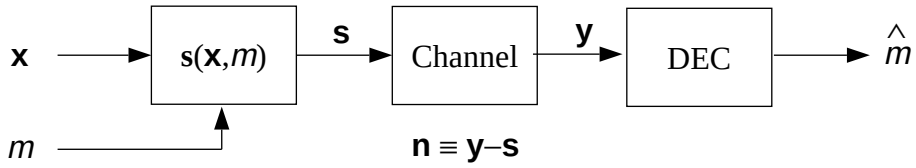

Figure 1: General information-embedding problem model. A message m is embedded in the host signal vector $\mathbf{x}$ using some embedding function $\mathbf{s}(\mathbf{x}, m)$. A perturbation vector $\mathbf{n}$ corrupts the composite signal $\mathbf{s}$. The decoder extracts an estimate $\hat{m}$ of m from the noisy channel output $\mathbf{y}$.

forms of QIM are optimal in an information theoretic sense. We then evaluate the methods of this paper in the context of Gaussian models for unintentional attacks in Sec. 7, and in the context of some general intentional attack models in Sec 8. Finally, Sec. 9 contains some concluding remarks.

2 Problem Model

Two mathematically equivalent models for the information embedding problem are useful in our development.

2.1 Distortion-Constrained Multiplexing Model

The information-embedding problem is naturally and generally described by Fig. 1. In this figure, there is a host signal vector $\mathbf{x} \in \mathbb{R}^N$ into which we wish to embed some information m.[1] We wish to embed at a rate of R_m bits per dimension (host signal sample) so we can think of m as an integer in the set $\{1, 2, \ldots, 2^{NR_m}\}$.

An embedding function maps the host signal $\mathbf{x}$ and embedded information m to a composite signal $\mathbf{s} \in \mathbb{R}^N$ subject to some distortion constraint. Various distortion measures may be of interest, an example of which is the squared-error distortion

$$D(\mathbf{s}, \mathbf{x}) = \frac{1}{N}\|\mathbf{s} - \mathbf{x}\|^2. \tag{1}$$

The composite signal $\mathbf{s}$ is subjected to various common signal processing manipulations such as lossy compression, addition of random noise, and resampling, as well as deliberate attempts to remove the embedded information. These manipulations occur in some channel, which produces an

[1]The vector $\mathbf{x}$ is any convenient representation of all or part of the host signal. In the case of a host image, it could be a vector of pixel values or Discrete Cosine Transform (DCT) coefficients, for example. In the case of a host audio waveform, this vector could be a vector of samples, spectral parameters, or linear prediction coding (LPC) coefficients, for example.

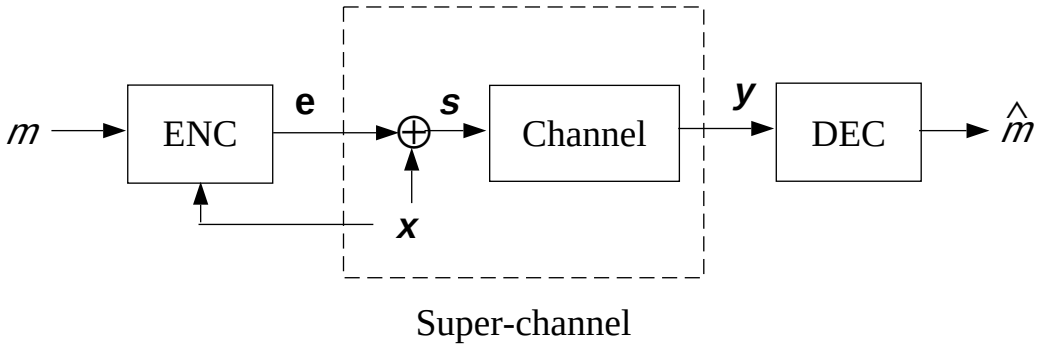

Figure 2: Equivalent super-channel model for information embedding. The composite signal is the sum of the host signal, which is the state of the super-channel, and a host-dependent distortion signal.

output signal $\mathbf{y} \in \mathbb{R}^N$. For future convenience, we define a perturbation vector to be the difference $\mathbf{n} \in \mathbb{R}^N$, as shown in Fig. 1; we consider cases of both signal-independent and signal-dependent, perturbation vectors in this paper.

A decoder extracts—i.e., forms an estimate $\hat{m}$ of—the embedded information m based on the channel output $\mathbf{y}$. We focus primarily on the "host-blind" case of interest in most applications, where $\mathbf{x}$ is not available to the decoder, in contrast to the "known-host" case, where the decoder can separately observe $\mathbf{x}$. (See, e.g., [14] [17] for information-theoretic treatments of some aspects of the known-host case.) Our interest is in decoders that produce reliable estimates whenever the channel is not too severe, where reliable means either that $\hat{m} = m$ deterministically or that $\Pr[\hat{m} \neq m] < \epsilon$ for sufficiently small ϵ. In such cases, the tolerable severity of the channel degradations is a measure of the robustness of an information embedding system.

2.2 Equivalent Super-channel Model

An alternative representation of the model of Fig. 1 is shown in Fig. 2. The two models are equivalent since any embedding function $\mathbf{s}(\mathbf{x}, m)$ can be written as the sum of the host signal $\mathbf{x}$ and a host-dependent distortion signal $\mathbf{e}(\mathbf{x}, m)$, i.e., $\mathbf{s}(\mathbf{x}, m) = \mathbf{x} + \mathbf{e}(\mathbf{x}, m)$, simply by defining the distortion signal to be $\mathbf{e}(\mathbf{x}, m) \stackrel{\Delta}{=} \mathbf{s}(\mathbf{x}, m) - \mathbf{x}$. Thus, one can view $\mathbf{e}$ as the input to a super-channel that consists of the cascade of an adder and the true channel. The host signal $\mathbf{x}$ is a state of this super-channel that is known at the encoder. The measure of distortion $D(\mathbf{s}, \mathbf{x})$ between the composite and host signals maps onto a host-dependent measure of the size $P(\mathbf{e}, \mathbf{x}) = D(\mathbf{x} + \mathbf{e}, \mathbf{x})$

of the distortion signal $\mathbf{e}$. For example, squared error distortion (1) equals the power of $\mathbf{e}$,

$$\frac{1}{N}\|\mathbf{s}-\mathbf{x}\|^2 = \frac{1}{N}\|\mathbf{e}\|^2.$$

Therefore, one can view information embedding problems as power-limited communication over a super-channel with a state that is known at the encoder.[2] As we will develop, this view will be convenient for determining achievable rate-distortion-robustness trade-offs of various information embedding and decoding methods.

2.3 Channel Models

In general, the channel model is either a characterization of the degradations that can actually occur to the composite signal, or alternatively, a description of the class of degradations to which the embedder and decoder must be robust, i.e., the system is designed to work against all degradations described by this particular model. The latter viewpoint is particularly useful in the context of intentional attacks.

We consider both probabilistic and deterministic channel models. In the probabilistic case, we specify the channel input-output relationship in terms of the conditional probability law $p_{\mathbf{y}|\mathbf{s}}(\mathbf{y}|\mathbf{s})$. Implicitly, this specification also describes the conditional probability law of the perturbation vectors against which the system must be robust since $p_{\mathbf{n}|\mathbf{s}}(\mathbf{n}|\mathbf{s}) = p_{\mathbf{y}|\mathbf{s}}(\mathbf{s}+\mathbf{n}|\mathbf{s})$. In the deterministic case, the channel input-output relationship is described most generally in terms of the set of possible outputs $\mathcal{P}\{\mathbf{y}|\mathbf{s}\}$ for every given input, or equivalently, in terms of the set of desired tolerable perturbation vectors $\mathcal{P}\{\mathbf{n}|\mathbf{s}\}$ for every given input.

3 Classes of Embedding Methods

An extremely large number of embedding methods have been proposed in the literature [22, 23, 1]. Broadly, for our purposes these can be divided into two classes: (1) host-interference non-rejecting methods and (2) host-interference rejecting methods.

Host-interference non-rejecting methods have the general property that the host signal is effectively a source of interference in the system, and generally result from system designs that do not allow the encoder in Fig. 2 to sufficiently exploit knowledge of the host signal $\mathbf{x}$.

[2] Cox, *et al.*, have also recognized that one may view watermarking as communications with side information known at the encoder [21].

The simplest of such methods have purely additive embedding functions of the form

$$\mathbf{s}(\mathbf{x}, m) = \mathbf{x} + \mathbf{w}(m), \tag{2}$$

where $\mathbf{w}(m)$ is typically a pseudo-noise sequence. Such embedding methods are often referred to as additive spread-spectrum methods, and some of the earliest examples are described in [24, 25, 10, 26, 11, 12]. Typically, $\mathbf{w}(m)$ takes the form

$$\mathbf{w}(m) = a(m)\mathbf{v} \tag{3}$$

where $\mathbf{v}$ is a unit-energy spreading vector and $a(m)$ is a scalar function of the message.[3]

It is often convenient to view additive spread-spectrum as perturbation of a projection. In particular, substituting (3) into (2) and using that $\mathbf{v}$ has unit energy, we obtain

$$\mathbf{s} = \mathbf{x} + a(m)\mathbf{v}, \tag{4}$$

which when projected onto $\mathbf{v}$ we obtain

$$\tilde{s} = \mathbf{s}^T\mathbf{v} = \tilde{x} + a(m) \tag{5}$$

where $\tilde{x}$ is the corresponding projection of the host signal, i.e.,

$$\tilde{x} = \mathbf{x}^T\mathbf{v}. \tag{6}$$

Finally, substituting (5) back into (4) yields the composite signal reconstruction from projections

$$\mathbf{s} = \mathbf{x} + (\tilde{s} - \tilde{x})\mathbf{v}. \tag{7}$$

From (2), we see that for this class of embedding methods, the host signal $\mathbf{x}$ acts as additive interference that inhibits the decoder's ability to estimate m. Consequently, even in the absence of any channel perturbations ($\mathbf{n} = \mathbf{0}$), one can usually embed only a small amount of information. Thus, these methods are useful primarily when either the host signal is available at the decoder

[3]Technically, spread-spectrum systems (2) for which (3) applies are classified as amplitude-modulation additive spread-spectrum methods, but since there is no risk of confusion in this paper we will use the term "additive spread-spectrum" to specifically mean those systems based on amplitude-modulation.

(as assumed in, e.g., [26]) or when the host signal interference is much smaller than the channel interference.

Information embedding systems can achieve host-interference rejection when knowledge of the host signal at the encoder is adequately exploited in system design. Examples include LBM and, more generally, quantize-and-replace systems. In LBM systems, the least significant bit(s) in the binary representation of a host sample are simply replaced with message bits. A class of quantize-and-replace systems that we refer to as generalized LBM systems implement a vector generalization of this embedding strategy. Generalized LBM embedding functions are of the form

$$\mathbf{s} = \mathbf{q}(\mathbf{x}) + \mathbf{d}(m), \tag{8}$$

where $\mathbf{q}(\cdot)$ represents the coarse quantizer that determines the most significant bits, and $\mathbf{d}$ is determined only by the (modulated) least significant bits. A defining characteristic of generalized LBM systems is that the embedding never alters the most significant bits of the host signal, which is expressed in terms of the constraint

$$\mathbf{q}(\mathbf{s}) = \mathbf{q}(\mathbf{x}). \tag{9}$$

Without loss of generality, we may assume that good generalized LBM quantizers are unbiased, i.e.,

$$E[\mathbf{q}(\mathbf{x}) - \mathbf{x}] = \mathbf{0}. \tag{10}$$

One example of a generalized LBM system is that developed in [7], where LBM is effectively applied to a pseudorandom projection of the form (6). Thus, the embedding is of the form (7) where $\tilde{s}$ is now of the form

$$\tilde{s} = \mathbf{s}^T\mathbf{v} = q(\tilde{x}) + d(m), \tag{11}$$

with $q(\cdot)$ a uniform, scalar quantization function of step size Δ and $d(m)$ a perturbation value. It is convenient to think of this class of generalized LBM systems as "spread LBM" systems.

While generalized LBM systems are host-interference rejecting, they are unnecessarily constrained in a way that makes them generally inefficient and vulnerable to various classes of attacks, which in turn limits the range of applications for which they can be used. Avoiding these constraints in the process of developing optimal information embedding systems naturally gives rise to a new and general class of host-interference rejecting embedding methods called quantization index modulation (QIM), which we develop in the sequel.

4 Quantization Index Modulation

To develop the quantization index modulation concept, we begin by viewing the embedding function $\mathbf{s}(\mathbf{x}, m)$ as an ensemble of functions of $\mathbf{x}$, indexed by m. We denote the functions in this ensemble as $\mathbf{s}(\mathbf{x}; m)$ to emphasize this view. If the embedding-induced distortion is to be small, then each function in the ensemble must be close to an identity function in some sense so that

$$\mathbf{s}(\mathbf{x}; m) \approx \mathbf{x}, \qquad \forall m. \tag{12}$$

That the system needs to be robust to perturbations suggests that the points in the range of one function in the ensemble should be far away in some sense from the points in the range of any other function. For example, one might desire at the very least that the ranges be non-intersecting. Otherwise, even in the absence of any perturbations, there will be some values of $\mathbf{s}$ from which one will not be able to uniquely determine m. In fact, it is precisely the non-intersection property that leads to host-signal interference rejection.

The non-intersection property along with the approximate-identity property (12), which suggests that the ranges of each of the functions "cover" the space of possible (or at least highly probable) host signal values $\mathbf{x}$, suggests that the functions be discontinuous. Quantizers are just such a class of discontinuous, approximate-identity functions. Then, "quantization index modulation (QIM)" refers to embedding information by first modulating an index or sequence of indices with the embedded information and then quantizing the host signal with the associated quantizer or sequence of quantizers.

Fig. 3 illustrates this QIM information-embedding technique. In this example, one bit is to be embedded so that $m \in \{1, 2\}$. Thus, we require two quantizers, and their corresponding sets of reconstruction points in $\mathbb{R}^N$ are represented in Fig. 3 with $\times$'s and $\circ$'s. If $m = 1$, the host signal is quantized with the $\times$-quantizer, i.e., $\mathbf{s}$ is chosen to be the $\times$ closest to $\mathbf{x}$. If $m = 2$, $\mathbf{x}$ is quantized with the $\circ$-quantizer.

As $\mathbf{x}$ varies, the composite signal value $\mathbf{s}$ varies from one $\times$ point ($m = 1$) to another or from one $\circ$ point ($m = 2$) to another, but it never varies between a $\times$ point and a $\circ$ point. Thus, even with an infinite energy host signal, one can determine m if channel perturbations are not too severe. The $\times$ points and $\circ$ points are both quantizer reconstruction points and signal constellation points,[4] and we may view design of QIM systems as the simultaneous design of an ensemble of source codes

[4] One *set* of points, rather than one individual point, exists for each value of m.

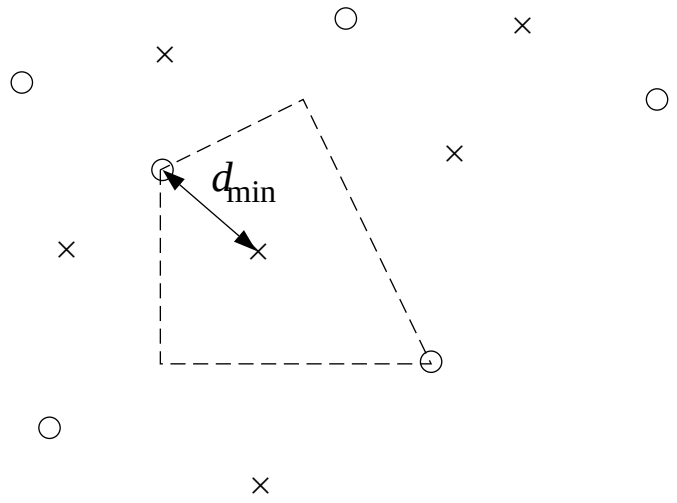

Figure 3: Quantization index modulation for information embedding. The points marked with $\times$'s and $\circ$'s belong to two different quantizers, each with its associated index. The minimum distance $d_{\min}$ measures the robustness to perturbations, and the sizes of the quantization cells, one of which is shown in the figure, determine the distortion. If $m = 1$, the host signal is quantized to the nearest $\times$. If $m = 2$, the host signal is quantized to the nearest $\circ$.

(quantizers) and channel codes (signal constellations).

Conveniently, properties of the quantizer ensemble can be related directly to the performance parameters of rate, distortion, and robustness. For example, the number of quantizers in the ensemble determines the information-embedding rate R_m. The sizes and shapes of the quantization cells determine the embedding-induced distortion, all of which arises from quantization error. Finally, for many classes of channels, the minimum distance

$$d_{\min} \stackrel{\Delta}{=} \min_{(i,j):i\neq j} \min_{(\mathbf{x}_i,\mathbf{x}_j)} \|\mathbf{s}(\mathbf{x}_i; i) - \mathbf{s}(\mathbf{x}_j; j)\| \tag{13}$$

between the sets of reconstruction points of different quantizers in the ensemble effectively determines the robustness of the embedding.[5]

It is important to emphasize that, in contrast to the case where the host signal $\mathbf{x}$ is known at the receiver, the minimum distance decoder needs to choose from all reconstruction points of the

[5]When the host signal is known at the decoder, as is the case in some applications of interest, then the more natural minimum distance is

$$d_{\min}(\mathbf{x}) \stackrel{\Delta}{=} \min_{(i,j):i\neq j} \|\mathbf{s}(\mathbf{x}; i) - \mathbf{s}(\mathbf{x}; j)\|, \qquad \text{or} \qquad d_{\min} \stackrel{\Delta}{=} \min_{\mathbf{x}} \min_{(i;j):i\neq j} \|\mathbf{s}(\mathbf{x}; i) - \mathbf{s}(\mathbf{x}; j)\|.$$

quantizers, not just those corresponding to the actual host signal $\mathbf{x}$. In particular, the minimum distance decoder makes decisions according to the rule[6]

$$\hat{m}(\mathbf{y}) = \arg\min_{m} \min_{\mathbf{x}} \|\mathbf{y} - \mathbf{s}(\mathbf{x}; m)\|. \tag{14}$$

If, which is often the case, the quantizers $\mathbf{s}(\mathbf{x}; m)$ map $\mathbf{x}$ to the nearest reconstruction point, then (14) can be rewritten as

$$\hat{m}(\mathbf{y}) = \arg\min_{m} \|\mathbf{y} - \mathbf{s}(\mathbf{y}; m)\|. \tag{15}$$

(While the minimum distance decoder is especially convenient to implement and analyze, a variety of other potentially useful decoders are discussed in [27].)

Intuitively, the minimum distance measures the size of perturbation vectors that can be tolerated by the system. For example, if channel perturbations are bounded according to[7]

$$\|\mathbf{y} - \mathbf{s}\|^2 = \|\mathbf{n}\|^2 \leq N\sigma_n^2. \tag{16}$$

then the minimum distance decoder is guaranteed to not make an error as long as

$$\frac{d_{\min}^2}{4N\sigma_n^2} > 1. \tag{17}$$

In the case of a classical additive white Gaussian noise channel with a noise variance of σ_n^2, at high signal-to-noise ratio (SNR) the minimum distance also characterizes the error probability of the minimum distance decoder [28],

$$\Pr[\hat{m} \neq m] \sim Q\left(\sqrt{\frac{d_{\min}^2}{4\sigma_n^2}}\right),$$

where $Q(\cdot)$ is the Gaussian Q-function,

$$Q(x) = \frac{1}{\sqrt{2\pi}} \int_x^\infty e^{-t^2/2}\, dt. \tag{18}$$

[6]Alternatively, if the host signal $\mathbf{x}$ is known at the decoder,

$$\hat{m}(\mathbf{y}, \mathbf{x}) = \arg\min_{m} \|\mathbf{y} - \mathbf{s}(\mathbf{x}; m)\|.$$

[7]We refer to this as the bounded perturbation channel and will revisit this deterministic channel in Sec. 8.2.1.

4.1 Distortion-Compensated QIM

Distortion compensation is a type of post-quantization processing that can improve the achievable rate-distortion-robustness trade-offs of QIM methods. To see this, we begin by noting that for a fixed rate and a given quantizer ensemble, scaling[8] all quantizers by $\alpha \leq 1$ increases $d^2_{\min}$ by a factor of $1/\alpha^2$, thereby increasing the robustness of the embedding. However, the embedding-induced distortion also increases by a factor of $1/\alpha^2$. Adding back a fraction $1-\alpha$ of the quantization error to the quantization value removes, or compensates for, this additional distortion. The resulting embedding function is

$$\mathbf{s}(\mathbf{x}, m) = \mathbf{q}(\mathbf{x}; m, \Delta/\alpha) + (1-\alpha)[\mathbf{x} - \mathbf{q}(\mathbf{x}; m, \Delta/\alpha)], \tag{19}$$

where $\mathbf{q}(\mathbf{x}; m, \Delta/\alpha)$ is the mth quantizer of an ensemble whose reconstruction points have been scaled by α so that two reconstruction points separated by a distance Δ before scaling are separated by a distance Δ/α after scaling. The first term in (19) represents normal QIM embedding. We refer to the second term as the distortion-compensation term.

The quantization error added back is a source of interference to the decoder. Typically, the probability density functions of the quantization error for all quantizers in the QIM ensemble are similar. Therefore, the distortion compensation term in (19) is effectively statistically independent of m and can be treated as independent noise. Thus, decreasing α leads to greater minimum distance, but for a fixed embedding-induced distortion, the distortion-compensation interference at the decoder increases. One optimality criterion for choosing α is to maximize the following "SNR" at the decision device:

$$\mathrm{SNR}(\alpha) = \frac{d_1^2/\alpha^2}{(1-\alpha)^2 \frac{D_{\mathbf{s}}}{\alpha^2} + \sigma_n^2} = \frac{d_1^2}{(1-\alpha)^2 D_{\mathbf{s}} + \alpha^2 \sigma_n^2},$$

where this SNR is defined as the ratio between the squared minimum distance between quantizers and the total interference energy from both distortion-compensation interference and channel interference. Here, d_1 is the minimum distance when $\alpha = 1$ and is a characteristic of the particular quantizer ensemble. One can easily verify that the optimal scaling parameter α that maximizes this SNR is

$$\alpha_{\mathrm{opt}} = \frac{\mathrm{DNR}}{\mathrm{DNR}+1}, \tag{20}$$

[8] If a reconstruction point is at $\mathbf{q}$, it is "scaled" by α by moving it to $\mathbf{q}/\alpha$.

where DNR is the (embedding-induced) distortion-to-noise ratio $D_{\mathbf{s}}/\sigma_n^2$.

As we will see, suitably coded versions of this distortion-compensated QIM with precisely the parameter setting (20) also have important asymptotic optimality properties. Before developing these properties, let us first investigate some constraints that are useful to impose on QIM systems to facilitate implementation.

5 Dither Modulation: An Implementation of QIM

A key aspect of the design of QIM systems involves the choice of practical quantizer ensembles for such systems, which we now explore. In the process, we obtain additional insights into the design, performance evaluation, and implementation of QIM embedding methods, particularly those of low-complexity. A convenient structure to consider is that of so-called dithered quantizers [29, 30], which have the property that the quantization cells and reconstruction points of any given quantizer in the ensemble are shifted versions of the quantization cells and reconstruction points of any other quantizer in the ensemble. In non-watermarking contexts, the shifts typically correspond to pseudorandom vectors called dither vectors. For information-embedding purposes, the dither vector can be modulated with the embedded signal, i.e., each possible embedded signal maps uniquely onto a different dither vector $\mathbf{d}(m)$. The host signal is quantized with the resulting dithered quantizer to form the composite signal. Specifically, we start with some base quantizer $\mathbf{q}(\cdot)$, and the embedding function is

$$\mathbf{s}(\mathbf{x}; m) = \mathbf{q}(\mathbf{x} + \mathbf{d}(m)) - \mathbf{d}(m).$$

We call this type of information embedding "dither modulation". We discuss several low-complexity realizations of such dither modulation methods in the sequel.

5.1 Coded Binary Dither Modulation with Uniform Scalar Quantization

Coded binary dither modulation with uniform, scalar quantization is one such realization.[9] We assume that $1/N \leq R_m \leq 1$. The dither vectors in a coded binary dither modulation system are constructed as follows:

i) The NR_m information bits $\{b_1, b_2, \ldots, b_{NR_m}\}$ representing the embedded message m are error correction coded using a rate-k_u/k_c code to obtain a coded bit sequence $\{z_1, z_2, \ldots, z_{N/L}\}$,

[9] By scalar quantization, we mean that the high dimensional base quantizer $\mathbf{q}(\cdot)$ is the Cartesian product of scalar quantizers.

where

$$L = \frac{1}{R_{\mathsf{m}}}(k_u/k_c). \tag{21}$$

(In the uncoded case, $\mathsf{z}_i = \mathsf{b}_i$ and $k_u/k_c = 1$.) We divide the host signal $\mathbf{x}$ into N/L non-overlapping blocks of length L and embed the ith coded bit z_i in the ith block, as described below.

ii) Two length-L dither sequences $d[k,0]$ and $d[k,1]$ and one length-L sequence of uniform, scalar quantizers with step sizes $\Delta_1, \ldots, \Delta_L$ are constructed with the constraint

$$d[k,1] = \begin{cases} d[k,0] + \Delta_k/2, & d[k,0] < 0 \\ d[k,0] - \Delta_k/2, & d[k,0] \geq 0 \end{cases}, \quad k = 1, \ldots, L,$$

This constraint ensures that the two corresponding L-dimensional dithered quantizers are the maximum possible distance from each other. For example, a pseudorandom sequence of $\pm\Delta_k/4$ and its negative satisfy this constraint. One could alternatively choose $d[k,0]$ pseudorandomly with a uniform distribution over $[-\Delta_k/2, \Delta_k/2]$.[10] Also, the two dither sequences need not be the same for each length-L block.

iii) The ith block of $\mathbf{x}$ is quantized with the dithered quantizer using the dither sequence $d[k, \mathsf{z}_i]$.

A detailed assessment of the complexity of this QIM realization is developed in [15, 27].

The minimum distance properties of coded binary dither modulation are readily deduced. In particular, any two distinct coded bit sequences differ in at least d_H places, where d_H is the minimum Hamming distance of the error correction code. For each of these d_H blocks, the reconstruction points of the corresponding quantizers are shifted relative to each other by $\pm\Delta_k/2$ in the kth dimension. Thus, the square of the minimum distance (13) over all N dimensions is

$$\begin{aligned} d^2_{\min} &= d_H \sum_{k=1}^{L} \left(\frac{\Delta_k}{2}\right)^2 \\ &= \left(d_H \frac{k_u}{k_c}\right) \frac{1}{4LR_{\mathsf{m}}} \sum_k \Delta_k^2 \\ &= \gamma_c \frac{1}{4LR_{\mathsf{m}}} \sum_k \Delta_k^2, \end{aligned} \tag{22}$$

[10] A uniform distribution for the dither sequence implies that the quantization error is statistically independent of the host signal and leads to fewer "false contours", both of which are generally desirable properties from a perceptual viewpoint [29].

where to obtain the second equality we have used (21), and where in the third line γ_c is the gain of the error correction code,

$$\gamma_c \stackrel{\Delta}{=} d_H(k_u/k_c). \tag{23}$$

In the high signal-to-distortion ratio (SDR) regime of primary interest for high-fidelity applications, the quantization cells are sufficiently small that the host signal can be modeled as uniformly distributed within each cell. In this case, the expected squared error distortion of a uniform, scalar quantizer with step size Δ_k is the familiar

$$\frac{1}{\Delta_k}\int_{-\Delta_k/2}^{\Delta_k/2} x^2\,dx = \frac{\Delta_k^2}{12}. \tag{24}$$

Thus, the overall average expected distortion (1) is

$$D_{\mathbf{s}} = \frac{1}{12L}\sum_k \Delta_k^2. \tag{25}$$

Combining (22) and (25) yields the "distortion-normalized" squared minimum distance,

$$d^2_{\text{norm}} \equiv \frac{d^2_{\min}}{D_{\mathbf{s}}} = \frac{3\gamma_c}{R_{\boldsymbol{m}}}, \tag{26}$$

a quantity that can be used to characterize the achievable performance of QIM realizations more generally, as we will develop.

5.2 Spread-transform Dither Modulation

One special class of coded binary dither modulation methods is what we refer to as spread-transform dither modulation (STDM). We now develop its properties and quantify its advantages over other forms of dither modulation, over additive spread-spectrum methods, and over spread LBM.

To introduce STDM, we begin by observing that the distortion-normalized squared minimum distance (26) of binary dither modulation with uniform scalar quantization does not depend on the sequence Δ_k, i.e., on the distribution of the distortion across samples within the length-L block. Thus, one is free to choose any distribution without sacrificing d^2_{norm}, so the Δ_k's can be chosen to optimize other characteristics of the embedding.

To understand this property, consider Figs. 4–6, each of which show the reconstruction points of two quantizers for embedding one bit in a block of two samples. For each of the three systems, the minimum distance—$\Delta/\sqrt{2}$—and the average squared error distortion—$\Delta^2/12$ per sample—are

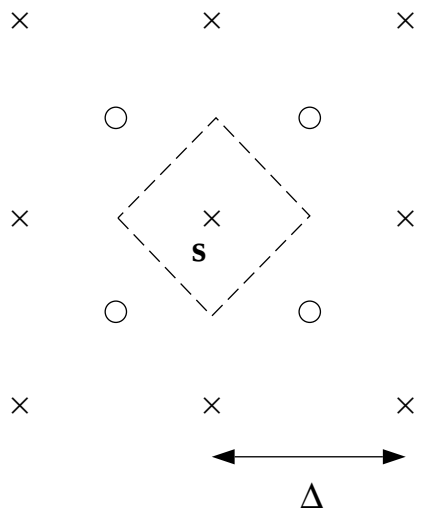

Figure 4: Dither modulation with uniform quantization step sizes.

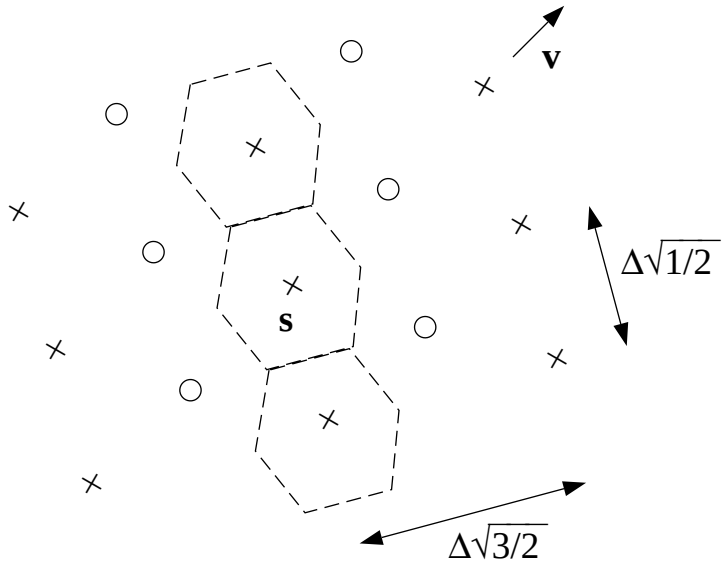

Figure 5: Transform dither modulation with non-uniform quantization step sizes.

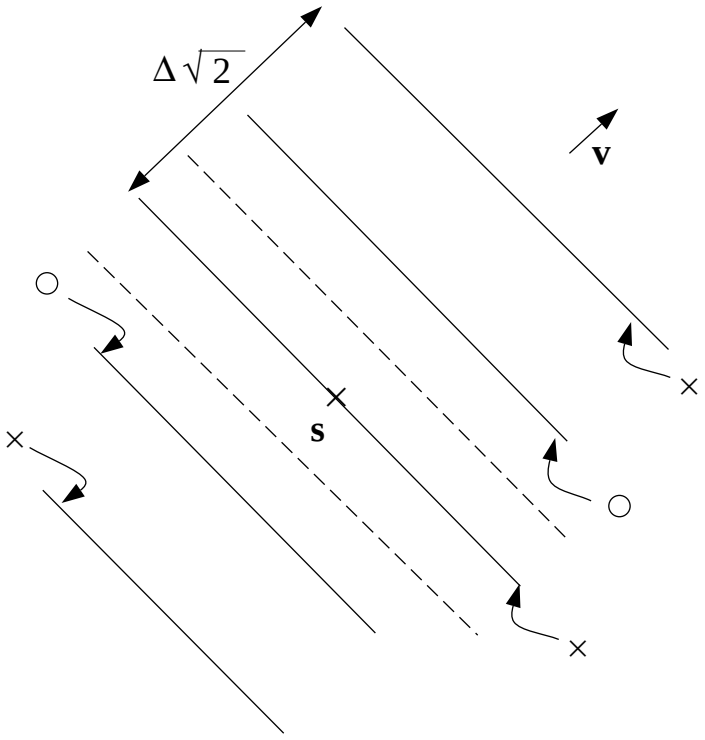

Figure 6: Transform dither modulation with quantization of only a single transform component. The quantization step size for the component of the host signal orthogonal to **v** is zero.

identical. Thus, the robustness against bounded perturbations is the same in each case. However, the quantization differs in each case. In Fig. 4, where scalar quantization is applied to each sample separately, the quantization step sizes are the same for both samples. In Figs. 5 and 6, the samples are first pretransformed and the resulting coefficients quantized unevenly. In particular, a unitary transform (coordinate rotation) is applied to the pair of samples before quantization; the first transform coefficient is the component of the host signal in the direction of $\mathbf{v}$ depicted. In Fig. 5, the step size for quantizing the first transform coefficient is larger than that used to quantize the second transform coefficient, which lies in the direction orthogonal to $\mathbf{v}$. Finally, in the extreme case of Fig. 6, the step size for the first coefficient is larger still, and that for the second coefficient is zero, i.e., all embedding occurs in the first coefficient. In this case, the reconstruction points become reconstruction lines, so to embed a 0-bit, the host signal is quantized to the nearest point on a line labeled with a $\times$. To embed a 1-bit, the host signal is quantized to the nearest point on a line labeled with a $\circ$.

While the three systems corresponding to Figs. 4–6 have the same minimum distance, the *number* of perturbation vectors of minimum length that cause decoding errors is higher for the case of Fig. 4 than for the case of Fig. 6. (For intermediate cases such as the one shown in Fig. 5, where quantization step sizes in different dimensions are different but non-zero, the number of perturbation vectors of minimum length that cause decoding errors is the same as in Fig. 4, but these vectors are not orthogonal.) Thus, for probabilistic channels such as additive noise channels, the *probability* of error is generally different in each case. For example, suppose a 0-bit is embedded and the composite signal is the $\times$ point labeled with $\mathbf{s}$ in Figs. 4 and 6. If the channel output lies in the decision region defined by the dashed box in Fig. 4 and defined by the two dashed lines in Fig. 6, then the decoder will correctly determine that a 0-bit was embedded. If the perturbation vector places the channel output outside the decision region, however, the decoder will make an error with very high probability. (There is some possibility that the channel output is outside the decision region but is still closer to a $\times$ point other than $\mathbf{s}$ than to the closest $\circ$. These events, however, are very unlikely for many perturbation probability distributions that are of practical interest.) Since the decision region of Fig. 6 contains the decision region of Fig. 4, it follows that the probability of a correct decision in the case of non-uniform quantization step sizes is higher.

The unitary transform in the case of Fig. 6 not only facilitates a comparison of Figs. 4 and 6, but also serves to spread any embedding-induced distortion over frequency and time/space when a peak distortion constraint is imposed, for example. Although, the distortion is concentrated in only one transform coefficient, if the energy of $\mathbf{v}$ is spread over space/time and frequency—for example,

if $\mathbf{v}$ is chosen pseudorandomly—then the distortion will also be spread.

As we will see in subsequent sections of this paper, dither modulation methods have considerable performance advantages over previously proposed additive spread-spectrum and spread LBM methods in a variety of contexts. However, much effort has already been invested in optimizing both additive spread-spectrum and spread LBM systems, for example, by exploiting perceptual properties of the human visual and auditory systems or designing receiver front-ends to mitigate effects of geometric and other distortions. An additional advantage of STDM specifically over other forms of dither modulation is that one can easily convert existing additive spread-spectrum and spread LBM systems into STDM systems while retaining the other optimized components of the system. In particular, it suffices to replace the addition step of additive spread-spectrum, i.e., (5), or the quantize-and-replace step of spread LBM, i.e., (11), with the dithered quantization step of STDM, i.e.,

$$\tilde{s} = \mathbf{s}^T \mathbf{v} = q(\tilde{x} + d(m)) - d(m). \tag{27}$$

5.2.1 SNR advantage of STDM

In this section, we quantify the performance gain of STDM over additive spread-spectrum and spread LBM from an SNR perspective that applies to a broad range of contexts. We focus our analysis on the representative case of embedding one bit in a length-L block $\mathbf{x}$ using a unit-length spreading vector $\mathbf{v}$. Because, as (5), (11), and (27) reflect, in each case the embedding occurs entirely in the projection of $\mathbf{x}$ onto $\mathbf{v}$, a one-dimensional problem results. In addition, because all of the embedding-induced distortion occurs only in the direction of $\mathbf{v}$, the distortion in each case also has the same temporal/spatial distribution and frequency distribution. Thus, one would expect that any perceptual effects due to time/space masking or frequency masking are the same in each case. Therefore, squared error distortion and SNR-type measures are more meaningful measure of distortion when comparing these embedding methods than one might expect in other more general contexts where squared error distortion may fail to capture certain perceptual effects.

SNR advantage of STDM over additive spread-spectrum Considering the case of additive spread-spectrum first, since $a(m) = \pm\sqrt{LD_{\mathbf{s}}}$ in (5), we have

$$|a(1) - a(2)|^2 = 4LD_{\mathbf{s}}. \tag{28}$$

For STDM (27),

$$\min_{(\tilde{x}_1,\tilde{x}_2)} |\tilde{s}(\tilde{x}_1,1) - \tilde{s}(\tilde{x}_2,2)|^2 = \Delta^2/4 = 3LD_{\mathbf{s}}, \tag{29}$$

where $\Delta = \sqrt{12LD_{\mathbf{s}}}$ so that the expected distortion in both cases is the same, and where we have used the fact that $d(1)$ and $d(2)$ are chosen such that $|d(1) - d(2)| = \Delta/2$.

The decoder in both cases makes a decision based on $\tilde{y}$, the projection of the channel output $\mathbf{y}$ onto $\mathbf{v}$. In the case of additive spread-spectrum, $\tilde{y} = a(m) + \tilde{x} + \tilde{n}$, while in the case of STDM, $\tilde{y} = \tilde{s}(\tilde{x}, m) + \tilde{n}$, where $\tilde{n}$ is the projection of the perturbation vector $\mathbf{n}$ onto $\mathbf{v}$. We let $P(\cdot)$ be some measure of energy. For example, $P(x) = x^2$ in the case of a deterministic variable x, or $P(\mathsf{x}) = \mathrm{var}\mathsf{x}$ when x is random. The energy of the interference or "noise" is $P(\tilde{x} + \tilde{n})$ for additive spread-spectrum, but only $P(\tilde{n})$ for STDM, i.e., the host signal interference for STDM is zero. Thus, the SNR at the decision device is

$$\mathrm{SNR}_{\mathrm{SS}} = \frac{4LD_{\mathbf{s}}}{P(\tilde{x} + \tilde{n})}$$

for additive spread-spectrum and

$$\mathrm{SNR}_{\mathrm{STDM}} = \frac{3LD_{\mathbf{s}}}{P(\tilde{n})}$$

for STDM, where the "signal" energies $P(a(1) - a(2))$ and $P\left(\min_{(\tilde{x}_1,\tilde{x}_2)} |\tilde{s}(\tilde{x}_1,1) - \tilde{s}(\tilde{x}_2,2)|\right)$ are given by (28) and (29). Thus, the advantage of STDM over additive spread-spectrum is

$$\frac{\mathrm{SNR}_{\mathrm{STDM}}}{\mathrm{SNR}_{\mathrm{SS}}} = \frac{3}{4}\frac{P(\tilde{x} + \tilde{n})}{P(\tilde{n})}, \tag{30}$$

which is typically very large since the channel perturbations $\tilde{n}$ are usually much smaller than the host signal $\tilde{x}$ if the channel output $\tilde{y}$ is to be of reasonable quality. For example, if the host signal-to-channel noise ratio is 30 dB and $\tilde{x}$ and $\tilde{n}$ are uncorrelated, then the SNR advantage (30) of STDM over additive spread-spectrum is 28.8 dB.[11]

SNR advantage of STDM over spread LBM Spread-transform dither modulation methods also have an SNR advantage over spread LBM methods. As we show in App. A, the distortion-normalized squared minimum distance (26) of LBM is $7/4 \approx 2.43$ dB worse than that of dither

[11] Note that while the high SDR approximation (30) predicts that STDM is worse than additive spread-spectrum by a factor of $4/3 = 1.25$ dB when $\tilde{x} \approx 0$ (as would be the case, for example, if the host signal $\mathbf{x}$ had very little energy in the direction of $\mathbf{v}$), in fact if one chooses $d(m) = \pm\Delta/4$, then it is straightforward to verify that STDM performs as well as additive spread-spectrum in this low SDR regime.

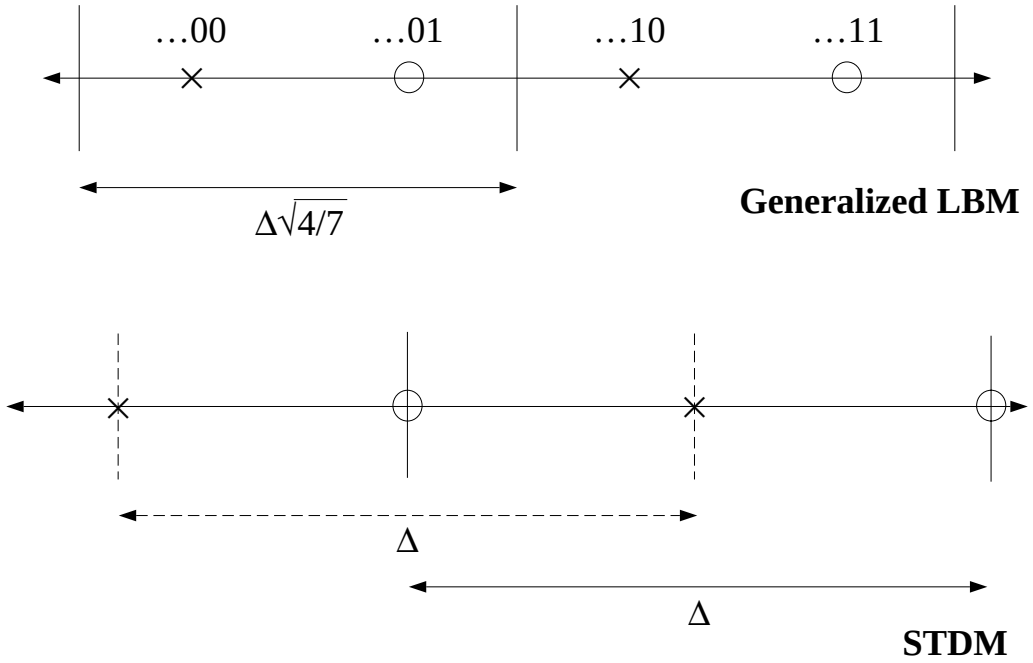

Figure 7: Spread-transform dither modulation vs. spread LBM. The embedding interval boundaries of spread LBM, which are shown with solid lines, are the same for both $\times$ points and $\circ$ points. In contrast, in the case of STDM, the $\times$-point embedding intervals, shown by solid lines, differ from the $\circ$-point embedding intervals, shown by dashed lines. An SNR advantage of $7/4 = 2.43$ dB for STDM results.

modulation in the case of coded binary embedding with uniform, scalar quantization. Thus, for a fixed rate and embedding-induced distortion, the squared-minimum distance, and hence the SNR at the decision device, for spread LBM will be 2.43 dB worse than that of STDM, i.e.,[12]

$$\frac{\text{SNR}_{\text{STDM}}}{\text{SNR}_{\text{SLBM}}} = \frac{7}{4} \approx 2.43 \text{ dB} \tag{31}$$

This SNR advantage is illustrated in Fig. 7, where the quantizer reconstruction points and embedding intervals for both spread LBM and STDM are shown, with the same embedding-induced squared error distortion for both cases.

The preceding analysis establishes some important advantages of QIM methods over common information embedding methods. In fact, it turns out that QIM methods are asymptotically optimal in many key scenarios of interest. To develop these results, we next examine information embedding within an information-theoretic framework.

[12]App. A also shows that for M-ary embedding the SNR gain grows to 2 (3 dB) as $M \to \infty$.

6 Information Theoretic Optimality of QIM

This section explores the best possible rate-distortion-robustness performance that one could hope to achieve with any information embedding system. Our analysis leads to insights about some properties and characteristics of good information embedding methods, i.e., methods that achieve performance close to the information-theoretic limits. In particular, a canonical "hidden QIM" structure emerges for information embedding that consists of (1) preprocessing of the host signal, (2) QIM embedding, and (3) postprocessing of the quantized host signal to form the composite signal. One incurs no loss of optimality by restricting one's attention to this simple structure. We also derive sufficient conditions under which only distortion compensation postprocessing is required. As we develop in Secs. 7 and 8, these conditions are satisfied in several important cases of practical interest.

6.1 Communication over Channels with Side Information

The super-channel model of Sec. 2.2 and Fig. 2 facilitates our analysis, i.e., we view information embedding as the transmission of a host-dependent distortion signal $\mathbf{e}$ over a super-channel with side information or state $\mathbf{x}$ that is known at the encoder. In this section we also restrict our attention to a squared error distortion constraint

$$\frac{1}{N}\sum_{i=1}^{N} e_i^2 \leq D_{\mathbf{s}},$$

and a memoryless channel with known probability density function (pdf)

$$p_{\mathbf{y}|\mathbf{s}}(\mathbf{y}|\mathbf{s}) = \prod_{i=1}^{N} p_{y|s}(y_i|s_i),$$

where y_i and s_i are the ith components of $\mathbf{y}$ and $\mathbf{s}$, respectively.[13] Then, the super-channel is also memoryless and has probability law

$$p_{\mathbf{y}|\mathbf{e},\mathbf{x}}(\mathbf{y}|\mathbf{e},\mathbf{x}) = p_{\mathbf{y}|\mathbf{s}}(\mathbf{y}|\mathbf{x}+\mathbf{e}) = \prod_{i=1}^{N} p_{y|s}(y_i|x_i+e_i) = \prod_{i=1}^{N} p_{y|e,x}(y_i|e_i,x_i).$$

[13] Extension of results in this section to the case where the channel is only blockwise memoryless is straightforward by letting y_i and s_i be the ith blocks, rather than ith scalar components, of $\mathbf{y}$ and $\mathbf{s}$. In this case, information rates are measured in bits per block, rather than bits per sample.

The capacity [31] of this super-channel is the reliable information-embedding rate R_m that is asymptotically achievable with long signal lengths N.

In non-watermarking contexts Gel'fand and Pinsker [32] and Heegard and El Gamal [33] have determined the capacity of such a channel in the case of a random state vector $\mathbf{x}$ with independent and identically distributed (iid) components when the encoder sees the entire state vector before choosing the channel input $\mathbf{e}$. In this case the capacity is

$$C = \max_{p_{u,e|x}(u,e|x)} I(u;y) - I(u;x), \tag{32}$$

where $I(\cdot;\cdot)$ denotes mutual information and u is an auxiliary random variable. Since $p_{u,e|x}(u,e|x) = p_{u|x}(u|x)p_{e|u,x}(e|u,x)$, we can think of u in (32) as being generated from x, and, in turn, e from u and x. While the mapping from x to u is in general probabilistic, from convexity properties of mutual information, one can deduce that the maximizing distribution in (32) always has the property that e is a deterministic function of (u,x) [32].

In the case of watermarking, the maximization (32) is subject to a distortion constraint $E[e^2] \leq D_{\mathbf{s}}$. A formal proof of the extension of (32) to include this constraint is developed in [20]. Others [18, 19, 16] are working on extending or have extended these results to the case where the channel law $p_{y|s}(y|s)$ is not fixed but rather is chosen by an attacker subject to a distortion constraint. A related information-theoretic formulation can be found in [14].

As we shall see in the next section, one way to interpret (32) is that $I(u;y)$ is the total number of bits per host signal sample that can be transmitted through the channel, and $I(u;x)$ is the number of bits per sample that are allocated to the host signal x. The difference between the two is the number of bits per host signal sample that can be allocated to the embedded information m.

6.1.1 Hidden QIM

As we show in this section, one can achieve the capacity (32) by a type of "hidden" QIM, i.e., QIM that occurs in a domain represented by the auxiliary random variable u. One moves into and out of this domain with pre- and post-quantization processing.

To develop this optimality of hidden QIM, we begin by adding an interpretation in terms of quantization (source coding) to the proof of the achievability of capacity by Gel'fand and Pinsker [32], the result of which is summarized as follows. Fig. 8 shows an ensemble of 2^{NR_m} quantizers, where $R_m = I(u;y) - I(u;x) - 2\epsilon$, where each source codeword (quantizer reconstruction vector) $\mathbf{u}$ is randomly drawn from the iid distribution $p_u(u)$, which is the marginal distribution correspond-

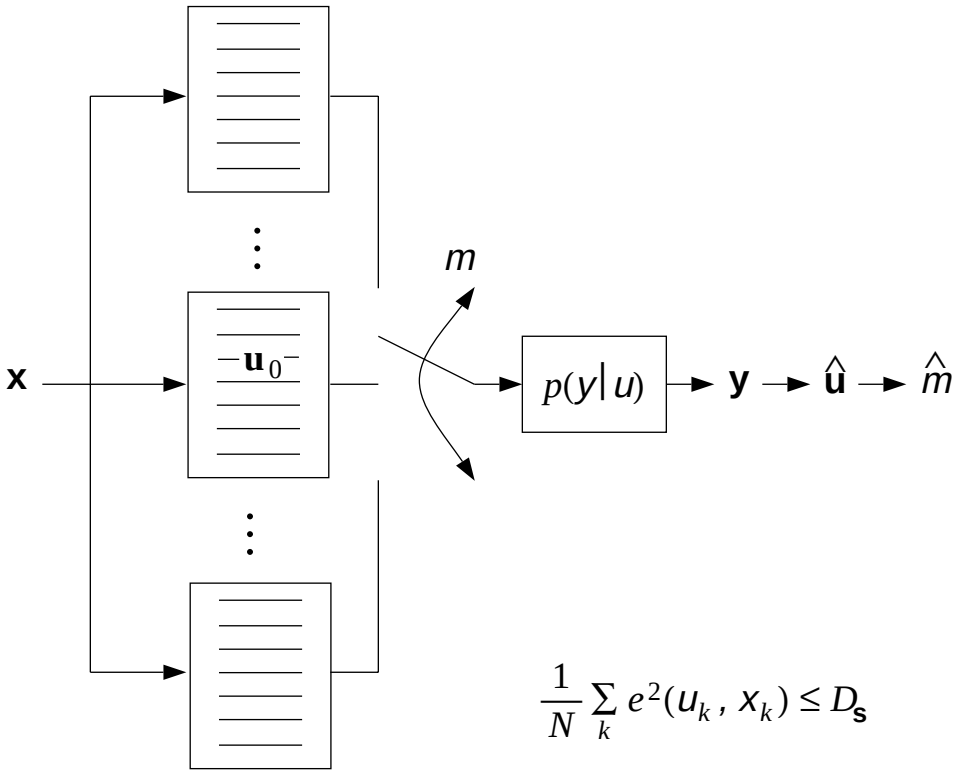

Figure 8: Capacity-achieving "hidden QIM". One embeds by choosing a codeword $\mathbf{u}_0$ that is jointly distortion-typical with $\mathbf{x}$ from the mth quantizer's codebook. The distortion function is $e^2(u, x)$. The decoder finds a codeword that is jointly typical with $\mathbf{y}$. If this codeword is in the ith subset, then $\hat{m} = i$.

ing to the host signal distribution $p_{\mathsf{x}}(x)$ and the maximizing conditional distribution $p_{\mathsf{u},\mathsf{e}|\mathsf{x}}(u,e|x)$ from (32). Although the source codebooks are therefore random, both the encoder and decoder, of course, know the codebooks. Each codebook contains $2^{N[I(\mathsf{u};\mathsf{x})+\epsilon]}$ codewords so there are $2^{N[I(\mathsf{u};\mathsf{y})-\epsilon]}$ codewords total.

QIM embedding in this $\mathbf{u}$-domain corresponds to finding a vector $\mathbf{u}_0$ in the mth quantizer's codebook that is jointly distortion-typical with $\mathbf{x}$ and generating

$$\mathbf{e}(\mathbf{u}_0, \mathbf{x}) = [e(u_{0,1}, x_1) \ \cdots \ e(u_{0,N}, x_N)]^T.$$

By distortion-typical, we mean that $\mathbf{u}_0$ and $\mathbf{x}$ are jointly typical and $\|\mathbf{e}(\mathbf{u}_0, \mathbf{x})\|^2 \leq N(D_{\mathbf{s}} + \epsilon)$, i.e., the function $e^2(u,x)$ is the distortion function in the $\mathbf{u}$-domain. Since the mth quantizer's codebook contains more than $2^{NI(\mathsf{u};\mathsf{x})}$ codewords, the probability that there is no $\mathbf{u}_0$ that is jointly distortion-typical with $\mathbf{x}$ is small.[14] Thus, the selection of a codeword from the mth quantizer is the quantization part of QIM, and the generation of $\mathbf{e}$, and therefore $\mathbf{s} = \mathbf{x} + \mathbf{e}$, from the codeword $\mathbf{u}_0$ and $\mathbf{x}$ is the post-quantization processing.

The decoder finds a $\mathbf{u}$ that is jointly typical with the channel output $\mathbf{y}$ and declares $\hat{m} = i$ if this $\mathbf{u}$ is in the ith quantizer's codebook. Because the total number of codewords $\mathbf{u}$ is less than $2^{NI(\mathsf{u};\mathsf{y})}$, the probability that a $\mathbf{u}$ other than $\mathbf{u}_0$ is jointly typical with $\mathbf{y}$ is small. Also, the probability that $\mathbf{y}$ is jointly typical with $\mathbf{u}_0$ is close to 1.[15] Thus, the probability of error $\Pr[\hat{m} \neq m]$ is small, and we can indeed achieve the capacity (32) with QIM in the $\mathbf{u}$-domain.

The remaining challenge, therefore, is to determine the right preprocessing and postprocessing given a particular channel (attack) $p_{\mathsf{y}|\mathsf{s}}(y|s)$. As mentioned above, for a number of important cases, it turns out that the only processing required is post-quantization distortion compensation. We discuss these cases in the next section.

6.1.2 Optimality of distortion-compensated QIM

When distortion-compensated QIM (DC-QIM) as introduced in Sec. 4.1 is viewed as an instance of hidden QIM, we obtain that $\mathbf{u}$ is a quantized version of $\alpha\mathbf{x}$. We show in this section that suitably coded versions DC-QIM can achieve capacity whenever the maximizing distribution $p_{\mathsf{u},\mathsf{e}|\mathsf{x}}(u,e|x)$ in (32) is of a form such that the postprocessing is linear, i.e., when, without loss of generality, e

[14] This principle is, of course, one of the main ideas behind the rate-distortion theorem [31, Ch. 13].
[15] These principles are, of course, two of the main ideas behind the classical channel coding theorem [31, Ch. 8].

is generated according to

$$e = u - \alpha x. \tag{33}$$

To see that DC-QIM can achieve capacity when the maximizing pdf in (32) satisfies (33), we show that one can construct an ensemble of random DC-QIM codebooks that satisfy (33). First, we observe that quantizing $\mathbf{x}$ is equivalent to quantizing $\alpha\mathbf{x}$ with a scaled version of the quantizer and scaling back the result, i.e.,

$$\mathbf{q}(\mathbf{x}; m, \Delta/\alpha) = \frac{1}{\alpha}\mathbf{q}(\alpha\mathbf{x}; m, \Delta). \tag{34}$$

where $\mathbf{q}(\cdot;\cdot,\cdot)$ is as defined following (19). Then, rearranging terms in the DC-QIM embedding function (19) and substituting (34) into the result, we obtain

$$\begin{aligned} \mathbf{s}(\mathbf{x}, m) &= \mathbf{q}(\mathbf{x}; m, \Delta/\alpha) + (1-\alpha)[\mathbf{x} - \mathbf{q}(\mathbf{x}; m, \Delta/\alpha)] \\ &= \alpha\mathbf{q}(\mathbf{x}; m, \Delta/\alpha) + (1-\alpha)\mathbf{x} \\ &= \mathbf{q}(\alpha\mathbf{x}; m, \Delta) + (1-\alpha)\mathbf{x}. \end{aligned} \tag{35}$$

We construct our random DC-QIM codebooks by choosing the codewords of $\mathbf{q}(\cdot; m, \Delta)$ from the iid distribution $p_u(u)$, the one implied by the maximizing pdf in (32) together with the host pdf $p_x(x)$. (Equivalently, we choose the codewords of $\mathbf{q}(\cdot; m, \Delta/\alpha)$ in (19) from the distribution of u/α, i.e., the iid distribution $\alpha p_u(\alpha u)$.) Our quantizers $\mathbf{q}(\cdot; m, \Delta)$ choose a codeword $\mathbf{u}_0$ that is jointly distortion-typical with $\alpha\mathbf{x}$. The decoder looks for a codeword in all of the codebooks that is jointly typical with the channel output. Then, following the achievability argument of Sec. 6.1.1, we can achieve a rate $I(u;y) - I(u;x)$. From (35), we see that

$$\mathbf{s}(\mathbf{x}, m) = \mathbf{x} + [\mathbf{q}(\alpha\mathbf{x}; m, \Delta) - \alpha\mathbf{x}] = \mathbf{x} + (\mathbf{u}_0 - \alpha\mathbf{x}).$$

Since $\mathbf{s}(\mathbf{x}, m) = \mathbf{x} + \mathbf{e}$, we see that $\mathbf{e} = \mathbf{u}_0 - \alpha\mathbf{x}$. Thus, if the maximizing distribution in (32) satisfies (33), our DC-QIM codebooks can also have this distribution and, hence, achieve capacity (32).

As a final comment, it is worth emphasizing that QIM systems are optimal in other important scenarios as well. As one example, in the noise-free case ($\mathbf{y} = \mathbf{s}$), which arises, for example, when a discrete-valued composite signal is transmitted over a digital channel with no errors, QIM is

optimal even without distortion compensation, and achieves capacity [27]

$$C_{\text{noise-free}} = \max_{p_{\mathsf{y}|\mathsf{x}}(y|x)} H(\mathsf{y}|\mathsf{x}). \tag{36}$$

As a second example, and as shown in [27], QIM is optimal even when the host signal is also available at the decoder achieving the capacity

$$C_{\text{known}} = \max_{p_{\mathsf{e}|\mathsf{x}}(e|x)} I(\mathsf{e};\mathsf{y}|\mathsf{x}) \tag{37}$$

determined by Heegard and El Gamal [33].

We next examine some key scenarios when the optimality condition (33) is met.

7 Gaussian Channels

In this section we examine the ultimate performance limits of information embedding methods when both the host signal is white and Gaussian, the channel is an additive white Gaussian noise (AWGN) channel, and the host and channel noise are independent of one another. Extensions to colored host and/or colored channel cases are developed in [15] [27]. Our main result of the section is that DC-QIM is optimal for this class of channels, and that in addition the optimum distortion compensation parameter α is also given by (20), which maximized SNR in uncoded DC-QIM systems.

In general, the embedding strategies optimized for Gaussian channel models can be expected to be good designs for a variety of applications in which one primarily requires robustness against unintentional attacks.[16] And while Gaussian host models are not always accurate, the better the host-signal interference rejection properties of an information embedding system, the smaller the role we might expect the host signal model to play in determining the ultimate performance of such systems.

7.1 Capacities and the optimality of DC-QIM

Specializing the formulation of Sec. 6.1 to the Gaussian scenario of interest, with the zero-mean, variance-σ_x^2 variables x_i denoting elements of the N-dimensional host signal vector $\mathbf{x}$, and similarly

[16] Indeed, these models can even apply to optimal, i.e., rate-distortion achieving [31], lossy compression of a Gaussian source, as discussed in [27].

the zero-mean, variance-σ_n^2 variables n_i denoting elements of the corresponding noise vector $\mathbf{n}$, the distortion constraint can be expressed as

$$\frac{1}{N}\sum_{i=1}^{N} e_i^2 \leq D_{\mathbf{s}},$$

with the corresponding constraint on $p_{u,e|x}(u,e|x)$ in (32) being $E[e^2] \leq D_{\mathbf{s}}$. We see that squared error distortion-constrained, Gaussian information embedding is equivalent to power-constrained communication over a Gaussian channel with Gaussian side information known at the encoder, a case for which Costa [34] has determined the capacity to be, expressed in terms of the (embedding induced) distortion-to-noise ratio (DNR),

$$C_{\text{Gauss}} = \frac{1}{2}\log_2(1+\text{DNR}), \qquad \text{DNR} = \frac{D_{\mathbf{s}}}{\sigma_n^2}. \tag{38}$$

Remarkably, the capacity is independent of the signal variance σ_x^2 and in fact, as we'll discuss later in this section, is the same as in the case when the host signal $\mathbf{x}$ is known at the decoder. Note that this implies that an infinite energy host signal causes no decrease in capacity in this Gaussian case, i.e., good information embedding systems can *completely* reject host-signal interference in the Gaussian case.

Based on our earlier results, to establish the optimality of DC-QIM for this channel, it suffices to verify that (33) is satisfied. This follows from the proof [34] of (38). In particular, as shown in [34], the pdf that maximizes (32) is indeed one implied by (33), for some parameter α, where u is chosen as a function of x so that $e \sim \mathcal{N}(0, D_{\mathbf{s}})$ and so that the pair e and x are independent. To see this, note that for a fixed value of α, an achievable rate $I(u;y) - I(u;x)$ is [34]

$$R(\alpha) = \frac{1}{2}\log_2\left(\frac{D_{\mathbf{s}}(D_{\mathbf{s}}+\sigma_x^2+\sigma_n^2)}{D_{\mathbf{s}}\sigma_x^2(1-\alpha)^2+\sigma_n^2(D_{\mathbf{s}}+\alpha^2\sigma_x^2)}\right),$$

which can also be written in terms of the DNR and the host SNR ($\text{SNR}_x = \sigma_x^2/\sigma_n^2$),

$$R(\alpha) = \frac{1}{2}\log_2\left(\frac{\text{DNR}(1+\text{DNR}+\text{SNR}_x)}{\text{DNR}\,\text{SNR}_x(1-\alpha)^2+(\text{DNR}+\alpha^2\text{SNR}_x)}\right). \tag{39}$$

This rate is maximized by setting [*cf.* (20)]

$$\alpha_{\text{cap}} = \frac{\text{DNR}}{\text{DNR}+1} \tag{40}$$

from which we conclude that the rate (38) is achievable. To establish that (38) is also the maximum achievable rate, it suffices to show that it is the capacity when x is known at the decoder, since one obviously cannot do better in the host-blind case.

To develop the known-host capacity, first recall that the capacity is given by (37). Again, the maximization is subject to a distortion constraint, which in the case of white noise is $E[e^2] \leq D_{\mathbf{s}}$. Because subtracting a known constant from y does not change mutual information, we can equivalently write

$$C = \max_{p_{e|x}(e|x)} I(e; y - x|x).$$

Noting that $y-x = e+n$, we immediately conclude that in the case of an AWGN channel the known-host capacity is indeed given by (38), where the maximizing distribution $p_{e|x}(e|x)$ is a zero-mean Gaussian distribution with variance $D_{\mathbf{s}}$.

In the known-host case, additive spread-spectrum is optimal, and optimal additive spread-spectrum systems superimpose zero-mean iid Gaussian sequences with variance $D_{\mathbf{s}}$ onto the host signal. However, it is important to note that QIM is also optimal in this case as well—as discussed in [15], quantizers of optimal QIM systems have reconstruction sequences s_i chosen iid from a zero-mean Gaussian distribution with variance $\sigma_x^2 + D_{\mathbf{s}}$. Hence, yet another attractive property of QIM methods is that they are optimal in more general Gaussian broadcast scenarios where some intended recipients of the embedded information know the host signal and some do not.

As a final comment, several of the methods we have discussed can be optimal in the small host signal interference scenario ($x \to 0$). In fact, the capacity (38) is rather immediate in this scenario: Fig. 2 reduces to the classical communication problem considered in, e.g., [31] since $s \to e$, so that the capacity is the usual mutual information between $e = s$ and y maximized over all $p_e(\cdot)$ such that $E[e^2] \leq D_{\mathbf{s}}$. In the additive white Gaussian noise channel case, specifically, (38) results. Examining (39) in the associated regime ($\mathrm{SNR}_x \to 0$), we see that distortion-compensated QIM with any α, including $\alpha = 1$ (regular QIM), is optimal in this small host interference scenario. As one might expect, additive spread spectrum systems can be capacity-achieving in this limit as well, which we will see more explicitly in Sec. 7.3.4.

7.2 Capacities for Hybrid Transmission

In this section, we consider scenarios corresponding to applications in which information embedding is part of a hybrid transmission scheme. We investigate two classes of such schemes: analog-digital digital-digital transmission. In the former class, the host is an analog signal, as arises in, for

example, the digital audio broadcasting application. In the latter class, the host signal is itself a digital signal, which has implications for broadcast transmission and related applications [31, Ch. 14].

In both cases, one is generally most concerned with the quality of the *received* signals, i.e., the channel output, rather than the channel input (composite signal).

7.2.1 Analog host signals

In this section, we determine how reliable embedding at a given rate impacts the quality with which an analog host signal is received and can be decoded with its conventional receiver from a noisy channel.

In general, the effect of the embedding is to create an additional noise source DNR times as strong as the channel noise, and therefore, the received signal quality drops by a factor of (1+DNR) or

$$10 \log_{10}(1 + \mathrm{DNR}) \text{ dB}. \tag{41}$$

For example, in the scenario analyzed in Sec. 7.1, optimum DC-QIM results in an embedding-induced distortion that looks like white noise with variance $D_{\mathbf{s}}$. With no embedding, one would have had a received host signal-to-noise ratio of $\mathrm{SNR}_x = \sigma_x^2/\sigma_n^2$. Due to the additional interference from the embedding-induced distortion, however, the received host SNR drops to

$$\frac{\sigma_x^2}{D_{\mathbf{s}} + \sigma_n^2} = \frac{\mathrm{SNR}_x}{1 + \mathrm{DNR}},$$

a drop of $1 + \mathrm{DNR}$.

Since the capacity in bits per dimension (bits per host signal sample) is given by (38), and there are two independent host signal samples per second for every Hertz of host signal bandwidth [28], the capacity in bits per second per Hertz is

$$C = \log_2(1 + \mathrm{DNR}) \text{ b/s/Hz}. \tag{42}$$

Taking the ratio between (42) and (41), we see that the "value" in embedded rate of each dB drop in received host signal quality is

$$C = \frac{\log_2(1 + \mathrm{DNR})}{10 \log_{10}(1 + \mathrm{DNR})} = \frac{1}{10} \log_2 10 \approx 0.3322 \text{ b/s/Hz/dB} \tag{43}$$

Thus, the available embedded digital rate in bits per second depends only on the bandwidth of the host signal and the tolerable degradation in received host signal quality, and is approximately 1/3 b/s for every Hz of bandwidth and every dB drop in received host SNR. It is worth noting that, as developed in [15, 27], these results carry over to the case of colored host and/or colored channel cases as well.

Additional insights into the performance limits of such systems when the digital signal is specifically information for refining the analog signal, as arises in applications involving the upgrading of analog infrastructure, are developed in [20].

7.2.2 Coded digital host signals

When the host signal is a coded digital signal, an alternative measure of the received host signal quality is the capacity of the corresponding host digital channel. For example, in the case of white noise and a white host signal,[17] if there were no embedding, the capacity corresponding to a host digital signal power of σ_x^2 and a noise variance of σ_n^2 is

$$R_0 = \frac{1}{2}\log_2(1 + \mathrm{SNR}_x).$$

Embedding an additional digital signal within the host digital signal drops the host digital capacity to

$$R_1 = \frac{1}{2}\log_2\left(1 + \frac{\mathrm{SNR}_x}{1 + \mathrm{DNR}}\right)$$

due to the drop in received host SNR of $1 + \mathrm{DNR}$. Unlike in the case of an analog host signal, if one must actually lower the rate of the coded host digital signal as a result of the embedding, then one may have to redesign both the digital encoder that generates this coded digital host signal and the corresponding decoder. Thus, depending on the designed noise margin of the original digital host signal, backwards-compatibility may or may not be possible.

However, even when digital-digital transmission cannot be backwards compatible, using information embedding for simultaneous transmission of two digital signals is potentially attractive from the point of view of complexity and privacy. In particular, the decoder for the host signal need not decode (nor know how to decode) the embedded signal, and vice-versa.

As discussed further in [27], this is qualitatively different behavior from the superposition coding

[17] As is well known [31], white Gaussian coded signals are capacity-achieving for transmission over additive white Gaussian noise channels, so this is a good model for the host signal in this case.

and successive cancellation decoding one might otherwise use for simultaneous transmission of two digital signals, where one of the receivers needs to decode both messages to receive its own.

Interestingly, the information embedding approach is equally efficient. To see this, we note that the embedded digital channel rate is given by (38),

$$R_2 = \frac{1}{2}\log_2(1 + \mathrm{DNR})$$

so that the combined rate of the two channels is

$$R_1 + R_2 = \frac{1}{2}\log_2(1 + \mathrm{DNR} + \mathrm{SNR}_x).$$

Since the associated expended power is $D_{\mathbf{s}} + \sigma_x^2$, we conclude that this digital-over-digital transmission strategy is indeed efficient: the combined rate $R_1 + R_2$ is as large as the achievable rate using a *single* digital signal with this same total power.

7.3 Gaps to Capacity

In Sec. 7.1, we saw that DC-QIM is a capacity-achieving strategy. In this section, for comparison we evaluate the degree to which specific strategies such as regular QIM (i.e., without distortion compensation), coded additive spread-spectrum, uncoded STDM, and uncoded generalized LBM can each approach capacity—and hence the performance of DC-QIM—when suitably optimized. We quantify the performance of these systems in terms of the additional DNR required to achieve the same rate as a capacity-achieving system.

7.3.1 Regular QIM gap to capacity

As we now show, the performance of the best QIM methods without distortion compensation can approach the Gaussian capacity at high rates and is within 4.3 dB of capacity at low rates, indicating that the QIM class is large enough to include very good embedding functions and decoders.

To develop a lower bound on the achievable rate of QIM without distortion-compensation, we begin by specializing (39) to the case $\alpha = 1$, resulting in

$$R_{\mathrm{QIM}} \geq \frac{1}{2}\log_2\left(\mathrm{DNR}\frac{1 + \mathrm{DNR} + \mathrm{SNR}_x}{\mathrm{DNR} + \mathrm{SNR}_x}\right), \tag{44}$$

where to achieve this bound we choose reconstruction points from the pdf implied by (33).[18] The righthand side of (44) is generally not the capacity of QIM, however—i.e., QIM systems can achieve a rate greater than the lower bound (44). Indeed, the righthand side of (44) actually approaches $-\infty$ in the limit of low DNR.

A tighter lower bound is obtained by developing a different lower bound on the capacity of a particular subclass of QIM methods we refer to as "spread-transform QIM." In spread-transform QIM, which is a generalization of STDM as developed in Sec. 5.2, the host signal vector $\mathbf{x} = [x_1 \cdots x_N]^T$ is projected onto $N/L_{\rm ST}$ orthonormal vectors $\mathbf{v}_1, \ldots, \mathbf{v}_{N/L_{\rm ST}} \in \mathbb{R}^N$ to obtain transformed host signal samples $\tilde{x}_1, \ldots, \tilde{x}_{N/L_{\rm ST}}$, which are quantized using QIM. Because projection onto the vectors $\mathbf{v}_i$ represents a change of orthonormal basis, the transformed host signal samples and the transformed noise samples $\tilde{n}_1, \ldots, \tilde{n}_{N/L_{\rm ST}}$, which are the projections of the original noise vector $\mathbf{n} = [n_1 \cdots n_N]^T$ onto the orthonormal vectors $\mathbf{v}_i$, are still independent, zero-mean, Gaussian random variables with the same variance as the original host signal and noise samples, respectively. However, if the distortion per original host signal sample is $D_{\mathbf{s}}$, then the distortion per transformed host signal sample is $L_{\rm ST} D_{\mathbf{s}}$. Thus, we obtain a "spreading gain" of $L_{\rm ST}$ in terms of DNR, but the number of bits embedded per original host signal sample is only $1/L_{\rm ST}$ times the number of bits embedded per transformed host signal sample. Thus, one can determine an achievable rate $R_{\rm STQIM}$ of spread-transform QIM by appropriately modifying (44) to obtain

$$\begin{aligned} R_{\rm STQIM} &\geq \frac{1}{2L_{\rm ST}} \log_2 \left(L_{\rm ST} \cdot {\rm DNR} \frac{1 + L_{\rm ST} \cdot {\rm DNR} + {\rm SNR}_x}{L_{\rm ST} \cdot {\rm DNR} + {\rm SNR}_x} \right) \\ &\geq \frac{1}{2L_{\rm ST}} \log_2 (L_{\rm ST} \cdot {\rm DNR}). \end{aligned} \tag{45}$$

To upper bound the gap between QIM and capacity we first recognize from (45) that the minimum DNR required for QIM to achieve a rate R asymptotically with large N is

$${\rm DNR}_{\rm QIM} \leq \frac{2^{2L_{\rm ST}R}}{L_{\rm ST}}, \tag{46}$$

which is minimized at $L_{\rm ST} = 1/(2R \ln 2)$.[19] However, $L_{\rm ST} \geq 1$ even in the limit of large N to have

[18] The pdf of the reconstruction points $u = s$ in this case is $\mathcal{N}(0, D_{\mathbf{s}} + \sigma_x^2)$, which is not the same as the well-known rate-distortion optimal pdf [31] for quantizing Gaussian random variables, which is $\mathcal{N}(0, \sigma_x^2 - D_{\mathbf{s}})$.

[19] Note that since

$$\frac{N}{\left(\frac{N}{L_{\rm ST}} + 0.5\right)} \leq \frac{N}{{\rm round}\left(\frac{N}{L_{\rm ST}}\right)} \leq \frac{N}{\left(\frac{N}{L_{\rm ST}} - 0.5\right)},$$

one can indeed approach this optimum spreading gain $L_{\rm ST}$ in the limit of large N even though $N/L_{\rm ST}$ need be a

$N/L_{\mathrm{ST}} \leq N$. Thus, if one sets

$$L_{\mathrm{ST}} = \max\left\{\frac{1}{2R\ln 2}, 1\right\}, \tag{47}$$

then (46) remains a valid upper bound on the required DNR for a QIM method to achieve a rate R. From (38) we see that the minimum DNR required for a capacity-achieving method to achieve a rate R is $\mathrm{DNR}_{\mathrm{opt}} = 2^{2R} - 1$, which when combined with (46) yields the following upper bound between QIM and the Gaussian capacity:

$$\frac{\mathrm{DNR}_{\mathrm{QIM}}}{\mathrm{DNR}_{\mathrm{opt}}} \leq \frac{2^{2L_{\mathrm{ST}}R}}{L_{\mathrm{ST}}\left(2^{2R}-1\right)}. \tag{48}$$

This expression is plotted in Fig. 9, where L_{ST} is given by (47).

We now examine the asymptotic limits of (48) at low and high rates. Eq. (47) implies $L_{\mathrm{ST}} = 1/(2R\ln 2)$ in the limit of small R, so in this limit (48) approaches

$$\begin{aligned}
\frac{\mathrm{DNR}_{\mathrm{QIM}}}{\mathrm{DNR}_{\mathrm{opt}}} &\leq \frac{2^{2L_{\mathrm{ST}}R}}{L_{\mathrm{ST}}\left(2^{2R}-1\right)} \\
&= \frac{2^{1/\ln 2}(2R\ln 2)}{2^{2R}-1} \\
&= e\,\frac{2R\ln 2}{2^{2R}-1} \to e, \qquad \text{as } R \to 0.
\end{aligned}$$

Thus, the gap is at most a factor of e (approximately 4.3 dB) in the limit of low rates. In the limit of large R, (47) implies $L_{\mathrm{ST}} = 1$ so (48) approaches

$$\frac{\mathrm{DNR}_{\mathrm{QIM}}}{\mathrm{DNR}_{\mathrm{opt}}} = \frac{2^{2R}}{2^{2R}-1} \to 1, \qquad \text{as } R \to \infty.$$

Thus, QIM asymptotically achieves capacity at high embedding rates.

As we described in Sec. 7.2, in hybrid transmission applications one may be concerned about the degradation to the received host signal, which is $(1+\mathrm{DNR})$ rather than DNR. The gap in DNR (48) is larger than the gap in $(1+\mathrm{DNR})$, which has a corresponding upper bound

$$\frac{1+\mathrm{DNR}_{\mathrm{QIM}}}{1+\mathrm{DNR}_{\mathrm{opt}}} \leq \frac{1+\frac{2^{2RL_{\mathrm{ST}}}}{L_{\mathrm{ST}}}}{2^{2R}}.$$

This gap is plotted in Fig. 10 as a function of $2R$, the rate in b/s/Hz. Again, L_{ST} is given by (47)

positive integer less than or equal to N.

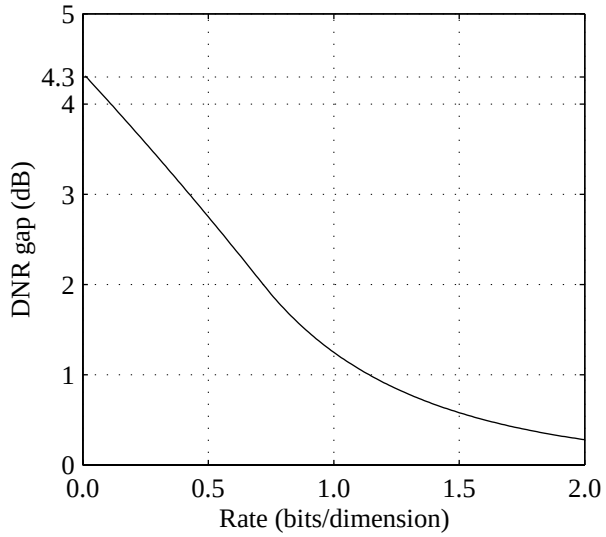

Figure 9: DNR gap between spread-transform QIM and Gaussian capacity (achieved by DC-QIM). The maximum gap is a factor of e ($\approx 4.3\,\mathrm{dB}$).

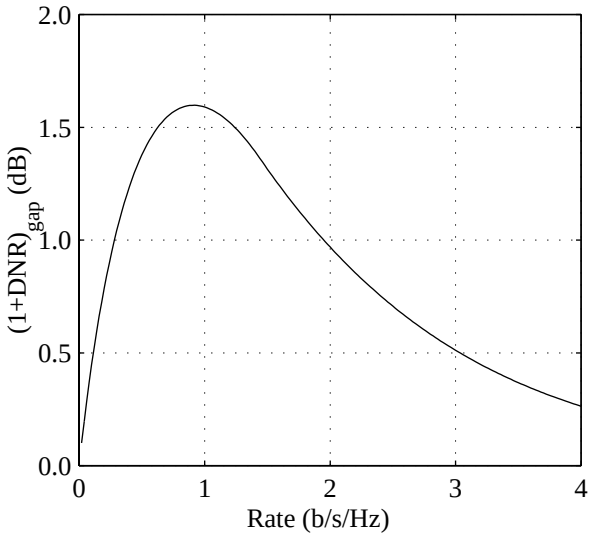

Figure 10: Received host SNR gap (1+DNR) between spread-transform QIM and capacity (achieved by DC-QIM).

since minimizing $\mathrm{DNR}_{\mathrm{QIM}}$ also minimizes $1+\mathrm{DNR}_{\mathrm{QIM}}$. Thus, for example, at the (near) worst-case digital rate of 1 b/s/Hz using QIM requires at most 1.6 dB more drop in analog channel quality than the approximately 3 dB drop required for DC-QIM (Sec. 7.2.1).

7.3.2 Uncoded STDM gap to capacity

The results above can be compared to the achievable performance of uncoded binary spread-transform dither modulation (STDM) with uniform scalar quantization as a minimal-complexity realization of QIM.

The gap between uncoded STDM and capacity can easily be quantified for low rates ($R_m \leq 1$), which are typical in many applications, at a given probability of error. A straightforward union bound on the bit-error probability of uncoded binary STDM with uniform scalar quantization is (see Fig. 6)

$$P_b \leq 2\,Q\left(\sqrt{\frac{d_{\min}^2}{4\sigma_n^2}}\right).$$

This bound is reasonably tight for low error probabilities, and from (26) we can write this probability of error in terms of the rate-normalized distortion-to-noise ratio $\mathrm{DNR}_{\mathrm{norm}} = \mathrm{DNR}/R_m$,

$$P_b \approx 2\,Q\left(\sqrt{\frac{3\cdot\mathrm{DNR}}{4R_m}}\right) = 2\,Q\left(\sqrt{\frac{3}{4}\mathrm{DNR}_{\mathrm{norm}}}\right). \tag{49}$$

A capacity-achieving method can achieve arbitrarily low probability of error as long as $R_m \leq C_{\mathrm{Gauss}}$, which using (38) can be expressed as

$$\frac{\mathrm{DNR}}{2^{2R_m}-1} \geq 1.$$

For low embedding rates R_m, $2^{2R_m}-1 \approx 2R_m \ln 2$ so the minimum required $\mathrm{DNR}_{\mathrm{norm}}$ for arbitrarily low probability of error is

$$\mathrm{DNR}_{\mathrm{norm}} \geq 2\ln 2 \approx 1.4 \text{ dB}. \tag{50}$$

The probability of error P_b of STDM is plotted as a function of $\mathrm{DNR}_{\mathrm{norm}}$ in Fig. 11. The required $\mathrm{DNR}_{\mathrm{norm}}$ for a given P_b can be compared to (50) to determine the gap to capacity. For example, at an error probability of 10^{-6}, uncoded STDM is about 13.6 dB from capacity. One can reduce this gap by at least 9.3 dB through channel coding, vector quantization, and non-dithered quantization. The remaining gap (at most 4.3 dB) is the gap between QIM and capacity and can be closed with distortion compensation. As shown in [15, 27], it is fairly easily to close the gap between uncoded

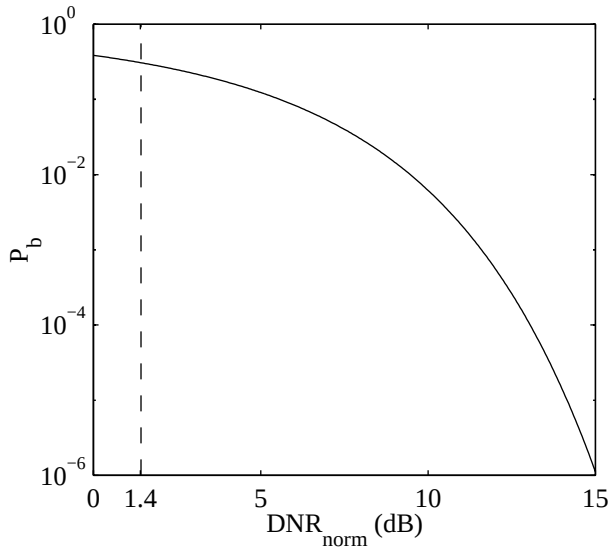

Figure 11: Uncoded spread-transform dither modulation (STDM) gap to Gaussian capacity. The solid curve shows the bit-error probability for uncoded STDM as a function of rate-normalized distortion-to-noise ratio ($\mathrm{DNR}_{\mathrm{norm}}$). The dashed curve is the minimum required $\mathrm{DNR}_{\mathrm{norm}}$ for reliable information-embedding for any embedding method.

STDM (with uniform scalar quantizers) and capacity by about 6 dB using practical channel codes and distortion compensation.

7.3.3 Uncoded spread LBM gap to capacity

The gap to capacity for uncoded binary spread LBM based on uniform, scalar quantization also follows readily from the results of App. A, which shows that the distortion-normalized minimum distance for this form of spread LBM is a factor of $7/4 \approx 2.43$ dB worse than that of STDM (26). Thus, the LBM counterpart to (49) is that the bit-error probability of uncoded spread LBM is

$$P_b \approx 2\,Q\left(\sqrt{\frac{3}{7}\mathrm{DNR}_{\mathrm{norm}}}\right). \tag{51}$$

Thus, the gap to capacity of uncoded binary spread LBM at an error probability of 10^{-6} is about 16 dB, 2.4 dB more than the 13.6-dB gap of uncoded binary STDM. Furthermore, as also discussed in App. A, for M-ary implementations the gap widens by an additional 0.6 dB as $M \to \infty$.

7.3.4 Coded additive spread-spectrum gap to capacity

For additive spread-spectrum, where $s = x + w(m)$, the distortion signal in Fig. 2 is not a function of the host signal: $e(x, m) = w(m)$. Thus, $y = s + n = e + x + n$. The distortion constraint is still $E[e^2] = D_{\mathbf{s}}$ so that in the Gaussian case considered here, the achievable rate of an additive spread-spectrum method is the well-known [31] Gaussian channel capacity, treating both x and n as interference sources,[20]

$$R_{\mathrm{SS}} = \frac{1}{2}\log_2\left(1 + \frac{D_{\mathbf{s}}}{\sigma_x^2 + \sigma_n^2}\right) = \frac{1}{2}\log_2\left(1 + \frac{\mathrm{DNR}}{\mathrm{SNR}_x + 1}\right), \tag{52}$$

where, again, SNR_x is the ratio between the host signal variance and the channel noise variance. Comparing (52) to (38), we see that the gap to capacity of additive spread-spectrum is

$$\frac{\mathrm{DNR}_{\mathrm{ss}}}{\mathrm{DNR}_{\mathrm{opt}}} = \mathrm{SNR}_x + 1, \tag{53}$$

which is typically large, since SNR_x must be large so that channel noise will not excessively degrade signal quality.

In fact, in the high signal-to-distortion (SDR) limit where $\sigma_x^2/D_{\mathbf{s}} \gg 1$, the achievable rate of additive spread-spectrum (52) clearly approaches zero, again reflecting the inability of additive spread-spectrum methods to reject host signal interference like other methods.

At the opposite extreme, when $\mathrm{SNR}_x \to 0$ the host interference is small so the gap (53) disappears, and indeed additive spread spectrum is an optimum embedding strategy for this case, along with both DC-QIM and QIM as discussed at the end of Sec. 7.1.

The other scenario in which additive spread-spectrum can be optimal is when the host is known at the decoder, which also corresponds to a non-interfering host situation.

7.3.5 Known-host case

As discussed at the end of Sec. 7.1, both capacity-achieving QIM and capacity-achieving additive spread-spectrum methods exist when the host signal is known at the decoder. Although QIM realizations in the form of coded dither modulation with uniform, scalar quantization are not optimal in this case, for AWGN channels one can achieve performance within $\pi e/6 \approx 1.53$ dB of capacity as we show below. We consider the case of dither signals with a uniform distribution over

[20] This rate is also the capacity when n is non-Gaussian, but still independent of s, and a correlation detector is used for decoding [35].

the interval $[-\Delta/2, \Delta/2]$. In this case,

$$s = q(x + d) - d = x + e,$$

where the quantization error e is uniformly distributed over the interval $[-\Delta/2, \Delta/2]$ and statistically independent of x [29]. Thus, the achievable rate $I(e; e+n)$ is slightly lower than the case where e is Gaussian. The entropy power inequality can be used to show that the decrease in achievable rate is bounded by [36]

$$C_{\text{Gauss,known}} - R_{\text{dith}} \leq \frac{1}{2} \log_2 \frac{1 + \text{DNR}}{1 + (6/\pi e)\text{DNR}}. \tag{54}$$

This gap approaches the upper limit of $\frac{1}{2} \log_2 \frac{\pi e}{6} \approx 0.2546$ bits/dimension as the DNR gets large. For any finite DNR, the gap is smaller. By subtracting the upper bound on the gap (54) from the capacity (38), one obtains a lower bound on the achievable rate of this type of dither modulation:

$$R_{\text{dith}} \geq \frac{1}{2} \log_2 \left(1 + \frac{6}{\pi e} \text{DNR} \right). \tag{55}$$

Thus, dither modulation with uniform scalar quantization in this case is at most $\pi e/6 \approx 1.53$ dB from capacity.

8 Intentional Attacks

We now turn our attention from AWGN channel models for unintentional attacks, to some alternative models for intentional attacks. Intentional, distortion-constrained attacks may be encountered in copyright, authentication, and covert communication applications. In these kinds of applications, attackers generally attempt to remove or alter the embedded information, and face a distortion constraint on their signal manipulations so that the integrity of the host signal is not compromised.

An attacker's ability to prevent reliable watermark decoding depends on the amount of knowledge that the attacker has about the embedding and decoding processes. To limit such knowledge, some digital watermarking systems use keys, parameters that allow appropriate parties to embed and/or decode the embedded signal. The locations of the modulated bits and the pseudo-noise vectors in an additive spread-spectrum and generalized LBM systems are examples of keys. If only certain parties privately share the keys to both embed and decode information, and no one else can do either of these two functions, then the watermarking system is a private-key system. Alterna-

tively, if some parties possess keys that allow them to either embed or decode, but not both, then the system is a public-key system since these keys can be made available to the public for use in one of these two functions without allowing the public to perform the other function. However, in some scenarios it may be desirable to allow everyone to embed and decode watermarks without the use of keys. For example, in a copyright ownership notification system, everyone could embed the ASCII representation of a copyright notice such as, "Property of ..." in their copyrightable works. Such a system is analogous to the system currently used to place copyright notices in (hardcopies of) books, a system in which there is no need for a central authority to store, register, or maintain separate keys—there are none—or watermarks—all watermarks are English messages—for each user. The widespread use of such a universally accessible "no-key" system requires only standardization of the decoder so that everyone will agree on the decoded watermark, and hence, the owner of the copyright.

We analyze both private-key and no-key systems in the sequel, and establish the attractiveness of QIM in both cases.

8.1 Attacks on Private-key Systems

Although the attacker does not know the key in a private-key scenario, he or she may know the basic algorithm used to embed the watermark. In [16], Moulin and O'Sullivan model such a scenario by assuming that the attacker knows the codebook distribution, but not the actual codebook. As we now develop, exploiting results results of Moulin and O'Sullivan in this private-key scenario, we determine that DC-QIM methods are optimal (capacity-achieving) against squared error distortion-constrained attackers.

Moulin and O'Sullivan have derived both the capacity-achieving distribution and an explicit expression for the capacity (32) in the case where the host is white and Gaussian and the attacker faces an expected perturbation energy constraint $E[\|\mathbf{n}\|^2] \leq \sigma_n^2$. In this case the capacity is [16]

$$C_{\text{Gauss,private}} = \frac{1}{2}\log_2\left(1 + \frac{\text{DNR}_{\text{attack}}}{\beta}\right), \qquad \beta = \frac{\text{SNR}_{x,\text{attack}} + \text{DNR}_{\text{attack}}}{\text{SNR}_{x,\text{attack}} + \text{DNR} - 1},$$

where $\text{DNR}_{\text{attack}} = D_{\mathbf{s}}/\sigma_n^2$ is the distortion-to-perturbation ratio and $\text{SNR}_{x,\text{attack}} = \sigma_x^2/\sigma_n^2$ is the host signal-to-perturbation ratio. The maximizing distribution is such that [16]

$$e = u - \alpha_{\text{Gauss,private}} x,$$

with $e \sim \mathcal{N}(0, D_{\mathbf{s}})$ statistically independent of x and

$$\alpha_{\text{Gauss,private}} = \frac{\text{DNR}_{\text{attack}}}{\text{DNR}_{\text{attack}} + \beta}. \tag{56}$$

Since this distribution satisfies the condition (33), we can infer from our analysis in Sec. 6.1.2 that distortion-compensated QIM can be used to achieve capacity against these attacks. Moreover, (56) gives the optimal distortion-compensation parameter.

Moulin and O'Sullivan have also considered the case of host signals that are not necessarily Gaussian but that have zero-mean, finite-variance, and bounded and continuous pdfs. In the limit of small $D_{\mathbf{s}}$ (high SDR) and σ_n^2, a limit of interest in high-fidelity applications, the capacity approaches

$$C_{\text{high-fidelity}} \to \frac{1}{2} \log_2 \left(1 + \text{DNR}_{\text{attack}}\right),$$

and the capacity-achieving distribution is such that

$$e \to u - \alpha_{\text{high-fidelity}} x,$$

where, again, $e \sim \mathcal{N}(0, D_{\mathbf{s}})$ is statistically independent of x [16]. Since this distribution satisfies the condition (33), we can again conclude that distortion-compensated QIM can achieve capacity in this high-fidelity limit. The capacity-achieving distortion-compensation parameter is [16]

$$\alpha_{\text{high-fidelity}} = \frac{\text{DNR}_{\text{attack}}}{\text{DNR}_{\text{attack}} + 1}.$$

8.2 Attacks on No-Key Systems

In contrast to the scenario above, in no-key systems an attacker has full knowledge of the embedding and decoding processes, including all codebooks. For this case, some deterministic models we develop in this section are better for characterizing the associated worst-case in-the-clear (i.e., fully informed) attacks. With these models, we show that QIM methods in general, and dither modulation in particular, are robust and achieve provably better rate-distortion-robustness trade-offs than both additive spread-spectrum and generalized LBM techniques.

We consider two models for such attackers: (1) a bounded perturbation channel model in which the squared error distortion between the channel input and channel output is bounded and (2) a bounded host-distortion channel model in which the squared error distortion between the host signal and channel output is bounded. In each case, we develop conditions under which error-

free decoding is possible with various implementations of QIM and DC-QIM, and quantify their advantages over the corresponding realizations of additive spread-spectrum and generalized LBM.

8.2.1 Bounded perturbation channel

The bounded perturbation channel is one in which the attacker can perturb the composite signal in any way it desires (based on its full knowledge of the composite signal and the embedding algorithm), provided the energy in the perturbation vector does not exceed a prescribed level, i.e., (16), which reflects a requirement that the attacker not excessively degrade the original composite signal. Thus, this channel model imposes only a maximum distortion[21] or minimum SNR constraint between the channel input and output.

Binary dither modulation with uniform scalar quantization One can combine the guaranteed error-free decoding condition (17) for a minimum distance decoder (15) with the distortion-normalized minimum distance (26) of binary dither modulation with uniform scalar quantization to compactly express its achievable performance as

$$\frac{(d_{\min}^2/D_{\mathbf{s}})D_{\mathbf{s}}}{4N\sigma_n^2} = \gamma_c \frac{3/4}{NR_m}\frac{D_{\mathbf{s}}}{\sigma_n^2} > 1, \tag{57}$$

or, equivalently, its achievable rate as

$$\sup R_m = \frac{3\gamma_c}{4N}\frac{D_{\mathbf{s}}}{\sigma_n^2}. \tag{58}$$

One can view the achievable rate (58) as the deterministic counterpart to the more conventional notions of achievable rates and capacities of random channels discussed in Secs. 6 and 7.

Additive spread-spectrum The nonzero minimum distance of QIM methods offers quantifiable robustness to perturbations, even when the host signal is not known at the decoder. In contrast, additive spread-spectrum methods offer relatively little robustness if the host signal is not known

[21] Some types of distortion, such as geometric distortions, can be large in terms of squared error, yet still be small perceptually. However, in some cases these distortions can be mitigated either by preprocessing at the decoder or by embedding information in parameters of the host signal that are less affected (in terms of squared error) by these distortions. For example, a simple delay or shift may cause large squared error, but the magnitude of the DFT coefficients are relatively unaffected.

at the decoder. As discussed in Sec. 3, these methods have linear embedding functions of the form

$$\mathbf{s}(\mathbf{x}, m) = \mathbf{x} + \mathbf{w}(m), \tag{59}$$

where $\mathbf{w}(m)$ is a pseudo-noise vector. From the definition of minimum distance (13),

$$\begin{aligned} d_{\min} &= \min_{(i,j):i\neq j} \min_{(\mathbf{x}_i,\mathbf{x}_j)} \|\mathbf{x}_i + \mathbf{w}(i) - \mathbf{x}_j - \mathbf{w}(j)\| \\ &= \min_{(i,j):i\neq j} \|\mathbf{x}_i + \mathbf{w}(i) - (\mathbf{x}_i + \mathbf{w}(i) - \mathbf{w}(j)) - \mathbf{w}(j)\| \\ &= 0, \end{aligned}$$

i.e., the minimum distance is zero.

Thus, although these methods may be effective when the host signal is known at the decoder, when the host signal is not known, they offer no guaranteed robustness to perturbations, i.e., no achievable rate expression analogous to (58) exists for additive spread-spectrum. As is evident from (59), in an additive spread-spectrum system, $\mathbf{x}$ is an additive interference, which is often much larger than $\mathbf{w}$ due to the distortion constraint. In contrast, the quantization that occurs with QIM provides immunity against this host signal interference, as discussed in Sec. 4.[22]

Generalized LBM As shown in App. A, the distortion-normalized minimum distance of generalized binary LBM with uniform scalar quantization is about 2.43 dB worse than that of the corresponding dither modulation strategy. Therefore, its achievable rate-distortion-robustness performance is also about 2.43 dB worse than (57). Again, as also developed in the appendix, for M-ary implementations, the gap grows to 3 dB for large M.

8.2.2 Bounded host-distortion channel

As an alternative to the bounded perturbation channel, some attackers may work with distortion constraint between the channel output and the host signal, rather than the channel input, since this distortion is the most direct measure of degradation to the host signal. For example, if attackers have partial knowledge of the host signal, which may be in the form of a probability distribution, so that they can calculate this distortion, then it may be appropriate to bound the expected distortion

[22] Another way to understand this host-signal interference rejection is to consider, for example, that a quantized random variable has finite entropy while a continuous random variable has infinite entropy.

Table 1: Attacker's distortion penalties. The distortion penalty is the additional distortion that an attacker must incur to successfully remove a watermark. A distortion penalty less than 0 dB indicates that the attacker can actually improve the signal quality and remove the watermark simultaneously.

Embedding Method	Distortion Penalty ($D_{\mathbf{y}}/D_{\mathbf{s}}$)
Regular QIM	$1 + \frac{d^2_{\text{norm}}}{4N} > 0$ dB
Binary Dith. Mod. w/uni. scalar quant.	2.43 dB $\geq 1 + \gamma_c \frac{3/4}{NR_m} > 0$ dB
DC-QIM	$-\infty$ dB
Additive spread-spectrum	$-\infty$ dB
LBM	≤ 0 dB
Binary LBM w/uni. scalar quant.	-2.43 dB

$D_{\mathbf{y}} = E[D(\mathbf{y}, \mathbf{x})]$, where this expectation is taken over the conditional probability density $p_{\mathbf{x}|\mathbf{s}}(\mathbf{x}|\mathbf{s})$.[23] We refer to this as the bounded host-distortion channel.

For this channel we measure robustness to attacks by the minimum expected distortion $D_{\mathbf{y}}$ for a successful attack, where the expectation is taken with respect to $p_{\mathbf{x}|\mathbf{s}}(\mathbf{x}|\mathbf{s})$. The ratio between $D_{\mathbf{y}}$ and the expected embedding-induced distortion $D_{\mathbf{s}}$ is the distortion penalty that the attacker must pay to remove the watermark and, hence, is a figure of merit measuring the robustness-distortion trade-off at a given rate. Distortion penalties for the primary methods of interest are derived below and summarized in Table 1 for the high SDR regime of primary interest. As this table reflects, among these methods considered, only QIM methods (including binary dither modulation with uniform scalar quantization) are robust enough such that the attacker must degrade the host signal quality to remove the watermark.

Regular QIM We first consider the robustness of regular quantization index modulation. For any distortion measure, as long as each reconstruction point $\mathbf{s}$ lies at the minimum distortion point of its respective quantization cell, the QIM distortion penalty is greater than or equal to 1 since

[23] Note that if the attacker has full knowledge of the host signal, he or she can trivially remove the embedded information by setting $\mathbf{y} = \mathbf{x}$, so $D_{\mathbf{y}} = 0$. We restrict our attention to the more realistic scenario in which an attacker has only partial knowledge of the host, in the form of a conditional pdf.

any output $\mathbf{y}$ that an attacker generates must necessarily lie away from this minimum distortion point. Equality occurs only if each quantization cell has at least two minimum distortion points, one of which lies in the incorrect decoder decision region. For expected squared-error distortion, the minimum distortion point of each quantization cell is its centroid, and one can express this distortion penalty in terms of the distortion-normalized minimum distance and the signal length N, as we show below.

We use $\mathcal{R}$ to denote the quantization cell containing $\mathbf{x}$ and $p_{\mathbf{x}}(\mathbf{x}|\mathcal{R})$ to denote the conditional probability density function of $\mathbf{x}$ given that $\mathbf{x} \in \mathcal{R}$. Again, for sufficiently small quantization cells, this probability density function can often be approximated as uniform over $\mathcal{R}$, for example. Since $\mathbf{s}$ is the centroid of $\mathcal{R}$,

$$\int_{\mathcal{R}} (\mathbf{s} - \mathbf{x}) p_{\mathbf{x}}(\mathbf{x}|\mathcal{R})\, d\mathbf{x} = \mathbf{0}. \tag{60}$$

Also, the expected squared-error per letter embedding-induced distortion given $\mathbf{x} \in \mathcal{R}$ is

$$D_{\mathbf{s}|\mathcal{R}} = \frac{1}{N} \int_{\mathcal{R}} \|\mathbf{s} - \mathbf{x}\|^2 p_{\mathbf{x}}(\mathbf{x}|\mathcal{R})\, d\mathbf{x}. \tag{61}$$

The most general attack can always be represented as $\mathbf{y} = \mathbf{s} + \mathbf{n}$, where $\mathbf{n}$ may be a function of $\mathbf{s}$. The resulting distortion is

$$\begin{aligned} D_{\mathbf{y}|\mathcal{R}} &= \frac{1}{N} \int_{\mathcal{R}} \|\mathbf{y} - \mathbf{x}\|^2 p_{\mathbf{x}}(\mathbf{x}|\mathcal{R})\, d\mathbf{x} \\ &= \frac{1}{N} \int_{\mathcal{R}} \|(\mathbf{s} - \mathbf{x}) + \mathbf{n}\|^2 p_{\mathbf{x}}(\mathbf{x}|\mathcal{R})\, d\mathbf{x} \\ &= \frac{1}{N} \int_{\mathcal{R}} \|\mathbf{s} - \mathbf{x}\|^2 p_{\mathbf{x}}(\mathbf{x}|\mathcal{R})\, d\mathbf{x} + \frac{1}{N}\|\mathbf{n}\|^2 \int_{\mathcal{R}} p_{\mathbf{x}}(\mathbf{x}|\mathcal{R})\, d\mathbf{x} + \frac{2}{N}\mathbf{n}^{\mathrm{T}} \int_{\mathcal{R}} (\mathbf{s} - \mathbf{x}) p_{\mathbf{x}}(\mathbf{x}|\mathcal{R})\, d\mathbf{x} \\ &= D_{\mathbf{s}|\mathcal{R}} + \frac{\|\mathbf{n}\|^2}{N}, \end{aligned}$$

where we have used (61), the fact that $p_{\mathbf{x}}(\mathbf{x}|\mathcal{R})$ is a probability density function and, thus, integrates to one, and (60) to obtain the last line. For a successful attack, $\|\mathbf{n}\| \geq d_{\min}/2$ so

$$D_{\mathbf{y}|\mathcal{R}} \geq D_{\mathbf{s}|\mathcal{R}} + \frac{d_{\min}^2}{4N}.$$

Averaging both sides of this expression over all quantization cells $\mathcal{R}$ yields $D_{\mathbf{y}} \geq D_{\mathbf{s}} + d_{\min}^2/4N$ so that our figure of merit for quantization index modulation methods is

$$\frac{D_{\mathbf{y}}}{D_{\mathbf{s}}} \geq 1 + \frac{d_{\min}^2/D_{\mathbf{s}}}{4N} = 1 + \frac{d_{\mathrm{norm}}^2}{4N}. \tag{62}$$

Thus, for any QIM method of nonzero distortion-normalized minimum distance d_{norm}, the attacker's distortion penalty is always greater than 1 (0 dB), indicating that to remove the watermark, the attacker must degrade the host signal quality beyond the initial distortion caused by the embedding of the watermark.

Binary dither modulation with uniform, scalar quantization In this case, (26) gives d^2_{norm} in (62). Moreover, due to the uniformity of the quantizers, the bound (62) is met with equality so that the attacker's distortion penalty specializes to

$$\frac{D_{\mathbf{y}}}{D_{\mathbf{s}}} = 1 + \gamma_c \frac{3/4}{NR_{\boldsymbol{m}}}. \tag{63}$$

Because the Hamming distance d_H of a block code cannot exceed the number of coded bits $NR_{\boldsymbol{m}}(k_c/k_u)$,

$$\frac{\gamma_c}{NR_{\boldsymbol{m}}} = \frac{d_H}{NR_{\boldsymbol{m}}(k_c/k_u)} \leq 1,$$

where the first equality follows from the definition (23) of γ_c. Thus, an upper bound for the distortion penalty (63) in this case is

$$1 + \gamma_c \frac{3/4}{NR_{\boldsymbol{m}}} \leq \frac{7}{4} \approx 2.43 \text{ dB}.$$

Although this penalty may seem modest, it is larger than that obtainable by either additive spread-spectrum or generalized LBM, as we show below. Larger distortion penalties are not possible because in-the-clear attackers can concentrate all their distortion in the minimum distance direction in N-dimensional space.

As a final note, (63) implies that binary dither modulation with uniform, scalar quantization can defeat any attacker as long as

$$\left(1 + \gamma_c \frac{3/4}{NR_{\boldsymbol{m}}}\right) \frac{D_{\mathbf{s}}}{D_{\mathbf{y}}} > 1,$$

an expression whose counterpart for the bounded perturbation channel was (57). Thus, the corresponding achievable rates are given by

$$\sup R_{\boldsymbol{m}} = \frac{3\gamma_c}{4N} \left(\frac{D_{\mathbf{y}}}{D_{\mathbf{s}}} - 1\right)^{-1}.$$

Distortion-compensated QIM An in-the-clear attacker of a DC-QIM system knows the quantizers and can determine the watermark m after observing the composite signal $\mathbf{s}$. If the quantization cells are contiguous so that the distortion-compensation term in (19) does not move $\mathbf{s}$ out of the cell containing $\mathbf{x}$, then an attacker can recover the original host signal with the following attack:

$$\begin{aligned}\mathbf{y} &= \frac{\mathbf{s} - \alpha\mathbf{q}(\mathbf{s}; m, \Delta/\alpha)}{1-\alpha} \\ &= \frac{\mathbf{s} - \alpha\mathbf{q}(\mathbf{x}; m, \Delta/\alpha)}{1-\alpha} \\ &= \mathbf{x},\end{aligned}$$

where the final line follows simply by inverting (19). Thus, the attacker's distortion penalty $D_{\mathbf{y}}/D_{\mathbf{s}}$ is $-\infty$ dB. We see that although DC-QIM is optimal against both independent additive Gaussian noise attacks and squared error distortion-constrained attacks in private-key scenarios, it is in some sense "maximally suboptimal" against in-the-clear attacks. Regular QIM, on the other hand, is almost as good as DC-QIM against additive Gaussian noise attacks (Sec. 7) and also resistant to in-the-clear attacks as discussed above. Thus, regular QIM methods may offer an attractive compromise when one requires resistance to both intentional attacks and unintentional attacks in a no-key system.

Additive spread-spectrum Since the embedding function of an additive spread-spectrum system is (2), the resulting distortion is $D_{\mathbf{s}} = \|\mathbf{w}\|^2/N > 0$. An attacker with full knowledge of the embedding and decoding processes can decode the message m, and hence, reproduce the corresponding pseudo-noise vector $\mathbf{w}$. Therefore, the attacker can completely remove the watermark by subtracting $\mathbf{w}$ from $\mathbf{s}$ to obtain the original host signal, i.e., $\mathbf{y} = \mathbf{s} - \mathbf{w}(m) = \mathbf{x}$. Hence, the resulting distortion penalty is $D_{\mathbf{y}}/D_{\mathbf{s}} = 0/D_{\mathbf{s}} = -\infty$ dB.

Because the additive spread-spectrum embedding function combines the host signal $\mathbf{x}$ and watermark $\mathbf{w}(m)$ in a simple linear way, anyone that can extract the watermark, can easily remove it. Thus, these methods are not suitable for universally accessible no-key digital watermarking applications. By contrast, the advantage of QIM is that it effectively hides the host signal even when the embedded information m is known.

Generalized LBM Recall that the embedding function of a generalized LBM system can be written as (8) with $\mathbf{q}(\cdot)$ having the property (9). Good generalized LBM systems also have the property that the reconstruction points of $\mathbf{q}(\cdot)$ are at the centroids of the quantization cells, as we'll

assume. One attack that completely removes information about m is to output these reconstruction points, i.e., $\mathbf{y} = \mathbf{q}(\mathbf{s}) = \mathbf{q}(\mathbf{x})$. Since $\mathbf{y}$ is at a minimum distortion point of the quantization cell, $D_{\mathbf{y}}/D_{\mathbf{s}} \leq 1 = 0$ dB, with equality only if both $\mathbf{s}$ and $\mathbf{y}$ are minimum distortion points. Thus, an attacker can remove the watermark without causing additional distortion to the host signal. This result applies regardless of whether error correction coding is used. Thus, in contrast to dither modulation, error correction coding does not improve LBM in this context.

When, in addition, it is the least significant bit of a uniform, scalar quantizer that is modulated, the results in App. A imply further that

$$D_{\mathbf{s}} = \frac{7}{48L} \sum_k \Delta_k^2,$$

while

$$D_{\mathbf{y}} = \frac{1}{12L} \sum_k \Delta_k^2.$$

Thus, $D_{\mathbf{y}}/D_{\mathbf{s}} = 4/7 \approx -2.43$ dB. Again, when many less significant bits are modulated, the results of the appendix can be used to establish that the penalty grows to -3 dB.

9 Concluding Remarks

We have seen that QIM methods are provably better than additive spread-spectrum and generalized LBM against bounded perturbation and in-the-clear attacks and are near-optimal for Gaussian channels, for which distortion-compensated QIM (DC-QIM) is optimal. Furthermore, dither modulation is a practical implementation of QIM that exhibits many of the attractive performance properties of QIM. The convenient structure of dither modulation, which is easily combined with error correction coding, allows the system designer to achieve different rate-distortion-robustness trade-offs by tuning parameters such as the quantization step size. Also, one can conveniently upgrade previously developed additive spread-spectrum and spread LBM systems to spread-transform dither modulation systems by replacing the respective addition and quantize-and-replace steps with a dithered quantization step.

In the course of our investigation, a number of rather intriguing results have emerged. For example, the information-embedding capacity in the Gaussian case does not depend at all on whether the host signal is available during decoding, and DC-QIM is optimal in both scenarios, and achieves perfect rejection of host signal interference, even in the high SDR regime.

Also somewhat surprisingly, the optimal embedding strategy for Gaussian channels and for typical attacks in private-key systems, DC-QIM, is "maximally suboptimal" against in-the-clear attacks. On the other hand regular QIM, which has performance within 4.3 dB of DC-QIM in the Gaussian case, performs better than any other currently known method against in-the-clear attacks, which arise in copyright notification applications where no-key architectures are used, for example. In particular, unlike additive spread-spectrum and generalized LBM methods, QIM and dither modulation methods force an attacker to pay a distortion penalty. Thus, QIM emerges as a universally good embedding strategy against a wide variety of intentional and unintentional attacks.

For hybrid transmission strategies, using DC-QIM for digital-over-analog transmission (in for example digital audio broadcasting applications) allows embedding rates of about 1/3 b/s/Hz for every dB drop in analog signal quality. In digital-over-digital transmission (in broadcast applications, for example), DC-QIM is as efficient as any single digital transmission, and thus as good as the alternative superposition coding and successive cancellation decoding approach.

Many important directions for further research remain. At one end of the spectrum, further insights into the fundamental principles and structure of information embedding and digital watermarking systems will come from the development of still better general attack models. Those emerging from game-theoretic formulations and arbitrarily-varying channel models appear to be an important starting point in this respect.

At the same time, many of the results in this paper have important implications for practical applications, and the most effective implementations of QIM and DC-QIM embedding systems for these applications will take into account in detail the specific types of geometric distortions and other attacks that typically arise. For example, in image watermarking applications, embedders and decoders ultimately need to be robust to a wide range of often surprisingly challenging attacks, ranging from scaling and rotation, to cropping and column replacement. A great deal of future work is needed in this area to enable the use of QIM techniques in watermarking applications, and indeed these represent some especially interesting design challenges.

A LBM Distortion-normalized Minimum Distance

In this appendix we calculate the distortion-normalized minimum distance of binary low-bit(s) modulation (LBM) with uniform, scalar quantization. We assume that the host signal and embedded signal are statistically independent.

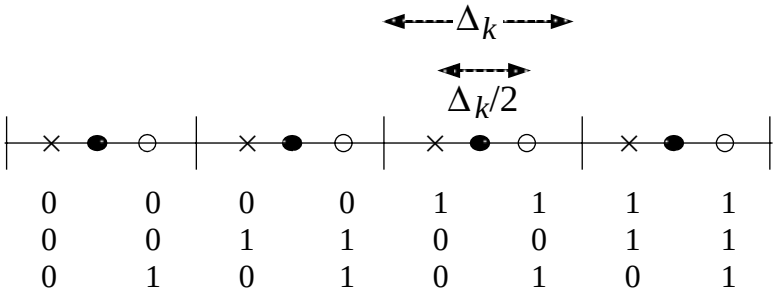

Figure 12: Low-bit modulation with a uniform, scalar quantizer. The quantizer has a step size of $\Delta_k/2$, and the least significant bit (lsb) is modulated. All reconstruction points marked with a $\times$ have a lsb of 0. Points marked with a $\circ$ have a lsb of 1. This process is equivalent to first quantizing using a quantizer with a step size of Δ_k, whose reconstruction points are marked with a $\bullet$, and adding $\pm\Delta_k/4$.

Since the embedding function of any good generalized LBM method can be written as (8) with (10), the expected distortion is

$$\begin{aligned}\frac{1}{N}E\left[\|\mathbf{s}-\mathbf{x}\|^2\right] &= \frac{1}{N}E\left[\|\mathbf{q}(\mathbf{x})-\mathbf{x}+\mathbf{d}(m)\|^2\right] \\ &= \frac{1}{N}E\left[\|\mathbf{q}(\mathbf{x})-\mathbf{x}\|^2+2(\mathbf{q}(\mathbf{x})-\mathbf{x})^T\mathbf{d}(m)+\|\mathbf{d}(m)\|^2\right] \\ &= \frac{1}{N}E\left[\|\mathbf{q}(\mathbf{x})-\mathbf{x}\|^2\right]+\frac{1}{N}E\left[\|\mathbf{d}(m)\|^2\right], \end{aligned} \tag{64}$$

where we have used (10) and the independence of $\mathbf{x}$ and m to obtain the final line.

We analyze coded binary LBM with uniform scalar quantization, an LBM system in which each in a sequence of coded bits is repeated L times and embedded in a length-L block with a sequence of uniform, scalar quantizers.

The embedding is accomplished by modulating the least significant bit of each quantizer. The kth uniform, scalar quantizer is illustrated in Fig. 12. The coarse quantizer $q_k(\cdot)$ has a step size of Δ_k, and the kth least significant bit adjustment element d_k equals $\pm\Delta_k/4$.

Comparing this scheme to coded binary dither modulation with uniform scalar quantization as described in Sec. 5.1, we see that this scheme has the same minimum distance, i.e., (22). Restricting attention to the high SDR regime in which $\mathbf{x}$ can be modeled as uniformly distributed within each cell of $\mathbf{q}(\cdot)$, as was used to develop (24) in Sec. 5.1, the first term in (64) is

$$\frac{1}{L}\sum_k E\left[\|q(x_k)-x_k\|^2\right] = \frac{1}{12L}\sum_k \Delta_k^2, \tag{65}$$

the same as the expected distortion (25) of the corresponding dither modulation system. The second term in (64), however, is

$$\frac{1}{L}\sum_k d_k^2 = \frac{1}{16L}\sum_k \Delta_k^2. \tag{66}$$

Thus, the overall expected distortion is

$$D_{\mathbf{s}} = \left(\frac{1}{12L} + \frac{1}{16L}\right)\sum_k \Delta_k^2 = \frac{7}{48L}\sum_k \Delta_k^2,$$

and the distortion-normalized squared minimum distance is

$$d^2_{\text{norm}} = \frac{12\gamma_c}{7R_{\boldsymbol{m}}}.$$

By comparing with (26), we see that binary coded LBM with uniform scalar quantization is worse than the corresponding dither modulation system by

$$\frac{3}{12/7} = \frac{7}{4} \approx 2.43\,\text{dB}. \tag{67}$$

Also, note the result (67) is invariant to the actual distribution of the Δ_k's, and invariant to any preprocessing of the host signal by a unitary transformation. Thus the gap between STDM and spread LBM is also given by (67).

In other variants of LBM, the gap can be worse. For instance, in the case of M-ary coded implementations of dither modulation and LBM based on uniform scalar quantization where the $M > 2$ sets of reconstruction points together form a regular lattice, then the minimum distances of the two schemes remain equal (but generally different from the binary case), and the first term in (64) remains (65). However, as M gets large, d_k becomes effectively uniformly distributed over the range $(-\Delta_k/2, \Delta_k)$, so the second term in (64) changes from (66) to

$$\frac{1}{L}\sum_k E\left[d_k\right] = \frac{1}{12L}\sum_k \Delta_k^2,$$

the same as (65). Thus the gap (67) grows to a factor of 2 (3 dB) in this large M limit.

Acknowledgment

The authors thank Prof. Amos Lapidoth for calling our attention to the paper by Costa, and an anonymous reviewer for helpful comments that improved the clarity of the paper.

References

[1] M. D. Swanson, M. Kobayashi, and A. H. Tewfik, "Multimedia data-embedding and watermarking technologies," *Proceedings of the IEEE*, vol. 86, pp. 1064–1087, June 1998.

[2] I. J. Cox and J.-P. M. G. Linnartz, "Some general methods for tampering with watermarks," *IEEE Journal on Selected Areas in Communications*, vol. 16, pp. 587–593, May 1998.

[3] J.-P. Linnartz, T. Kalker, and J. Haitsma, "Detecting electronic watermarks in digital video," in *Proc. of the 1999 IEEE International Conference on Acoustics, Speech, and Signal Processing*, vol. 4, (Phoenix, AZ), pp. 2071–2074, Mar. 1999.

[4] G. Arce, C. G. Boncelet, Jr., R. F. Graveman, and L. M. Marvel, "Applications of information hiding," in *Proc. of Third Annual Federated Laboratory Symposium on Advanced Telecommunications & Information Distribution Research Program*, (College Park, MD), pp. 423–427, Feb. 1999.

[5] H. C. Papadopoulos and C.-E. W. Sundberg, "Simultaneous broadcasting of analog FM and digital audio signals by means of adaptive precanceling techniques," *IEEE Transactions on Communications*, vol. 46, pp. 1233–1242, Sept. 1998.

[6] B. Chen and C.-E. W. Sundberg, "Broadcasting data in the FM band by means of adaptive contiguous band insertion and precancelling techniques," in *Proceedings of 1999 IEEE International Conference on Communications*, vol. 2, (Vancouver, Canada), pp. 823–827, June 1999.

[7] M. D. Swanson, B. Zhu, and A. H. Tewfik, "Data hiding for video-in-video," in *Proceedings of the 1997 IEEE International Conference on Image Processing*, vol. 2, (Piscataway, NJ), pp. 676–679, 1997.

[8] J. M. Barton, "Method and apparatus for embedding authentication information within digital data." United States Patent #5,646,997, July 1997.

[9] K. Tanaka, Y. Nakamura, and K. Matsui, "Embedding secret information into a dithered multi-level image," in *Proceedings of the 1990 IEEE Military Communications Conference*, pp. 216–220, 1990.

[10] W. Bender, D. Gruhl, N. Morimoto, and A. Lu, "Techniques for data hiding," *IBM Systems Journal*, vol. 35, no. 3-4, pp. 313–336, 1996.

[11] I. J. Cox, J. Killian, T. Leighton, and T. Shamoon, "A secure, robust watermark for multimedia," in *Information Hiding. First International Workshop Proceedings*, pp. 185–206, June 1996.

[12] J. R. Smith and B. O. Comiskey, "Modulation and information hiding in images," in *Information Hiding. First International Workshop Proceedings*, pp. 207–226, June 1996.

[13] J. R. Hernandez, F. Perez-Gonzalez, J. M. Rodriguez, and G. Nieto, "Performance analysis of a 2-D-multipulse amplitude modulation scheme for data hiding and watermarking of still images," *IEEE Journal on Selected Areas in Communications*, vol. 16, pp. 510–524, May 1998.

[14] J. A. O'Sullivan, P. Moulin, and J. M. Ettinger, "Information theoretic analysis of steganography," in *Proc. of the 1998 IEEE International Symposium on Information Theory*, (Cambridge, MA), p. 297, Aug. 1998.

[15] B. Chen and G. W. Wornell, "Implementations of quantization index modulation methods for digital watermarking and information embedding," to appear in *Journal of VLSI Signal Processing Systems for Signal, Image, and Video Technology, Special Issue on Multimedia Signal Processing*, 2000.

[16] P. Moulin and J. A. O'Sullivan, "Information-theoretic analysis of information hiding," *Preprint*, Oct. 1999.

[17] N. Merhav, "On random coding error exponents of watermarking systems," *IEEE Trans. Inform. Theory*, vol. 46, pp. 420–430, Mar. 2000.

[18] A. Cohen and A. Lapidoth, "On the gaussian watermarking game," in *IEEE Int. Symp. Inform. Theory*, p. 48, June 2000.

[19] P. Moulin and J. O'Sullivan, "Information-theoretic analysis of information hiding," in *IEEE Int. Symp. Inform. Theory*, p. 19, June 2000.

[20] R. J. Barron, B. Chen, and G. W. Wornell, "The duality between information embedding and source coding with side information and its implications and applications," submitted to *IEEE Transactions on Information Theory*, Jan. 2000.

[21] I. J. Cox, M. L. Miller, and A. L. McKellips, "Watermarking as communications with side information," *Proceedings of the IEEE*, vol. 87, pp. 1127–1141, July 1999.

[22] F. Hartung and M. Kutter, "Multimedia watermarking techniques," *Proceedings of the IEEE*, vol. 87, pp. 1079–1107, July 1999.

[23] F. A. P. Petitcolas, R. J. Anderson, and M. G. Kuhn, "Information hiding — a survey," *Proceedings of the IEEE*, vol. 87, pp. 1062–1078, July 1999.

[24] A. Z. Tirkel, G. A. Rankin, R. van Schyndel, W. J. Ho, N. R. A. Mee, and C. F. Osborne, "Electronic water mark," in *Proceedings of Digital Image Computing, Technology and Applications*, (Sydney, Australia), pp. 666–672, Dec. 1993.

[25] R. van Schyndel, A. Z. Tirkel, and C. F. Osborne, "A digital watermark," in *Proceedings of the First IEEE International Conference on Image Processing*, vol. 2, (Austin, TX), pp. 86–90, Nov. 1994.

[26] I. J. Cox, J. Killian, F. T. Leighton, and T. Shamoon, "Secure spread spectrum watermarking for multimedia," *IEEE Transactions on Image Processing*, vol. 6, pp. 1673–1687, Dec. 1997.

[27] B. Chen, *Design and Analysis of Digital Watermarking, Information Embedding, and Data Hiding Systems*. PhD thesis, Massachusetts Institute of Technology, Cambridge, MA, June 2000.

[28] E. A. Lee and D. G. Messerschmitt, *Digital Communication*. Kluwer Academic Publishers, second ed., 1994.

[29] N. S. Jayant and P. Noll, *Digital Coding of Waveforms: Principles and Applications to Speech and Video*. Prentice-Hall, 1984.

[30] R. Zamir and M. Feder, "On lattice quantization noise," *IEEE Transactions on Information Theory*, vol. 42, pp. 1152–1159, July 1996.

[31] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. John Wiley & Sons, Inc., 1991.

[32] S. I. Gel'fand and M. S. Pinsker, "Coding for channel with random parameters," *Problems of Control and Information Theory*, vol. 9, no. 1, pp. 19–31, 1980.

[33] C. Heegard and A. A. E. Gamal, "On the capacity of computer memory with defects," *IEEE Transactions on Information Theory*, vol. IT-29, pp. 731–739, Sept. 1983.

[34] M. H. M. Costa, "Writing on dirty paper," *IEEE Transactions on Information Theory*, vol. IT-29, pp. 439–441, May 1983.

[35] A. Lapidoth, "Nearest neighbor decoding for additive non-Gaussian noise channels," *IEEE Transactions on Information Theory*, vol. 42, pp. 1520–1529, Sept. 1996.

[36] F.-W. Sun and H. C. A. van Tilborg, "Approaching capacity by equiprobable signaling on the Gaussian channel," *IEEE Transactions on Information Theory*, vol. 39, pp. 1714–1716, Sept. 1993.