

Comparative performance of watermarking schemes using M-ary modulation with binary schemes employing error correction coding

Brett Bradley, Hugh Brunk
Digimarc Corporation, 19801 SW 72nd Ave., Suite 250, Tualatin, OR 97062

ABSTRACT

A common application of digital watermarking is to encode a small packet of information in an image, such as some form of identification that can be represented as a bit string. One class of digital watermarking techniques employs spread spectrum like methods where each bit is redundantly encoded throughout the image in order to mitigate bit errors. We typically require that all bits be recovered with high reliability to effectively read the watermark. In many watermarking applications, however, straightforward application of spread spectrum techniques is not enough for reliable watermark recovery. We therefore resort to additional techniques, such as error correction coding.

As proposed by M. Kutter¹, M-ary modulation is one such technique for decreasing the probability of error in watermark recovery. It¹ was shown that M-ary modulation techniques could provide performance improvement over binary modulation, but direct comparisons to systems using error correction codes were not made. In this paper we examine the comparative performance of watermarking systems using M-ary modulation and watermarking systems using binary modulation combined with various forms of error correction. We do so in a framework that addresses both computational complexity and performance issues.

Keywords: Digital watermarking, Error Correction Coding, M-ary Modulation, Digimarc

1. INTRODUCTION

In still image digital watermarking, a variety of factors can make it very difficult to recover the watermark message. Clearly, one such factor is that the watermark message power must be much less than that of the host image in order to maintain the original image fidelity. Another factor is that the watermarked image is likely to undergo a series of degradations: lossy compression, analog printing and digital recapture, and filtering are some examples. A third problem is that, in some watermarking schemes, the watermark message shares the channel with a synchronization signal, which may inherently interfere with it. Finally, some applications demand that we are able to recover the watermark message from a very small image, or a very small sub-section of the original image.

In some applications of digital watermarking, the objective is to embed N bits in a digital image that can later be recovered after various image degradations. In this paper we explore possibilities for the encoding of data that will be subject to the effects of the watermark channel. The channel, which is inextricably connected to our choice of a variant on spatial spread spectrum watermarking schemes, is very often characterized by extremely low signal to noise ratios (SNR). In order to get even a modest number of bits as throughput, one is forced to resort to somewhat complicated methods of data embedding. We will begin the paper by drawing a distinction between the watermark pixel domain, and the raw watermark bit domain. The two are related through a variant of classic direct sequence spread spectrum communications, and we will use this relationship to develop the parameters for possible error correction schemes. We will also address uncoded biorthogonal M-ary modulation and determine whether or not it is a viable alternative to traditional error correction coding techniques in the context of watermarking. A description of the experimental setup will ensue. Next, we will provide simulation results, and address issues of robustness and computational complexity. Finally, conclusions will be drawn.

2. THE WATERMARK PIXEL DOMAIN AND ITS RELATIONSHIP TO DIRECT SEQUENCE SPREAD SPECTRUM

2.1 Embedder

This paper will concern itself entirely with techniques based on Direct Sequence Spread Spectrum. We will use, for the most part, the format described by Dr. Adnan Alattar in his “Smart Images” paper². In section 3.1 of that paper, a mathematical description of how each bit is spread throughout the image is given. To summarize the relevant aspects as they pertain to this paper, a pseudo random binary key of length J is used to represent each bit. In other words, one watermark bit maps to J watermark chips. Depending upon the sign of the bit, the watermark key or its inverse is used. Each of the J components of the key is associated with a physical pixel location in the $N \times M$ image. The process is analogous to direct sequence spread spectrum communications in that a single bit is mapped to a chipping sequence of length J . In this paper we will concern ourselves with the transition from the watermark pixel domain (chips) to the Raw Bit watermark domain where we can apply error correction. The characteristics of the transition process are important because they dictate what is possible in the raw domain.

2.2 Model for noise in the watermark pixel domain

We begin our analysis by composing a model that describes how the watermark reader sees the watermarking channel. To this end, we will ignore issues related to the very important initial steps of watermark reading, which include pre-filtering, synchronization, etc. so that we may concentrate on the process of watermark chip extraction and bit recovery. Suppose that we are given, by magic, an area of $N \times M$ pixels in the watermark domain. i.e. we have synchronized to the watermark coordinate system. The total received signal at each pixel location, denoted by parameters i and k , can be described by the following equation:

$$r_i^{(k)} = w_i^{(k)} + I_i^k + g_i^k + n_i^k, \quad (1)$$

where I is the host image, w is the watermark message, g is the synchronization signal, and n is additional noise due to the “channel.” The variables i and k denote that the pixel location in question is associated with the k^{th} chip of watermark bit i . In general there are N total bits in the watermarked image, and there are J chips that belong to each bit. The total noise power at each pixel, which includes everything but the watermark itself, is typically enormous compared with the watermark chip power.

2.3 Classic correlation receiver

The successful retrieval of a watermark bit depends upon the noise present at each of the watermark chips. In classic direct sequence spread spectrum communications the relationship between chips and bits is defined in terms of signal to noise ratio and processing gain. The signal to noise ratio at the chip level is SNR_0 and the corresponding probability of error is

$$Q(\sqrt{SNR_0}) \quad (2)$$

if one is to attempt to determine the polarity of one of the chips. In such classic direct sequence spread spectrum systems, the noise at the chip level is typically close to white Gaussian, and therefore it behooves us not to make hard decisions at the chip level in order to determine the value of a particular bit. Rather, all chips are taken into account together by using a linear correlation receiver. In white Gaussian noise, such a receiver is optimum^{5,6}. The parameters of the receiver are described as follows: if K_i is the chipping sequence of length J that is used to represent the i^{th} bit, then the received signal, R_i for that bit can be described by

$$R_i = b_i K_i + N_i \quad (3)$$

R_i in equation (3) is really just the vector form of equation (1). The watermark vector is now represented by $b_i K_i$ and all noise components including the image and synchronization signal have been collapsed into the noise vector N_i . The linear correlation receiver performs the inner product between the received signal, R_i , and the chipping sequence K_i . This procedure results in the detection statistic

$$s_i = Jb_i + \varphi, \quad (4)$$

where the scalar value, $\varphi = \langle K_b, N_i \rangle$, represents the output noise of the correlator. The estimated value of the bit, b_i , is obtained by taking the polarity of s_i . The estimated polarity of a bit, after linear correlation, is much more reliable than the estimated polarity of an individual chip. The net effect of the linear correlation process is that the signal to noise ratio is increased by a factor equal to the number of chips per bit ($SNR_0 \rightarrow J * SNR_0$). The factor J , which corresponds to the number of chips, is called the processing gain. The probability of error is decreased from that at the chip level by substituting the new SNR into (2). Note: if the bit is of a particular type of error correction bits, we do not make hard decisions on the detector statistic. Rather, the decoding process uses the magnitude of the resulting information to make better decisions downstream.

2.4 A receiver with pre-filtering

Due to the host of factors listed in the introduction, straightforward application of the linear correlation receiver simply will not suffice. The probability of error is substantial enough for a watermark bit that the payload is unrecoverable even when using error correction coding. In order to mitigate the problem, a filter that takes advantage of local image pixel correlation is applied to the neighborhood of each of the chips belonging to the bit we are interested in receiving^{1,2}. By filtering the chips prior to applying the linear correlation receiver, the probability of error is reduced significantly. For the purpose of this paper, we will use a filter that has a simple response for all input values within the range of interest; i.e. any local neighborhood of pixels. In our notation N_x is an arbitrary neighborhood of image pixels centered on pixel x , $F()$ denotes the filter, and R_i' is the filtered version of R_i , which is defined in equation (3).

$$F(N_x) \in \{-1, 0, 1\} \quad (5)$$

$$R_i' = F(N_{R_i}) \quad (6)$$

The filter makes hard decisions about the polarity of the message at the chip level taking into account each chip pixel's neighborhood. We could just as well call this a binary symmetric channel with error probability, p , and nonzero erasure probability. For the present, we will ignore erasures reintroducing the state at a later time for completeness. By dropping the zero state of the output of the tri-level watermark pixel domain filter, its response to the neighborhood can be described by the binary set of outputs $\{-1, 1\}$. Given that our watermark message at the chip level is a binary antipodal signal, the probability of successfully recovering the message chip polarity is described by a Bernoulli random variable. After extracting the chip, the correlation receiver is used to obtain a soft estimate of the corresponding bit.

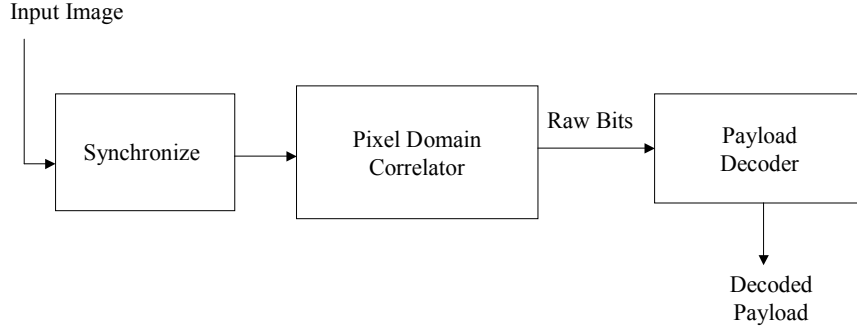
Due to the type of noise present at each watermark pixel, and the type of receiver filter used, we do not benefit from the same amount of processing gain that conventional spread spectrum systems obtain. Because we make hard decisions at the pixel or chip level prior to using the correlation receiver, we suffer an approximate 2dB penalty compared with the correlation receiver in white gaussian noise (*The watermark together with image and other noise is decidedly NOT described by a Gaussian white channel, and therefore the 2dB loss does not apply. Rather, we have found that we benefit substantially in most situations when such a filter is used*). After application of the linear correlation receiver, the probability of error in the Raw Bit watermark domain is described by a distribution that is the sum of J Bernoulli random variables. If the Bernoulli probability p is the same for all chips that comprise the bit of interest, a Binomial random variable will describe the probability of error. Specifically, we will have a watermark bit error when less than $J/2$ of the filtered chips agree with their corresponding key entries.

$$P_{err} = \sum_{k=0}^{J/2} \frac{N!}{k!(N-k)!} p^k (1-p)^{N-k} \quad (7)$$

In truth, the Bernoulli probability, p , varies with each chip's position, and hence the above equation only roughly describes the situation. If we define the mean Bernoulli probability for the chips, p_{mean} , then the probability of error in (7) is pessimistically described by p_{mean} replacing p . The reason is that the mean of the bit estimate will be the same, but the variance will be smaller than that of the Binomial distribution when p varies across the watermarked media.

The characteristics of the chip filter lend itself to easy analysis at later stages. By specifying a range of Bernoulli probabilities for making a chip error and including the dimensions of the total watermarked area in terms of pixels, we can calculate the net probability of error for the entire system (spread spectrum through error correction). However, we will find it even more convenient to break things down at a level beyond that of the chip. We will refer to this level as the Raw Bit domain.

Fig.1, Key Aspects of Watermark Detection



2.5 The Raw Bit domain

We will find it convenient to introduce an intermediate domain that lies between the watermark reader, and the final decoded payload bits called the Raw Bit domain. The relationship between the Raw Bit domain, and the other domains mentioned is shown in figure 1. The figure encapsulates all error correction coding methods that we will consider; it will need to be adjusted for M-ary modulation. In the Raw Bit domain, a bit is composed of J chips. As we shall see, the conversion to Raw Bits allows us to express our channel as a white Gaussian noise channel governed by input SNR instead of by Bernoulli probabilities. The importance of the conversion is that the gains from various coding schemes are easily determined when the Raw Bits are subject to white Gaussian noise. Similar observations were reported under the context of a different embedding scheme³. For large enough J, the modified binomial distribution, which describes our bit error statistic, approximates a Gaussian distribution with mean and variance described by the equations below.

$$\mu = 2Jp - J \quad (8)$$

$$\sigma^2 = 4Jp(1 - p) \quad (9)$$

We define the signal to noise ratio as

$$SNR = \frac{\mu^2}{\sigma^2} = J \frac{(p - \frac{1}{2})^2}{p(1 - p)} \quad (10)$$

The mean and variance are different than those of a normal Binomial distribution because the Bernoulli alphabet we use is $\{-1, 1\}$ instead of $\{0, 1\}$. The SNR is proportional to J, which makes it very easy to make simple adjustments to the SNR when J is changed. For the case of chips described by Bernoulli random variables with varying p, the resultant bit error probability is not Binomial, but it too is approximately Gaussian for large enough J. In this case, the resulting SNR will be slightly better than equation (10).

Another reason it is a good idea to define things at the Raw Bit level is that any change in the pixel filter would change the behavior and probability of error at the pixel level, but in the Raw Bit domain the statistics would remain approximately Gaussian. The net change at the Raw Bit level would be a shift in SNR, which would not require any sort of re-analysis. Generally, error corrected bits will consist of a different number of chips than J. We can accommodate such cases by adjusting the signal to noise ratio appropriately. This can be done very easily because, as described above, the processing

gain (SNR) is very simply related to the number of chips used in the bit. The overall system works as follows. Given the performance of the watermark reader at the pixel level we can define a range of Bernoulli parameters, p_{mean} , for the gamut of watermarking channel conditions within interest. We can then translate this number to a signal to noise ratio at the Raw Bit level, and see what the probability of error for the candidate error correction scheme is operating at that level of SNR and define our probability of error for a given decoding scheme. Again, for each p_{mean} within the range of interest, the corresponding SNR is a lower bound.

At this time we reintroduce our watermark chip filter as a binary symmetric channel with nonzero erasure probability. Erasures will reduce the number of watermark chips in a Raw Bit by an amount proportional to the erasure probability. If the erasure probability is p_e then the number of chips, J , per Raw Bit will be reduced, on average, to $J(1-p_e)$. The net effect is a decrease in SNR in the raw domain – typically a marginal one because p_e is small. Since we have converted all components of the watermark domain into a Raw Bit domain representation, we will drop our treatment of the former in what follows.

3. ERROR CORRECTION CODES AND M-ARY MODULATION

3.1. Convolutional codes and Reed-Solomon codes

Suppose we have a payload of L bits that we would like to embed in an image. Error correcting codes increase the payload size by adding redundant information. It is the redundant information that allows the code to do its job, to correct errors. In watermarking the cost associated with increasing the payload size is that there will be a smaller SNR per *coded* bit. Error correcting codes are useful if and only if the coding gain achieved by introducing redundancy more than makes up for the loss in SNR. The amount of redundant information, and hence the expansion in payload size, is expressed by the code's rate. For every k bits of uncoded data, n bits are embedded in the watermarked media. The code's rate is defined as $R = k/n$.

Error correction decoders operate upon the Raw Bit values (output of the correlator defined above.) Each of the Raw Bits takes on a number from $-J$ to J , where J is the length of the chipping sequence. Some decoders make hard binary decisions on each of the Raw Bits where the sign of the value is the basis for the decision. The decoder then works its magic on the resulting sequence. Generally speaking, there is some loss of information in making hard decisions on the Raw Bits prior to decoding. It is possible to implement decoders that operate on the Raw Bits themselves. Such decoders, termed soft decoders, typically perform better than their counterparts – hard decoders. However, in most cases, the extra computational cost of soft decoding is prohibitive. In this paper we consider two distinct classes of error correction codes, Reed-Solomon block codes and Convolutional codes.

Reed-Solomon codes are a particular type of block code. A typical (n,k) block coder maps blocks of k bits into blocks of n coded bits. The coder is said to be memory-less because the n coded bits depend only on the k source bits⁵. Hard decoding is almost always used in practice with block codes, and in such cases the block decoder will correct up to t bit errors. Reed-Solomon codes can be thought of as a generalization of block codes where symbols, instead of bits, make up block elements. In other words, the coder maps blocks of k symbols into blocks of n coded symbols. Each of the n coded symbols is taken from an alphabet of 2^m orthogonal symbols. For most cases that we will consider symbols correspond to groups of bits; for example, L bits will be represented by k symbols composed of n bits each. One reason Reed Solomon codes are often used in practice is because they have good minimum distance properties. The minimum distance, $d_{\text{min}} = n - k + 1$, is directly related to the code's capacity to correct symbol errors. The code is guaranteed to correct up to $t = \frac{1}{2}(d_{\text{min}} + 1)$ symbol errors. Another reason these codes are often used is that efficient hard-decoding algorithms are available⁵.

Convolutional codes are used more often than block codes because they are conceptually and practically simpler to implement, and their performance is often superior. Convolutional codes have memory; passing an L bit information sequence through a linear shift register generates encoded data. The input bits are shifted into the register, k bits at a time, and n coded bits are output to produce the n/k increase in redundancy. The maximum delay (memory) of the shift register, and hence the code, is called the constraint length. The fact that convolutional codes are linear codes with memory makes them suitable for efficient soft decoding algorithms, e.g. the Viterbi algorithm⁶.

One peculiar thing about convolutional codes, in particular, when they are used for short payload lengths, is their behavior in terms of the total decode at both low and high SNR. The success of the total decode is a much more important quantity to monitor than the bit error rate. For watermarking applications, we typically will accept the result of the decode only when

that result is error free. Any number of bit errors is too many. For the uncoded case, the probability of correct payload retrieval is directly related to the probability of a bit error. For example, we know that in uncoded signaling the probability of no bit errors (complete message retrieval) is $(1-p_b)^N$, where N is the number of bits. However, for convolutional codes, the probability of a bit error does not tie in neatly to the probability of correctly decoding. At low SNR the probability of a bit error will be high, but the probability of decoding the payload correctly will be much higher than one might expect. Errors are not independent; they almost always occur in groups. Another curious behavior occurs at high SNR. To see this, consider a Viterbi decoder operating on the k^{th} bit of a payload of arbitrary length. The coded version of the k^{th} bit is represented by r encoded bits beginning at index $r(k-1) + 1$ of the coded sequence. It is well known that a final decision on any given payload bit should be made much later than the above correspondence would indicate⁶. Specifically, if the group of r coded bits beginning at index $r(k-1) + 1$ is the latest to enter the decoder, then decisions on the $k - jD^{\text{th}}$ payload bit will be near optimal, where D is the constraint length of the code and j is an integer – typically j equal to 4 or 5 will suffice⁶. In terms of decoding a watermark message at high SNR, the probability of a bit error should be very low. However, because we have a small payload, decisions regarding the final bits in the payload must be made prematurely, effectively increasing the error probability. Decisions are made prematurely because there is no decoding delay available.

3.2 Concatenated Codes

In some applications convolutional codes are combined with Reed Solomon codes to form what is called a *concatenated code*. A Reed-Solomon encoder is used on the payload bit string; it is referred to as the outer code. The Reed-Solomon encoded data is then itself encoded using a convolutional code; the inner code. The method has been used for deep space communications (definitely a low SNR environment!). The idea is to exploit the strengths and cater to the weaknesses of each type of code. Reed-Solomon codes offer performance superior to that of convolutional codes, but only for reasonably good channels. For very poor channels convolutional codes work better. Using the convolutional code to “clean up” the raw channel allows the Reed-Solomon code to work where it can perform its best. Another synergy between the two codes is that the convolutional decoder tends to produce bursty errors, which is what the Reed-Solomon code is best at correcting.

3.3 M-ary methods

In his paper, M. Kutter introduced a signaling scheme for watermarking that was a generalization of binary signaling at the bit level¹. This scheme, called biorthogonal M-ary signaling, maps groups of $\log_2(M)$ bits to one of M symbols. The benefit of doing so is that under certain conditions the probability of symbol error becomes arbitrarily small as M goes to infinity. Specifically, if the energy per bit is greater than -1.6dB, the Shannon limit, the above statement holds⁵. The larger the value of M the fewer symbols we are required to hide in an image to convey the same number of bits. The fewer symbols we have to hide in an image, the more locations we can use per symbol¹. In other words, the symbol energy increases for increasing M (Although the energy per bit must remain fixed).

We will find it useful to relate the net symbol SNR to Raw Bit SNR, developed earlier. One repetition of a symbol requires $M/2$ chips or watermark pixel locations. Just as there are many chips per watermark bit in the binary case, each symbol is typically repeated many times in the watermark domain to increase its aggregate SNR. We will use an example to ease the process of development. Suppose we have an $L \times L$ pixel area to be watermarked; call this the “*TotalPixelArea*.” Further suppose that we would like to embed N bits. In the binary case it is easy to see that each bit will get $(TotalPixelArea / N)$ chip repetitions. The SNR per bit is related to the Raw Bit SNR by a multiplication factor,

$$SNR_b = aSNR_{raw}, \quad (11)$$

where $a = (TotalPixelArea / N) / J$. For example, if $TotalPixelArea = 128 \times 128$, $N = 64$ bits, and $J = 32$, then $a = 8$, a shift of +9dB. The more general case is described by the equations below.

$$a = \frac{TotalPixelArea}{\frac{M}{2} \left(\frac{N}{\log_2(M)} \right) J} \quad (12)$$

In equation (12), $M/2$ is the number of pixels required per symbol repetition, $N/\log_2(M)$ is the total number of symbols needed to represent the required number of embedded bits, and J is the number of pixels (chips) per Raw Bit.

$$SNR_{Symb} = aSNR_{Raw} \frac{M}{2} \quad (13)$$

Results using the above equations are tabulated for some sample parameters, below.

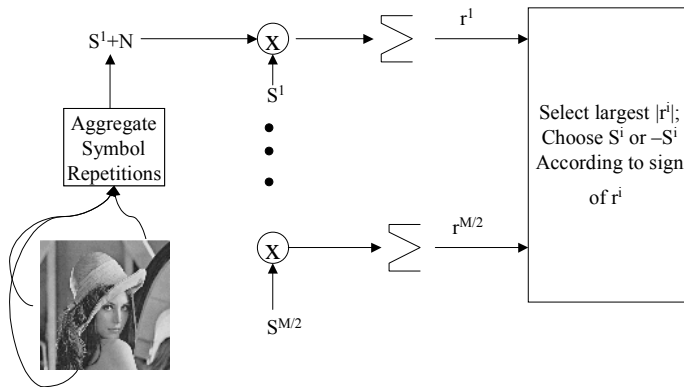
Table 1
N = 60-64bits, rawbitSnr = -2dB, chips/rawBit = 32, Symbol Size = M/2

M	Num Symbols	Symbol Reps	Symbol SNR(dB)
2	64	256	7
8	21	195	11
16	16	128	13
32	12	85	14
64	10	51	15

Here we summarize the process of embedding and detecting M-ary symbols.

M distinct signals are produced by computing an M/2 order Hadamard matrix, $H_{M/2}$. To generate the additional M/2 biorthogonal entries, we append the sign-reversed version of $H_{M/2}$ to itself¹. The detection process works by garnering all available symbol repetitions from the watermark pixel domain to form an aggregate noisy estimate of the symbol, $S^i + N$. We would use, for example, the same chip filtered described by equations (5, 6) to prefilter each contribution to a symbol element. M/2 correlators of length M/2 are used in the detection process, which is described in the figure below.

Fig.2 M-ary Symbol Detection



The “SymbolSNR” is the SNR of each symbol element prior to match filtering with the bank of M/2 correlation receivers, $R = S + N$ and $snr = s_i^2/\text{var}(N_i)$. Selecting the symbol is a two-step process: It requires largest magnitude determination followed by sign determination.

Making an error at the symbol level can be accomplished by confusing the correct symbol with its antipodal version, or mistaking one of the other possible M/2 symbols for the correct one. Errors almost always occur due to the second condition mentioned because the distance in signal space between orthogonal signals is half that of antipodal signals^{1,5}. The probability of error for biorthogonal M-ary signals, P_M , is given in Proakis’ book as a function of symbol signal to noise ration, SNR_S ⁵. We reproduce the equation here.

$$P_M = 1 - \frac{1}{\sqrt{2\pi}} \int_{-\sqrt{2SNR_S}}^{\infty} \left(\frac{1}{\sqrt{2\pi}} \int_{-(v+\sqrt{2SNR_S})}^{v+\sqrt{2SNR_S}} e^{-x^2/2} dx \right)^{M/2-1} e^{-v^2/2} dv \quad (14)$$

The equation may be evaluated numerically for different values of M and SNR_s . We are more interested in the probability of not decoding the entire watermark, P_W . Every symbol is represented independently in the image and each has the same SNR_s (the reason for equal SNR_s is a result of an interleaving procedure²). We require the successful demodulation of all symbols that make up the bit payload to be successful. The equation governing our success is

$$P_W = 1 - (1 - P_M)^{N/\log_2(M)} \quad (15)$$

3.4 Combining M-ary signaling with error-correction coding

So far we have treated error correction coding and M-ary modulation as if they are mutually exclusive methods. Yet, since we have entertained the idea of concatenated codes, it seems reasonable to think of possibilities employing M-ary modulation with error correction. We do not treat any such techniques in the simulation, but for the sake of completeness we develop the basic premise. We consider two examples of hybrid M-ary and error correction techniques.

Reed Solomon codes are particularly suited to M-ary modulation when the full alphabet of 2^k symbols is used⁵. To embed L bits with standard M-ary modulation would require K , M-ary symbols, where $K \log_2(M)$ is equal to L . Using Reed Solomon codes, the K , M-ary symbols are converted to a larger number, N , for embedding. Since more symbols are required to embed the same number of bits, the SNR per symbol will decrease. The probability of symbol error at the output of the decoder is upper-bounded by

$$P_{es} = \frac{1}{N} \sum_{i=t+1}^N i(N, i) P_M^i (1 - P_M)^{N-i} \quad (16)$$

where P_M is given in equation (14), but the quantity must be adjusted for the loss in SNR.

$$SNR_{dB} \rightarrow SNR_{dB} - 10 \log_{10}(N/K) \quad (17)$$

The probability of decoding the payload is given by (15). As we shall see, at low SNR the uncoded case is superior. The situation is reversed for higher SNR.

In limited bandwidth situations, trellis coding is often used to enhance performance. Trellis coding consists of first convolutionally encoding the information bits, and then using line coding to select a symbol to transmit from the available constellation. For example, in digital communications, one might use a rate 1/3 convolutional code, and then map the triplet of bits to a symbol in an 8-PSK constellation. On the decoding side a Viterbi soft decoder would map the symbols to received bits. The symmetry in the constellation makes for a large minimum distance error event, and hence the gain over an uncoded system is substantial.

It is possible to construct a quasi-trellis coding scheme using convolutional codes and M-ary modulation. In the M-ary orthogonal scheme, however, all symbol errors barring the antipodal symbol are equally likely, and therefore the minimum distance error event will be closer than in a more conventional trellis-coding scheme. Nevertheless, let us proceed with an illustration of such a scheme. Suppose we apply 8-ary modulation to rate 1/3 convolutionally encoded bits. Coded bits, in groups of three, are mapped to a unique symbol. The detector will perform M-ary detection as normal using the bank of correlation receivers. Instead of selecting the symbol where the correlation is maximum and making hard decisions on the coded bits, the decoder will use the strength of that correlation value with reference to the next highest in order to assign a soft reliability metric to the chosen symbol. The resulting soft values are fed to a Viterbi soft decoder.

4. ERROR CORRECTION CODING AND UNCODED M-ARY SIMULATION RESULTS

In the discussion that follows, the word “protocol” is synonymous with a data-encoding scheme. In fact, it supersedes it since M-ary modulation is not truly an error-correction encoding scheme. Table 2 shows a listing of the various protocols that have been considered in this study. Protocols with convolutional codes, Reed-Solomon codes, and also concatenated Reed-Solomon and convolutional codes have been examined. In addition, we simulated several different levels of M-ary signaling. For the concatenated codes, an effort was made to look at a variety of code rate allocations between the Reed-Solomon and convolutional code. For all types of protocols, the number of chips, or equivalently Raw Bits, per protocol bit was also varied to identify the best compromise between coding and repetition. In all protocols the total number of chips used is 128x128. All convolutional codes used were memory 8 codes (with 128 states), except a single protocol where a memory 7

code was used to see if it paradoxically could perform better due to the small decoding trellis. Most protocols were designed to allow approximately 64 bits for the payload; small deviations sometimes had to be made to accommodate the available Reed-Solomon codes, and various M-ary levels.

In order to measure the performance of each of the protocols we used simulation. Since only a small number of variables need to be in the model, the simulation task is tractable. Refer to figure 1. It shows a diagram of the reader broken into three components: synchronization, watermark reading, and payload decoding. The only input to the payload-decoding block is the sequence of raw values, one for each bit or symbol of the protocol. The factor that determines the performance of the payload-decoding block is the statistical relationship between the original protocol bits and the Raw Bits that are input to the decoding block. We have chosen a range of SNR that truly stresses the various methods considered.

Before going into detail regarding simulation results, we perform a cursory comparison with expected theoretical results involving the bit error rate, which is a conventional measure of reliability in digital communications, and point out a few problems with relying entirely on the measurement. In this example the payload is 56 bits, and three different methods are compared: binary modulation, 32-ary biorthogonal modulation, and binary modulation with rate 1/3 convolutional codes employing soft Viterbi decoding. The three methods are compared over the same range of Raw Bit SNR, where a Raw Bit is comprised of 32 chips. For binary modulation, the actual SNR per bit is the Raw Bit SNR + 9.47dB due to the extra space available per bit for processing gain. The rate 1/3 convolution coding method gets half that, an additional 4.5db of processing gain per coded bit; the other 4.5dB was used for coding gain. The SNR per bit of the M-ary scheme is essentially the same as that for the binary scheme, but the SNR per symbol is $\log_2(M)$ greater.

Figure 3a, Bit Error Rate from Simulation

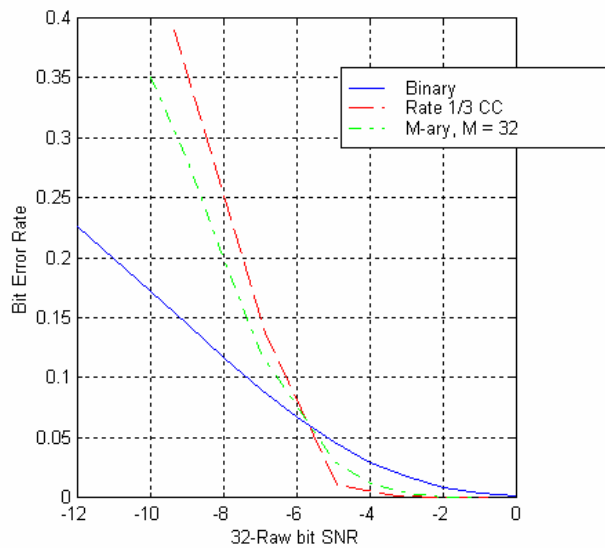
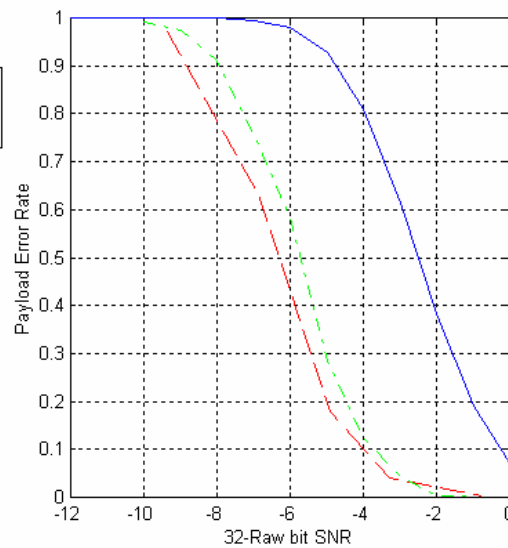


Figure 3b, Payload Error Rate



Comparing the binary method with that using convolutional codes, we see that the former has a lower bit error rate until ~ -6 dB, after which it is supplanted by the coding scheme. The M-ary symbol error rate is higher than the binary bit error rate until ~ -4 dB. The M-ary bit error rate is always lower than its symbol error rate, since each symbol consists of a group of bits. It is approximated by $\frac{1}{2}$ the symbol error rate for large M^5 .

Viewing Figure 3b, which depicts the probability of a correct decode for each of the methods, we notice that for low Raw Bit SNR the probability of error for the binary scheme is enormous. For the range of SNR where it appeared better than the convolutional coding technique in terms of bit error rate, the method is essentially useless if there is no tolerance for bit errors – a stipulation that seems reasonable when the payload is small. Hernandez, et al. have pointed out this problem³.

More interesting, and less obvious, is that the probability of a payload error for the convolutional coding method is always less than that of the binary method, even in the region where bit error probability would suggest otherwise. In fact the

probability of a decode error is arguably reasonable, for some applications, when the bit error rate is still marginally high. For example at -6dB the bit error rate is roughly 0.1 and the corresponding probability of decoding the entire payload correctly is as high as 0.6. The reason for this phenomenon is that, using our example, 60% of the noisy payload data will succeed in being decoded as a whole. The other 40% will be characterized, in general, by bursts of errors. This all or nothing result is a consequence of considering the data in one lumped sum. A conclusion that can be drawn from this is that bit error rate is not a very interesting statistic when applied to convolutional coding of small payloads. Another observation about convolutional codes is that the probability of error for high SNR is worse than experience with such a technique would indicate. The reason for this is that Viterbi soft decoders need a delay of about 4-5 times the constraint length prior to making bit decisions in order to operate at their optimum rate⁶. We, however, are forced to make immediate decisions for data at the tail of the bit stream because of the very small payload size, a fact that significantly increases the probability of error.

The performance of each of the protocols was simulated for SNR ranging from -7dB to -2dB in steps of 1 dB. The SNR is given in terms of a Raw Bit with 32 chips; for simulation purposes the SNR was adjusted as appropriate for other number of chips/bit. Each data point was obtained from simulating 8000 randomly chosen payloads and passing them through the AWGN channel. For reasons mentioned in the preceding paragraphs, we quote the performance of the total decode and ignore the bit error rate. Results showing the probability of correct payload decoding are shown in Table 3.

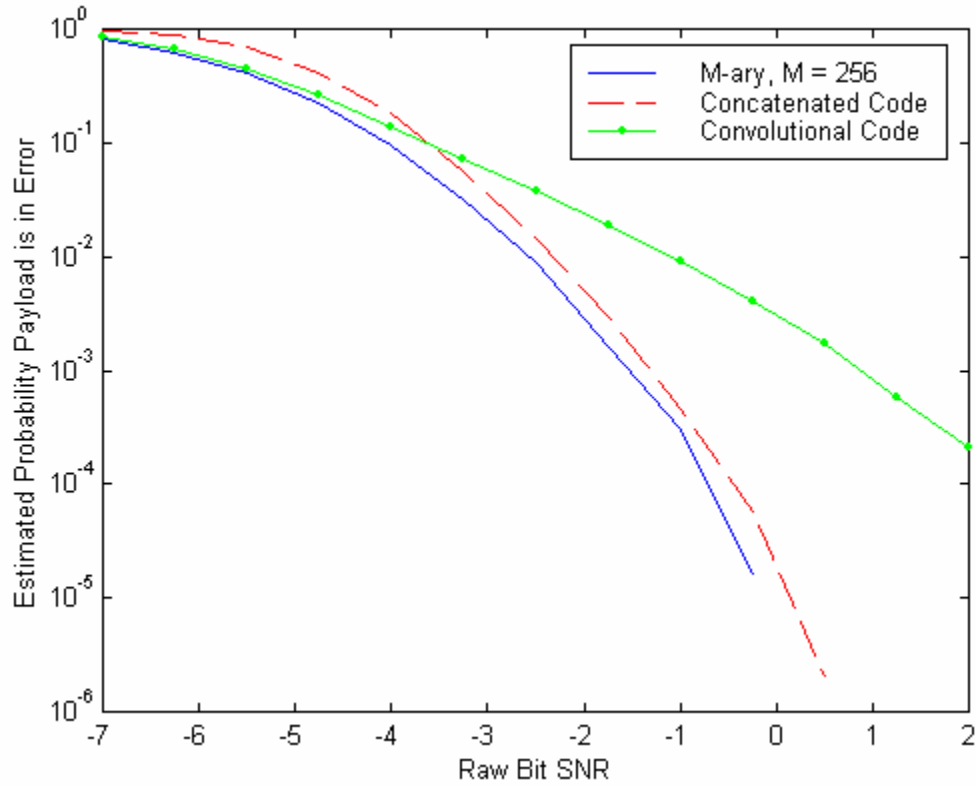
A few very general remarks can be made concerning the various classes of protocols. Protocols employing pure Reed-Solomon block codes did not perform well over the chosen range of SNR. Our result is consistent with that posted by others who have tried BCH codes, a binary block code³. The two protocols that used pure convolutional codes of differing rates performed well, better than those that used Reed-Solomon codes for all SNR in the simulation. Several of the concatenated codes achieved very good rates of decoding the payload correctly at high SNR. The higher order M-ary schemes did reasonably well at all SNR, and they were among the best at higher SNR in particular.

Protocols that use convolutional codes as their primary ingredient are not a bad choice for watermarking. Comparing Protocols 29 and 30, the two that used convolutional codes of rates $1/3$ and $1/4$ respectively, we can say that the overall performance is about the same. The higher rate ($1/3$) is slightly better at low SNR, but at a higher and more useable SNR, the rate $1/4$ code might be marginally superior. Protocols 11, 12, and 14-16 used concatenated codes and they performed at least as well as Protocols 29 and 30 at high SNR. The reason for this is that at approximately -4dB we begin to see a transition, whereas the error characteristics of protocols 29 and 30 change from error patterns characterized by bursts that plague the entire payload to shorter error events at the tail of a payload. The short error events, which are a result of early trellis truncation in the Viterbi Algorithm, are corrected if an outer code is used, provided we can accommodate the loss in SNR required by the extra coding. In light of this discussion, one might argue that a concatenated code operating upon the entire payload is overkill. If instead of using a traditional concatenated code, we apply a BCH code to the tail of a convolutionally encoded payload we would protect the most error prone part of the payload. The impact in terms of SNR would be less than that of the concatenated code because we would be required to embed fewer coded bits.

Protocols that use pure M-ary modulation, with large M, are also a reasonable choice for watermarking under the simulated conditions. Protocol 36, which has $M = 256$, is arguably the best of all protocols for the simulated range of SNR. It is, of course, possible to consider values of M larger than 256 if processing time is of a lesser concern. The appendix shows that Protocol 36 requires almost 5 times the number of operations that Protocol 29 does. In the appendix, however, we did not consider the possibility of using a Fast Walsh Transform⁷, which can significantly curtail the number of required operations. The Fast Walsh Transform is possible when a Walsh-Hadamard set of basis functions is used – the case we simulated here.

Having identified some promising protocols, further simulation was performed. Protocols #11, #30, and #36 were simulated over a range from -10dB to -1dB in steps of $.75\text{dB}$. Each data point is the result of 500000 simulation trials. In figure 5, we plot the estimated probability of not correctly decoding the payload in its entirety. Comparing Protocols #11 and #30, the two based on error-correction schemes, we see that which of the two protocols is best depends on the SNR. For poor effective channels, the protocol with a convolutional code outperforms the protocol with a concatenated code. For better channels the situation is reversed. Protocol #36, which is an M-ary method, performs better than the two error-correction coding techniques, throughout the simulated range of SNR. We conclude that M-ary modulation is certainly a suitable method for watermarking.

Figure 5, Detailed Comparison of Promising Protocols



5. CONCLUSIONS

The class of digital image watermarking techniques based on spread spectrum methodology is often characterized by very low SNR. Methods have been proposed that either enhance the basic technique by applying error correction codes, or generalize the spatial spread spectrum principle using M-ary modulation to obtain better results^{3,4,1}. Comparisons between different error correction schemes have been made previously³, but this paper expands on the theme by testing multiple code rates, exploring parameter allocation, and introducing concatenated codes. Furthermore, M-ary schemes are compared against the various error correction coding methods. In the context of our comparison, we have chosen to introduce a framework called the Raw Bit domain that allows us to distill down the elements of competing schemes so that we can make direct comparisons given channel conditions (SNR), and payload size.

As a result of our simulations, we believe that convolutional codes alone or concatenated with Reed-Solomon codes can be a good choice to increase payload robustness. M-ary methods, for M greater than or equal to 256, are at least as good of a choice, considerations of computational complexity aside. Soft Convolutional decoding techniques suffer a bit more than one would expect at high SNR when the payload is small, a fact that makes them inferior to higher order M-ary techniques. Some types of concatenated codes are able to mitigate this problem at high enough SNR. In watermarking the channel noise characteristics are likely to vary substantially, a fact that should be kept in mind when choosing a coding scheme.

Table 2a, Protocol Descriptions and Parameters (Error Correction Coding)

Protocol #	# Payload Bits	#chips/protocol bits	RS symbol	RS N	RS K	CC rate	CC memory
1	60	32	6	42	10	0.5	8
2	63	32	7	36	9	0.5	8
3	60	32	6	28	10	0.33	8
4	63	32	7	24	9	0.33	8
5	60	32	5	25	12	0.25	8
6	60	32	6	21	10	0.25	8
7	63	32	7	19	9	0.25	8
8	60	64	5	25	12	0.5	8
9	60	64	6	21	10	0.5	8
10	63	64	7	18	9	0.5	8
11	60	64	5	17	12	0.33	8
12	60	64	6	14	10	0.33	8
13	63	64	7	12	0	0.33	8
14	60	96	5	17	12	0.5	8
15	60	96	6	14	10	0.5	8
16	63	96	7	12	9	0.5	8
17	63	32	7	73	9	-	-
18	64	32	8	64	8	-	-
19	60	64	6	42	10	-	-
20	63	64	7	36	9	-	-
21	64	64	8	32	8	-	-
22	60	96	6	28	10	-	-
23	63	96	7	24	9	-	-
24	64	96	8	21	8	-	-
25	60	128	5	25	12	-	-
26	60	128	6	21	10	-	-
27	63	128	7	18	9	-	-
28	64	128	8	16	8	-	-
29	64	85	-	-	-	0.33	8
30	64	64	-	-	-	0.25	8
31	60	64	5	17	12	0.33	7

Table 2b, Protocol Descriptions and Parameters (M-ary)

Protocol #	Number of Bits	M-ary Level
32	64	2
33	64	16
34	66	64
35	63	128
36	64	256

Table 3, Comparative Rates of Correct Payload Decoding

Protocol #	-7dB	-6dB	-5dB	-4dB	-3dB	-2dB
1	0	0	0	0	0	0
2	0	0	0	0	0	0
3	0	0	0	.03	.25	.74
4	0	0	0	.02	.2	.67
5	0	0	.04	.28	.69	.95
6	0	0	.04	.25	.68	.95
7	0	0	.03	.21	.64	.94
8	0	0	.01	.12	.47	.85
9	0	0	.01	.12	.46	.86
10	0	0	.01	.1	.44	.84
11	.03	.17	.5	.82	.96	>.99
12	0	0	.53	.85	.98	>.99
13	.02	.08	.23	.37	.42	.42
14	0	.08	.31	.66	.9	.98
15	0	.1	.34	.71	.93	.99
16	0	0	0	.66	.91	.98
17	0	0	0	0	0	0
18	0	0	0	0	0	0
19	0	0	0	0	0	.01
20	0	0	0	0	0	0
21	0	0	0	0	0	0
22	0	0	0	0	.04	.23
23	0	0	0	0	0	.07
24	0	0	0	0	0	0
25	0	0	0	.06	.23	.57
26	0	0	0	.02	.12	.4
27	0	0	0	.01	.06	.24
28	0	0	0	0	0	0
29	.17	.43	.7	.87	.94	.975
30	.16	.4	.69	.86	.94	.98
31	.03	.17	.46	.8	.96	.99
32	.001	.006	.02	.07	.22	.44
33	.07	.18	.34	.62	.78	.94
34	.1	.27	.5	.74	.92	.986
35	.16	.39	.66	.88	.97	.994
36	.21	.45	.70	.92	.98	.999

APPENDIX

In the appendix we address issues of computational complexity for each method.

- Convolutional Coding using the Viterbi Algorithm– complexity increases linearly with payload size. Refer to references^{5,6} for implementation details. K is the constraint length and L is the introduced coding redundancy. The following is for rate 1/L convolutional codes.
 - 2^L Euclidean distance calculations per payload bit.
 - 3L operations per Euclidean distance calculation. 3L or 3L -1(L adds + L mults + L-1 adds).
 - 2×2^k adds per payload bit.
 - 2^k Compares per payload bit.

$$\text{NumOps} = N \times (2^L \times 3L + 2 \times 2^k + 2^k)$$

Example: Protocol number 29 has N = 64, L = 3, and K = 8. The number of operations is on the order of 54000.

2. Reed Solomon Codes-
 - Varies greatly depending upon algorithm used. In general, the complexity is less than that of convolutional codes.
3. M-ary Signaling-
 - $M/2$ adds plus $M/2$ multiplies per correlation with 1 of $M/2$ symbols.
 - $M/2$ symbols to be correlated against.
 - $M/2 + 1$ operations to determine the embedded symbol after correlation (choose the symbol based on the sign of the maximum correlation magnitude).
 - $N/\log_2(M)$ symbols required.

$\text{NumOps}_{\text{M-ary}} = N/\log_2(M) * (M \times M/2 + M/2 + 1)$
 Example: Protocol number 36 has $M = 256$, and $N = 64$. The number of operations required is on the order of 263,000, almost 5 times that of the Convolutional code example.

REFERENCES

1. M. Kutter, "Performance Improvements of Spread Spectrum Based Image Watermarking Schemes Through M-ary Modulation", *Preliminary Proceedings of the Third International Information Hiding Workshop*, pp. 245-260, Dresden, 1999.
2. A. Alattar, "Smart Images", *Proceedings of SPIE Vol. 3971*, pp. ?, San Jose, 2000.
3. J.R. Hernández, J. Delaigle, B. Macq, "Improving Data Hiding by Using Convolutional Codes and Soft-Decision Decoding," *Proceedings of SPIE*, Ping Wah Wong, Edward J. Delp, Editors, Vol. 3971, pp. ?, San Jose, 2000.
4. S. Baudry, P. Nguyen, H. Maître, "Channel Coding in Video Watermarking: Use of Soft Decoding to Improve the Watermarking Retrieval," *Proceeding ICIP 2000*, Vancouver, Canada, 2000.
5. J.G. Proakis, *Digital Communications, 3rd Edition*, Chapters 5, 8, McGraw-Hill, 1995.
6. E.A. Lee, D.G. Messerschmitt, *Digital Communication, 2nd Edition*, Chapters 9, 13, 14, Kluwer Academic Publishers, MA, 1994.
7. R.C. Gonzalez, R.E. Woods, *Digital Image Processing*, Chapter 3, Addison-Wesley, 1992.