

On Steganographic Embedding Efficiency

Jessica Fridrich¹, Petr Lisoněk², David Soukal¹

¹ Binghamton University, Binghamton, NY 13902–6000

² Simon Fraser University, Burnaby, British Columbia, V5A 1S6, Canada
{fridrich, dsoukal1}@binghamton.edu, plisonek@cecm.sfu.ca

Abstract. In this paper, we study embedding efficiency, which is an important attribute of steganographic schemes directly influencing their security. It is defined as the expected number of embedded random message bits per one embedding change. Constraining ourselves to embedding realized using linear covering codes (so called matrix embedding), we show that the quantity that determines embedding efficiency is not the covering radius but the average distance to code. We demonstrate that for linear codes of fixed block length and dimension, the highest embedding efficiency (the smallest average distance to code) is not necessarily achieved using codes with the smallest covering radius. Nevertheless, we prove that with increasing code length and fixed rate (i.e., fixed relative message length), the relative average distance to code and the relative covering radius coincide. Finally, we describe several specific examples of q -ary linear codes with q matched to the embedding operation and experimentally demonstrate the improvement in steganographic security when incorporating the coding methods to digital image steganography.

1 Introduction

Steganography is the art of undetectable communication. It was originally formalized by Simmons [1] as the prisoners’ problem. Alice and Bob are prisoners in separate cells who want to develop an escape plan. Their communication is monitored by a warden. Alice and Bob resort to steganography and hide the details of the escape plot in cover objects, such as digital images, by slightly modifying them. Their goal is to not raise the warden’s suspicion. In the simplest case, the warden is passive in that he just observes the traffic without modifying the messages in any way.

The main requirement of any steganographic technique is undetectability—the warden should not be able to distinguish between cover and stego objects (cover embedded with data) with success better than random guessing. A formal definition of steganographic security was given by Cachin [2]. The detectability of data hidden in a stego object is influenced by many factors, such as the choice of the cover object, the selection rule used to identify individual elements of the cover that could be modified during embedding, the type of embedding operation that modifies the cover elements, and the number of embedding changes (directly related to the secret message length). Assuming two embedding methods share

the same source of cover objects, the same selection rule and embedding operation, the one that introduces fewer embedding changes will be less detectable as it decreases the chance that any statistics used by the warden will be sufficiently disturbed to mount a successful steganalysis attack. The expected number of random message bits embedded per one embedding change is called embedding efficiency. This concept has been introduced by Westfeld [3] and has since been accepted as an important attribute of steganographic schemes [4, 5].

In 1998, Crandall [6] and Bierbrauer [7, page 195–197] showed that embedding efficiency of steganographic schemes can be improved by applying covering codes to the embedding process. This fact has been later independently rediscovered by van Dijk et al. [8] and Galland et al. [9]. In particular, a linear code can be used to construct an embedding scheme³ whose embedding capacity is the code redundancy, while the covering radius corresponds to the maximal number of embedding changes necessary for embedding any message.

In this paper, we first show that the expected number of embedding changes, which is directly related to the concept of embedding efficiency as used in current steganographic literature, corresponds to the average distance to code rather than the covering radius. Moreover, we show that in the class of linear codes of fixed length and dimension the highest embedding efficiency may not always be attained for a code with the smallest covering radius. However, with increasing code length and fixed rate (i.e., fixed relative message length), the relative covering radius and the relative distance to code asymptotically coincide.

In Section 2, we review selected known facts about embedding schemes realized using q -ary linear codes and state bounds on embedding efficiency. In Section 3, we study the properties of the average distance to code. Examples of specific coding schemes that can substantially improve the embedding efficiency of steganographic schemes are given in Section 4, where we experimentally demonstrate the benefit of using the proposed coding techniques for steganography. The paper is concluded in Section 5.

2 Covering codes in steganography

In this section, we briefly review some known results about steganographic schemes and covering codes including bounds on achievable embedding efficiency. We do so for a rather general definition of an embedding scheme in which message symbols from some finite field (rather than bits) are embedded at each pixel. The reason for this more general approach will become clear in Section 4 when we discuss the importance of ternary codes for steganography. Throughout the text, boldface symbols stand for vectors or matrices and the calligraphic font is used for sets. Italicized text highlights definitions of new concepts.

We will assume that the cover image \mathbf{X} is an element of \mathcal{G}^n , where \mathcal{G} is the set of all possible pixel values. For example, in steganography using 8-bit grayscale digital images, \mathcal{G} is the set of all integers in the range $[0, 255]$ and n is the number

³ In steganographic literature, such embedding schemes realized using linear codes are called matrix embedding [3, 6, 10].

of pixels. Data embedding consists of modifying the values of selected pixels so that the modified (stego) image \mathbf{Y} conveys the desired secret message. The impact of embedding is captured by a distortion metric $D : \mathcal{G}^n \times \mathcal{G}^n \rightarrow [0, \infty)$.

We further assume that there is a symbol-assignment function $s : \mathcal{G} \rightarrow \mathbb{F}_q$ that assigns an element of a finite field⁴ \mathbb{F}_q to each possible pixel value. The most common symbol-assignment function used in steganography is the least significant bit (LSB) of pixel values

$$s(i) = i \bmod 2. \quad (1)$$

Examples of other symbol-assignment functions are given in Section 4.

Writing the pixels of image \mathbf{X} as a one-dimensional vector, its vector of symbols $s(\mathbf{X}) = \mathbf{x} \in \mathbb{F}_q^n$ is obtained by applying s to each element. Everywhere in this paper, we measure the impact of embedding in the symbol space \mathbb{F}_q^n using the Hamming distance $d : \mathbb{F}_q^n \times \mathbb{F}_q^n \rightarrow \{0, 1, \dots, n\}$ between the corresponding symbol vectors, which is the number of embedding changes

$$D(\mathbf{X}, \mathbf{Y}) = d(s(\mathbf{X}), s(\mathbf{Y})) \text{ for all } \mathbf{X}, \mathbf{Y} \in \mathcal{G}^n. \quad (2)$$

Let \mathcal{M} be the set of all messages that can be communicated. An embedding scheme with a distortion bound R is a pair of embedding and extraction functions Emb and Ext ,

$$Emb : \mathbb{F}_q^n \times \mathcal{M} \rightarrow \mathbb{F}_q^n \text{ and } Ext : \mathbb{F}_q^n \rightarrow \mathcal{M}, \quad (3)$$

$$d(\mathbf{x}, Emb(\mathbf{x}, \mathbf{m})) \leq R \text{ for all } \mathbf{m} \in \mathcal{M} \text{ and all } \mathbf{x} \in \mathbb{F}_q^n, \quad (4)$$

such that for all messages $\mathbf{m} \in \mathcal{M}$ and all $\mathbf{x} \in \mathbb{F}_q^n$, $Ext(Emb(\mathbf{x}, \mathbf{m})) = \mathbf{m}$. In other words, (3) means that we can embed any message from \mathcal{M} in any \mathbf{x} and (4) states that we can do it by imposing at most R changes.

The value $h = \log_2 |\mathcal{M}|$ is called the embedding capacity of the scheme (in bits) and $\alpha = h/n$ the relative embedding capacity (or relative payload). We have an obvious upper bound

$$|\mathcal{M}| \leq q^n \text{ or } \alpha \leq \log_2 q. \quad (5)$$

We further define $\underline{e} = \frac{h}{R}$ as the lower embedding efficiency and $e = \frac{h}{R_a}$ as the embedding efficiency, where R_a is the expected number of changes over uniformly distributed cover objects $\mathbf{x} \in \mathbb{F}_q^n$ and messages $\mathbf{m} \in \mathcal{M}$. Note that since R is the upper bound on the number of embedding changes, for any embedding scheme $\underline{e} \leq e$.

We next review some known facts about embedding schemes and covering codes and state a bound on embedding efficiency. More details and proofs can be found in [9, 12, 13]. Throughout this article, we will use some standard concepts and results from Coding Theory that can be found for example in [11]. Unless

⁴ Here, q is a prime power. For background on finite fields, see for example Chapters 3 and 4 in [11].

stated otherwise, all codes considered in this article are linear codes, and we use the notation “[n, k, d] code” for a k -dimensional linear code with block length n and minimal distance d . If the minimal distance d is not important for our considerations, we may omit it and only speak of an [n, k] code. We note that the covering radius R of a q -ary code \mathcal{C} is defined as

$$R = \max_{\mathbf{x} \in \mathbb{F}_q^n} d(\mathbf{x}, \mathcal{C}), \quad (6)$$

where $d(\mathbf{x}, \mathcal{C}) = \min_{\mathbf{c} \in \mathcal{C}} d(\mathbf{x}, \mathbf{c})$ is the distance between \mathbf{x} and the code \mathcal{C} . An R -covering of \mathbb{F}_q^n is any subset \mathcal{C} of \mathbb{F}_q^n such that $\bigcup_{\mathbf{x} \in \mathcal{C}} \mathcal{B}(\mathbf{x}, R) = \mathbb{F}_q^n$, where $\mathcal{B}(\mathbf{x}, R)$ is the ball with center \mathbf{x} and radius R .

We now state and prove the matrix embedding theorem. It gives a recipe how to use an [n, k] code to communicate $n - k$ symbols using at most R changes in n pixels. Examples of specific matrix embedding schemes for binary and ternary codes are given in Section 4.

Theorem 1. (Matrix embedding) Let \mathcal{C} be an [n, k] code with a parity check matrix \mathbf{H} and covering radius R . The embedding scheme below can communicate $n - k$ symbols in n pixels with pixel symbols \mathbf{x} using at most R changes:

$$\begin{aligned} Emb(\mathbf{x}, \mathbf{m}) &= \mathbf{x} + \mathbf{e}_L = \mathbf{y}, \\ Ext(\mathbf{y}) &= \mathbf{H}\mathbf{y}, \end{aligned}$$

where $\mathbf{m} \in \mathbb{F}_q^{n-k}$ is a sequence of $n - k$ message symbols and \mathbf{e}_L is a coset leader of the coset $\mathcal{C}(\mathbf{m} - \mathbf{H}\mathbf{x})$ for the syndrome $\mathbf{m} - \mathbf{H}\mathbf{x}$.

Proof. Since \mathcal{C} has covering radius R , we know that $d(\mathbf{x}, \mathbf{y}) = w(\mathbf{e}_L) \leq R$, which proves that the embedding scheme has (a tight) distortion bound R . To prove that $Ext(Emb(\mathbf{x}, \mathbf{m})) = \mathbf{m}$, note that $Ext(Emb(\mathbf{x}, \mathbf{m})) = \mathbf{H}\mathbf{y} = \mathbf{H}\mathbf{x} + \mathbf{H}\mathbf{e}_L = \mathbf{H}\mathbf{x} + \mathbf{m} - \mathbf{H}\mathbf{x} = \mathbf{m}$.

Because there are $\sum_{i=0}^R \binom{n}{i} (q-1)^i$ ways in which one can make up to R changes in n pixels, we have

$$h = \log_2 |\mathcal{M}| \leq \log_2 \sum_{i=0}^R \binom{n}{i} (q-1)^i = \log_2 V_q(n, R) \leq nH_q(R/n), \quad (7)$$

where $V_q(n, R)$ is the volume of a ball of radius R in \mathbb{F}_q^n and $H_q(x) = -x \log_2 x - (1-x) \log_2 (1-x) + x \log_2 (q-1)$ is the q -ary entropy function⁵. Inequality (7) also gives us an upper bound on the lower embedding efficiency $\underline{\epsilon} = \frac{h}{R}$ for a given relative payload $\alpha = \frac{h}{n}$:

$$H_q^{-1}(\alpha) \leq \frac{R}{n} \implies \underline{\epsilon} = \frac{h}{R} = \alpha \cdot \frac{n}{R} \leq \frac{\alpha}{H_q^{-1}(\alpha)}, \quad (8)$$

⁵ We note that this definition of q -ary entropy function is slightly different from how this concept is usually defined in the literature. The difference is the multiplicative factor $\log_2 q$. This is because we define the relative payload α in bits per pixel, which is more common in steganography, rather than in q -ary symbols per pixel.

where $H_q^{-1}(\alpha) \in [0, (q-1)/q]$. We note that this upper bound on \underline{e} is asymptotically achievable using linear codes because the relative redundancy $(n-k)/n = h/n$ of almost all random $[n, k]$ codes asymptotically achieves $H_q(R/n)$ for a fixed $R/n < (q-1)/q$ and $n \rightarrow \infty$ (see, e.g., Theorem 12.3.5 in [14] for the binary case). Thus, there exist embedding schemes based on linear codes whose lower embedding efficiency is asymptotically optimal.

3 Average distance to code

From the Matrix Embedding Theorem 1, for fixed block length n and embedding capacity $n-k$, the highest lower embedding efficiency is achieved using an $[n, k]$ code with the smallest covering radius R . However, as argued in the Introduction, steganographers are more interested in the embedding efficiency and thus the average number of embedding changes. In this section, we first show that this concept is related to the average distance to code and then we demonstrate that a code with the smallest average distance to code does not have to have the smallest covering radius.

For an embedding scheme from Theorem 1, the expected number of embedding changes for messages uniformly distributed in \mathbb{F}_q^{n-k} is equal to the average weight of all coset leaders of \mathcal{C} . It is reasonable to assume that the messages are drawn uniformly at random from \mathbb{F}_q^{n-k} since typically they will be encrypted before embedding. We now show that the expected number of embedding changes is equal to the average distance to the code defined as

$$R_a = \frac{1}{q^n} \sum_{\mathbf{x} \in \mathbb{F}_q^n} d(\mathbf{x}, \mathcal{C}). \quad (9)$$

Because any two words \mathbf{x}, \mathbf{y} from the same coset \mathcal{C}_i have the same distance from \mathcal{C} : $d(\mathbf{x}, \mathcal{C}) = d(\mathbf{y}, \mathcal{C}) = w(\mathbf{e}_i)$, the weight of a coset leader of \mathcal{C}_i , we have

$$R_a = \frac{1}{q^n} \sum_{\mathbf{x} \in \mathbb{F}_q^n} d(\mathbf{x}, \mathcal{C}) = \frac{1}{q^n} \sum_{i=1}^{q^{n-k}} \sum_{\mathbf{x} \in \mathcal{C}_i} d(\mathbf{x}, \mathcal{C}) = \frac{1}{q^n} \sum_{i=1}^{q^{n-k}} q^k w(\mathbf{e}_i) = \frac{1}{q^{n-k}} \sum_{i=1}^{q^{n-k}} w(\mathbf{e}_i),$$

which is the average number of embedding changes for messages uniformly chosen from \mathbb{F}_q^{n-k} .

The remaining results in this section are formulated for binary codes. We first study codes of small dimension $k = 1, 2$ because such codes allow calculating the average distance to code analytically. Moreover, matrix embedding with codes of small dimension was recently proposed as a means to improve steganographic security when embedding large payloads close to the embedding capacity [13].

Theorem 2. For a binary $[n, 1]$ code

$$R_a \geq \frac{n}{2} \left(1 - 2^{-n+1} \binom{n-1}{\lceil \frac{n-1}{2} \rceil} \right). \quad (10)$$

Proof. Consider the matrix $\mathbf{H} = [\mathbf{I}, \mathbf{1}]$, where \mathbf{I} is the $(n-1) \times (n-1)$ identity matrix and $\mathbf{1}$ is the column of $n-1$ ones. It is easy to see that for $i \leq \lfloor (n-1)/2 \rfloor$ all $\binom{n}{i}$ possible sums of i columns of \mathbf{H} produce all syndromes of weight i and $n-i$. Thus, for n odd, $R_a = 2^{-n+1} \sum_{i=1}^{\lfloor (n-1)/2 \rfloor} i \binom{n}{i}$ and no other code can have a smaller R_a . For n even, we need to include $\binom{n-1}{\lceil (n-1)/2 \rceil}$ sums of $\lceil (n-1)/2 \rceil$ columns of the identity matrix \mathbf{I} . Thus, $R_a = 2^{-n+1} \sum_{i=1}^{\lfloor (n-1)/2 \rfloor} i \binom{n}{i} + \lceil (n-1)/2 \rceil \binom{n-1}{\lceil (n-1)/2 \rceil}$ and, again, no code can have a smaller R_a . Both expressions simplify to the right hand side of (10) after simple algebra. Note that the proof also shows that the inequality (10) is tight.

To present the analogue of Theorem 2 for 2-dimensional codes, we first need to introduce some notation. For an $[n, 2]$ code \mathcal{C} with basis $\{\mathbf{x}, \mathbf{y}\}$, let us define $\psi(\mathcal{C})$ to be the multiset (set with possibly repeated elements) $\{\alpha, \beta, \gamma\}$, where

$$\alpha = |\{i : x_i = y_i = 1\}|, \quad \beta = |\{i : x_i = 1, y_i = 0\}|, \quad \gamma = |\{i : x_i = 0, y_i = 1\}|.$$

Notice that the mapping ψ is well defined, that is, $\psi(\mathcal{C})$ is independent of the choice of a basis for \mathcal{C} .

Theorem 3. Let n be fixed, $n \geq 4$, and let \mathcal{C} be a binary $[n, 2]$ code. Then \mathcal{C} achieves the minimum average distance to code among all binary $[n, 2]$ codes if and only if no coordinate of \mathcal{C} is identically zero, and $\psi(\mathcal{C})$ is in one of the following forms:

$$\{\alpha, \alpha, \alpha + 1\}, \{\alpha, \alpha + 1, \alpha + 3\}, \{\alpha, \alpha + 1, \alpha + 2\}, \{\alpha, \alpha + 3, \alpha + 3\}, \{\alpha, \alpha + 1, \alpha + 1\}.$$

It is quite interesting to note that the most symmetric $[3\alpha, 2]$ codes \mathcal{C} defined by $\psi(\mathcal{C}) = \{\alpha, \alpha, \alpha\}$ are never optimal unless $\alpha = 1$. This theorem is taken from [15].

To show that a code minimizing the average distance to code among all $[n, k]$ codes with given n, k does not need to minimize the covering radius in this class, we now present the following example.

Let \mathbf{M} be the 4×15 binary matrix whose columns are all nonzero vectors from \mathbb{F}_2^4 . Let \mathbf{M}' be a matrix obtained from \mathbf{M} by deleting a single column. Let \mathcal{C} be the $[14, 4]$ code generated by \mathbf{M}' . The average distance to \mathcal{C} is $3548/2^{10}$. We have proved by an exhaustive classification of all $[14, 4]$ binary codes up to isomorphism that, for any $[14, 4]$ code \mathcal{C}' not isomorphic to \mathcal{C} , the average distance to \mathcal{C}' is at least $3602/2^{10}$, which is at least 1.5% more than that of \mathcal{C} . Since the maximum weight of \mathcal{C} is 8, the distance of the all-one vector from \mathcal{C} is 6. However, there are $[14, 4]$ codes with covering radius 5 (see Table 7.1 on page 193 in [14]) and thus \mathcal{C} does not minimize the covering radius among all $[14, 4]$ codes.

Even though the average distance to code and the covering radius are two different values that are not necessarily optimized by the same code, we prove that in the binary case these two concepts asymptotically coincide with increasing length of the code and fixed rate. Let us suppose that we are embedding relative

payload α , $0 \leq \alpha \leq 1$, in an n -element cover object. Thus, the message consists of αn bits and the code that realizes the embedding is a binary $[n, (1 - \alpha)n]$ code⁶. The following theorem states that for almost all such codes the relative covering radius $\rho = R/n$ and the relative distance to code $\rho_a = R_a/n$ converge with $n \rightarrow \infty$.

Theorem 4. For any $0 < \alpha < 1$ and any $\epsilon > 0$, the fraction of all binary $[n, (1 - \alpha)n]$ codes for which $|\rho - \rho_a| \leq \epsilon$ tends to 1 as n goes to infinity.

The proof of this theorem is in the appendix. We note that this result implies that the bound (8) is also an asymptotic bound on the embedding efficiency e

$$e \lesssim \frac{\alpha}{H^{-1}(\alpha)}. \quad (11)$$

4 Practical embedding schemes

In this section, we first explain the reasons for constructing steganographic schemes using q -ary codes with q matched to the embedding operation and then we give several examples of codes suitable for practical applications. Finally, we demonstrate how the codes improve steganographic security of ± 1 embedding in the spatial domain.

Let us start with the simple LSB embedding paradigm frequently employed in steganographic schemes for images, audio, and other digital media objects. To be specific, we assume that the cover is a grayscale digital image and we also assume that the sender can use all pixels for embedding, i.e., the embedding is not constrained to any selection channel [5]. The message bits are embedded as LSBs of pixels along a pseudo-random path determined by a secret stego key. The recipient reads the message from LSBs of pixels obtained by scanning the image in the same pseudo-random order as during embedding.

LSB flipping is a very unnatural operation that is quite detectable by modern steganalytic tools (see [16] and references therein). The fundamental reason for this is the special character of the LSB flipping operation that pairs up grayscale values $2i$ and $2i + 1$ for $i = 0, \dots, 127$. In other words, during embedding the value $2i$ is either left unchanged or changed to $2i + 1$. In particular, it is never changed to $2i - 1$. All reliable LSB detectors rely on this fact in some way or another.

An obvious and quite simple countermeasure is to make the embedding operation symmetrical and allow changes in both directions for all pixel values (with the obvious exception of the boundary values 0 and 255). For example, to modify the LSB of the grayscale value i of a given pixel, the embedder may flip a coin and with probability 1/2 increase the value of i by one and with probability 1/2 decrease its value by one. Note that this process introduces the same distortion

⁶ Statements involving the quantity αn hold whenever this value is an integer, and are void otherwise.

to the image as LSB embedding. This type of embedding is known as ± 1 embedding [17, 18] or LSB matching [19, 20]. In this paper, we will call this method binary ± 1 embedding.

The embedding efficiency of LSB embedding and binary ± 1 embedding is the same and equal to 2—assuming we are embedding a random bit-stream with uniform distribution of 0’s and 1’s, we embed 1 bit by making a change with probability 1/2. However, in the case of ± 1 embedding, we have three possibilities for each pixel—either leave it unchanged or modify by ± 1 . Obviously, we can use the following symbol-assignment function

$$t = s(i) = i \bmod 3 \quad (12)$$

and embed a ternary symbol $t \in \{0, 1, 2\} = \mathbb{F}_3$ in each pixel. We call this method ternary ± 1 embedding.

Assuming the embedded stream of ternary symbols is random with uniform distribution on \mathbb{F}_3^n , the probability that the pixel value i will stay unchanged, be modified by 1, or -1 is the same and equal to 1/3. Thus, we make a change with probability 2/3 and embed $\log_2 3$ bits. The embedding efficiency is thus $\log_2 3 / (\frac{2}{3}) \doteq 2.3774$. This is already larger than the embedding efficiency of binary ± 1 embedding. We can do, obviously, much better because we can now embed up to $\log_2 3$ bits per pixel (bpp) and thus the relative payload α shortens by the same factor. This means that we can further increase the embedding efficiency by applying matrix embedding with ternary codes.

4.1 Examples of good covering codes

Probably the simplest case of matrix embedding is based on q -ary $[\frac{q^m-1}{q-1}, \frac{q^m-1}{q-1} - m, 3]$ Hamming codes, which are perfect codes with minimum distance 3 and covering radius $R = 1$. Since there are q^{n-m} codewords whose distance to code is 0 and $q^n - q^{n-m}$ words $\mathbf{x} \in \mathbb{F}_q^n$ whose distance to code is 1, the average distance to code is $R_a = (q^n - q^{n-m})/q^n = 1 - q^{-m}$. Using Theorem 1, we can embed m q -ary symbols in $\frac{q^m-1}{q-1}$ pixels using at most one change. In other words, we can embed a relative payload $\alpha = m \frac{q-1}{q^m-1} \log_2 q$ bpp with lower embedding efficiency $\underline{e} = m \log_2 q$ and embedding efficiency $e = m \log_2 q / (1 - q^{-m})$.

Note that for $q = 2$, $m = 1$, and the symbol-assignment function (1), we obtain the classical LSB embedding. With increasing m , the payload α decreases while the embedding efficiency increases. The binary Hamming code was used for the first time in the JPEG steganographic algorithm F5 [3].

In Figure 1, we show the upper bound (8) on embedding efficiency for $q = 2, 3, 4$ as a function of relative payload α (in bpp). The embedding efficiency of binary and ternary Hamming codes for different values of m is shown with “+” and “ \times ” signs, respectively. Note that the curves start at the point $\alpha = \log_2 q$, $e = \frac{q}{q-1} \log_2 q$, which corresponds to embedding at the largest relative payload of $\log_2 q$ bpp. We also want to point out the benefit of using q -ary codes for a fixed relative payload α . For example, for $\alpha = 1$, the ternary ± 1 embedding can theoretically achieve embedding efficiency $e \simeq 4.4$, which is significantly




Fig. 1. Embedding efficiency of various q -ary codes with the upper bound (8) for $q = 2, 3, 4$.

higher than 2—the maximal efficiency of LSB embedding at this relative message length.

The remaining non-trivial perfect codes, the binary $[23, 12, 7]$ Golay code and the ternary $[11, 6, 5]_3$ Golay code, also provide very good performance (see Figure 1). The average distance to the binary Golay code is $R_a = \frac{1}{2^{23}}(1 \cdot \binom{23}{1} + 2 \cdot \binom{23}{2} + 3 \cdot \binom{23}{3}) \cdot 2^{12} \doteq 2.8525$, which gives $e = 11/R_a \doteq 3.8562$ at relative payload $\alpha = 11/23 \doteq 0.4783$. The average distance to the ternary Golay code is $R_a = \frac{1}{3^{11}}(1 \cdot \binom{11}{1} \cdot 2 + 2 \cdot \binom{11}{2} \cdot 4) \cdot 3^6 \doteq 1.9012$, giving $e = 5 \log_2 3/R_a \doteq 4.1683$ at relative payload $\alpha = 5 \log_2 3/11 \doteq 0.7204$.

For large payloads close to $\log_2 3$ bpp, the following simple one-dimensional ternary code can greatly improve embedding efficiency. Let \mathcal{C} be the ternary $[n, 1]$ code (1-dimensional subspace of \mathbb{F}_3^n) spanned by the all-one vector. Suppose that we use matrix embedding defined by the code \mathcal{C} , and that the embedding of ternary symbols into the grayscale image is realized using the ± 1 embedding as explained earlier in this section. If we denote the number of 0's, 1's and 2's in an arbitrary vector of \mathbb{F}_3^n by a , b and c , respectively, then the average distance

to \mathcal{C} can be computed as

$$R_a = \frac{1}{3^n} \sum (n - \max\{a, b, c\}) \binom{n}{a} \binom{n-a}{b}, \quad (13)$$

where the sum extends over all triples (a, b, c) of non-negative integers such that $a + b + c = n$. For the matrix embedding part, we can use the ternary matrix $\mathbf{H} = [\mathbf{I}, \mathbf{u}]$ where \mathbf{I} is the $(n-1) \times (n-1)$ identity matrix and \mathbf{u} is the column vector of 2's. The number of bits embedded per n pixels is $\log_2 3^{n-1}$, which gives relative payload $\alpha = \frac{n-1}{n} \log_2 3$. The points $[\alpha, R_a]$ are shown in Figure 1 as “ \star ” signs. For example, we can embed 1.188 bpp with embedding efficiency of almost 3 bits per change.

Binary matrix embedding schemes for large payloads were discussed in [13]. The authors proposed a class of random linear codes of small dimension and codes derived from simplex codes.

Finally, as shown in [5] random linear codes in \mathbb{F}_2^n with small codimension can also be used to construct computationally tractable embedding schemes with improved embedding efficiency (the triangle signs in Figure 1 correspond to codes with codimension $n - k = 19$). In this case, due to the small code codimension the coding can be done using efficient search techniques. Note that these random linear codes outperform binary Hamming codes. Another advantage of this approach is that we obtain a parametrized family of codes rather than a few instances of individual coding schemes, which greatly simplifies implementation.

4.2 Experiments

Even though it is clear that increased embedding efficiency should improve steganographic security, it would be useful to obtain a quantitative statement for a specific embedding scheme applied to real images. We evaluate the steganographic security using the current state-of-the-art blind feature-based classifier [21] on 2500 cover images obtained with 22 different digital cameras. The images include a mixture of indoor and outdoor shots taken under varying light conditions with and without flash, landscapes, and closeups. All images were taken in the raw (uncompressed) format, converted to grayscale and cropped to their central 1000×1000 region. We chose a database of raw images intentionally because previously JPEG compressed images should not be used for spatial domain steganography [22].

For our tests, we used three methods: (1) uncoded binary ± 1 embedding, (2) binary ± 1 embedding with binary Hamming codes, and (3) ternary ± 1 embedding with ternary Hamming codes. Note that in order to embed a message of relative length α bpp, we need to choose the parameter m of the Hamming code so that $(m+1) \frac{q-1}{q^{m+1}-1} \log_2 q < \alpha \leq m \frac{q-1}{q^m-1} \log_2 q$. Obviously, we are most efficient when α is close to $m \frac{q-1}{q^m-1} \log_2 q$. Thus, we chose the payloads for our tests in such a manner so that α is close to the upper bound for both binary and ternary codes.

We ran the following experiment for each payload and each embedding technique. Half of the images from the database were chosen as cover images and the other half were embedded using the corresponding method and payload. Then, using the blind classifier [21] we calculated the Receiver Operating Characteristic curve (ROC) as a measure of separability between the clusters of features of cover and stego images. To obtain a numerical characteristic of the performance of the detector, we used two quantities that are frequently used in current steganalysis literature—false alarms (cover images incorrectly detected as stego) at stego image detection accuracy 50% and 80%. Table 1 shows both numerical characteristics for the three embedding methods and two relative payloads. The parameters for payloads α_1 and α_2 were $m = 2$ and 3, respectively, for both binary and ternary Hamming codes.

Table 1. False alarms at 50% and 80% stego image detection for three embedding methods and two relative payloads.

Embedding method	$\alpha_1 = 0.666$		$\alpha_2 = 0.365$ bpp	
	FA50%	FA80%	FA50%	FA80%
Uncoded binary	1.3%	15%	3.9%	21%
Binary Hamming	2.5%	19%	8.1%	29%
Ternary Hamming	3.9%	21%	12.7%	38%

The results in Table 1 demonstrate that methods that use matrix embedding can be detected less reliably than the uncoded method. For example, for relative payload $\alpha_1 = 0.666$ bpp applying a ternary Hamming code triples the false alarm rate when compared to the uncoded binary ± 1 embedding.

We close this section with some general considerations about the limitations of the applicability of matrix embedding to steganography. It is not clear what improvement in embedding efficiency can be expected from using q -ary codes with $q > 3$ because in this case the act of embedding will have to start making changes with amplitude more than 1. It is an open and little researched area in steganography whether it is beneficial to decrease the number of embedding changes by allowing embedding changes of higher amplitude. In other words, it is not clear whether it is better to make more changes of low amplitude or fewer changes with larger amplitude. The answer to this question likely depends on other properties of the steganographic scheme, such as placement of embedding changes, the type of embedding operation, and the cover object. Recent studies [17] suggest that with increasing amplitude of embedding changes, the detection of steganography becomes more reliable quite rapidly. Because the improvement in embedding efficiency becomes increasingly smaller with increasing q (see Figure 1), it is not likely that incorporating q -ary codes for $q > 3$ would improve steganographic security.

Also, not all steganographic algorithms can benefit from ternary encoding. For example, in the F5 algorithm for JPEG images [3], the absolute value of quantized DCT coefficients is always decreased when necessary to change the LSB. If changes in both directions were allowed in F5, severe artifacts would be introduced in the histogram. Thus, the embedding operation in F5 does not allow applying ternary codes. Another example is Perturbed Quantization [23]. In this case, the direction of embedding changes is determined by side-information provided by a high resolution version of the cover object to minimize the combined distortion due to quantization and embedding. The character of the embedding operation here is also inherently binary.

5 Conclusions

Matrix embedding using linear codes (syndrome coding) is a general approach to improving embedding efficiency of steganographic schemes. The covering radius of the code corresponds to the maximal number of embedding changes needed to embed any message. Steganographers, however, are more interested in the average number of embedding changes rather than the worst case. In fact, the concept of embedding efficiency—the average number of bits embedded per embedding change—has been frequently used in steganography to compare and evaluate performance of steganographic schemes.

In this paper, we showed that the embedding efficiency is determined by the average distance to code rather than the covering radius. Thus, designers of steganographic systems should minimize the average distance to code rather than the covering radius. We demonstrated on an example that, within the class of linear codes of fixed dimension and length, the code with the minimal average distance to code does not have to have the smallest covering radius. However, with increasing code length and fixed rate, we proved that the average distance to code and the covering radius coincide.

In the second part of this paper, we demonstrated that embedding efficiency can be dramatically improved using q -ary codes with q matched to the steganographic embedding operation. We also briefly studied specific coding methods that can be used to realize embedding schemes in practice. In particular, we compared the performance of binary and ternary Hamming codes. Additionally, we proposed a simple one-dimensional ternary code suitable for improving embedding efficiency when embedding large payloads.

An important open problem is how to find families of binary or ternary codes with efficient coding procedures with embedding efficiency close to the theoretical bound. The recently proposed computationally efficient quantizers based on sparse generator matrices [24] look especially relevant to this problem. Alternatively, we plan to investigate random ternary linear codes and development of computationally efficient algorithms similar to those reported in [5].

6 Acknowledgements

The work on this paper was supported by Air Force Research Laboratory, Air Force Material Command, USAF, under the research grants number FA8750-04-1-0112 and F30602-02-2-0093. Research of Petr Lisoněk was partially supported by the Natural Sciences and Engineering Research Council of Canada (NSERC).

The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation there on. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of Air Force Research Laboratory, or the U.S. Government. The authors would like to thank Miroslav Goljan for providing the steganalysis results.

A Proof of Theorem 4

Before we give a proof of the theorem, we formulate two auxiliary lemmas ($H(x)$ is the binary entropy function).

Lemma 1. For any $0 \leq \rho < 1/2$ there exists an integer sequence k_n with

$$k_n/n \leq 1 - H(\rho) + f(n),$$

where $f(n) \in O(n^{-1} \log n)$, such that the fraction of all binary $[n, k_n]$ codes that are $\lfloor \rho n \rfloor$ -coverings tends to 1.

Proof. This lemma is proved in [14, page 325] (Theorem 12.3.5).

Lemma 2. For any $H^{-1}(\alpha) < \rho < 1/2$, the fraction of all binary $[n, (1 - \alpha)n]$ codes with covering radius at most $\lfloor \rho n \rfloor$ tends to 1 as $n \rightarrow \infty$.

Proof. Let us denote $\rho^* = H^{-1}(\alpha)$. Because $1 - H(\rho) < 1 - H(\rho^*)$ and $f(n) \rightarrow 0$ as n goes to infinity, there exists n_0 such that for any $n > n_0$,

$$1 - H(\rho) + f(n) \leq 1 - H(\rho^*) = 1 - \alpha.$$

Applying Lemma 1 to ρ , we obtain an integer sequence k_n for which

$$k_n/n \leq 1 - H(\rho) + f(n) \leq 1 - H(\rho^*) = 1 - \alpha,$$

for $n > n_0$. Thus, $k_n \leq (1 - \alpha)n$ and the fraction of all $[n, k_n]$ codes whose covering radius is at most $\lfloor \rho n \rfloor$ tends to one. However the same is true for at least the same fraction of $[n, (1 - \alpha)n]$ codes as well. This is so because for any two codes $\mathcal{C}_1 \subset \mathcal{C}_2$, \mathcal{C}_1 an $[n, k_1]$ code with covering radius R_1 and \mathcal{C}_2 an $[n, k_2]$ code with covering radius R_2 , we have $R_2 \leq R_1$.

Proof of Theorem 4. Let $\rho^* = H^{-1}(\alpha)$ and let \mathcal{C} be an $[n, (1 - \alpha)n]$ code. From (7) applied to \mathcal{C} (note that $h = \alpha n$), we have for its relative covering radius ρ , $\rho^* = H^{-1}(\alpha) \leq R/n = \rho$. On the other hand, from Lemma 2 it follows that

$\rho \leq \rho^* + \epsilon$ for all $n > n_0$, for a fraction of all $[n, (1 - \alpha)n]$ codes that goes to 1 as $n \rightarrow \infty$.

The average distance to such codes is $R_a = \frac{1}{2^{\alpha n}} \sum_{l=0}^{\rho n} l c_l$, where c_l is the number of coset leaders of weight l . Because $\rho_a \leq \rho$, we need a lower bound on ρ_a . Writing

$$R_a = \frac{1}{2^{\alpha n}} \sum_{l=0}^{\lfloor (\rho^* - \epsilon)n \rfloor} l c_l + \frac{1}{2^{\alpha n}} \sum_{l=\lfloor (\rho^* - \epsilon)n \rfloor + 1}^{\rho n} l c_l, \quad (14)$$

we will find a lower bound on the second sum. To do so, we first derive an upper bound on c_l for l satisfying $l < (\rho^* - \epsilon)n$. We start with

$$c_l \leq \binom{n}{l} \leq 2^{nH(l/n)}. \quad (15)$$

The second inequality follows from Lemma 2.4.2 in [14] and holds for any $l < n/2$ for sufficiently large n (e.g., $n > n_1$). Using the fact that $H(x)$ is increasing on $[0, 1/2]$, from Taylor expansion of $H(x)$ at ρ^* ,

$$2^{nH(l/n)} \leq 2^{nH(\rho^* - \epsilon)} = 2^{n(\alpha - \epsilon H'(\xi))}, \quad (16)$$

where $\rho^* - \epsilon < \xi < \rho^*$. Finally, because H' is decreasing on the same interval,

$$c_l \leq 2^{\alpha n} 2^{-n\epsilon H'(\xi)} < 2^{\alpha n} 2^{-n\epsilon H'(\rho^*)}, \quad (17)$$

for any $l < (\rho^* - \epsilon)n$.

We now obtain a lower bound for R_a . Writing $l_0 = \lfloor (\rho^* - \epsilon)n \rfloor$, from (14)

$$R_a \geq \sum_{l=l_0+1}^{\rho n} \frac{l c_l}{2^{\alpha n}} \geq (\rho^* - \epsilon)n \sum_{l=l_0+1}^{\rho n} \frac{c_l}{2^{\alpha n}} = (\rho^* - \epsilon)n \left(1 - \sum_{l=0}^{l_0} \frac{c_l}{2^{\alpha n}} \right)$$

because $\sum_{l=0}^R c_l = 2^{\alpha n}$. Using (17)

$$R_a \geq (\rho^* - \epsilon)n \left(1 - (\rho^* - \epsilon)n \cdot 2^{-n\epsilon H'(\rho^*)} \right) = (\rho^* - \epsilon)n(1 - \delta(n)), \quad (18)$$

where $\delta(n) \rightarrow 0$ exponentially fast with $n \rightarrow \infty$. Combining this result with $\rho_a \leq \rho \leq \rho^* + \epsilon$, we obtain the following bounds for the average distance to code in terms of the relative quantities (for $n > \max(n_0, n_1)$)

$$(\rho^* - \epsilon)(1 - \delta(n)) \leq \rho_a \leq \rho \leq \rho^* + \epsilon, \quad (19)$$

which proves the claim because $\epsilon > 0$ was arbitrary and $\delta(n) \rightarrow 0$ for $n \rightarrow \infty$.

References

1. Simmons, G.J.: The prisoners' problem and the subliminal channel. In Chaum, D., ed.: *Advances in Cryptology, Proceedings of CRYPTO '83*, Santa Barbara, CA, August 22–24, Plenum Press, New York (1984) 51–67

2. Cachin, C.: An information-theoretic model for steganography. In Aucsmitz, D., ed.: *Information Hiding*, 2nd International Workshop. Volume 1525 of LNCS., Springer-Verlag, New York (1998) 306–318
3. Westfeld, A.: High capacity despite better steganalysis (F5—a steganographic algorithm). In Moskowitz, I.S., ed.: *Information Hiding*, 4th International Workshop. Volume 2137 of LNCS., Springer-Verlag, New York (2001) 289–302
4. Sallee, P.: Model-based methods for steganography and steganalysis. *International Journal of Image Graphics* 5 (2005) 167–190
5. Fridrich, J., Goljan, M., Soukal, D.: Steganography via codes for memory with defective cells. (In: 43rd Conference on Coding, Communication, and Control, September 28–30, 2005)
6. Crandall, R.: Some notes on steganography. Steganography Mailing List, available from <http://os.inf.tu-dresden.de/westfeld/crandall.pdf> (1998)
7. Bierbrauer, J.: *Introduction to Coding Theory*. Chapman & Hall/CRC (2004)
8. van Dijk, M., Willems, F.: Embedding information in grayscale images. (In: Proceedings of the 22nd Symposium on Information and Communication Theory in the Benelux, Enschede, The Netherlands, May 15–16, 2001) 147–154
9. Galand, F., Kabatiansky, G.: Information hiding by coverings. (In: Proceedings ITW2003, Paris, France, 2003) 151–154
10. Fridrich, J., Goljan, M., Soukal, D.: Wet paper codes with improved embedding efficiency. *IEEE Transactions on Information Security and Forensics* 1 (2006) 102–110
11. Williams, F.J.M., Sloane, N.J.: *The Theory of Error-Correcting Codes*. North-Holland, Amsterdam (1977)
12. Bierbrauer, J.: On crandall’s problem. Personal Communication, (available from <http://www.ws.binghamton.edu/fridrich/covcodes.pdf>) (1998)
13. Fridrich, J., Soukal, D.: Matrix embedding for large payloads. In Delp, E., Wong, P.W., eds.: *Proceedings SPIE Electronic Imaging, Security, Steganography, and Watermarking of Multimedia Contents VIII*, San Jose, CA. (2006) W1–W15
14. Cohen, G.D., Honkala, I., Litsyn, S., Lobstein, A.: *Covering Codes*. Volume 54. Elsevier, North-Holland Mathematical Library (1997)
15. Khatirinejad, M., Lisoněk, P.: Linear codes for high payload steganography. (In: presented at AAEECC-16, Las Vegas, Nevada, February 20–24, 2006)
16. Ker, A.: A general framework for structural analysis of LSB replacement. In et al., M.B., ed.: *Information Hiding*, 7th International Workshop, IH 2005, Barcelona, Spain, June 6–8, 2005. Volume 3727 of LNCS., Springer-Verlag, Berlin (2005) 296–311
17. Soukal, D., Fridrich, J., Goljan, M.: Maximum likelihood estimation of secret message length embedded using PMK steganography in spatial domain. In Delp, E., Wong, P.W., eds.: *Proceedings SPIE, Electronic Imaging, Security, Steganography, and Watermarking of Multimedia Contents VII*, San Jose, CA, January 16–20, 2005. (Volume 5681.) 595–606
18. Wong, P.W., Chen, H., Tang, Z.: On steganalysis of plus-minus one embedding in continuous-tone images. In Delp, E., Wong, P.W., eds.: *Proceedings SPIE, Electronic Imaging, Security, Steganography, and Watermarking of Multimedia Contents VII*, San Jose, CA, January 16–20, 2005. (Volume 5681.) 643–652
19. Ker, A.: Resampling and the detection of LSB matching in color bitmaps. In Delp, E., Wong, P.W., eds.: *Proceedings SPIE, Electronic Imaging, Security, Steganography, and Watermarking of Multimedia Contents VII*, San Jose, CA, January 16–20, 2005. (Volume 5681.) 1–15

20. Ker, A.D.: Steganalysis of LSB matching in grayscale images. *IEEE Signal Processing Letters* 12 (2005) 441–444
21. Goljan, M., Fridrich, J., Holotyak, T.: New blind steganalysis and its implications. In Delp, E., Wong, P.W., eds.: *Proceedings SPIE, Electronic Imaging, Security, Steganography, and Watermarking of Multimedia Contents VIII*, San Jose, CA, January 16–19. Volume 6072. (2006) 1–13
22. Fridrich, J., Goljan, M., Du, R.: Steganalysis based on JPEG compatibility. In Tescher, A.G., ed.: *Special Session on Theoretical and Practical Issues in Digital Watermarking and Data Hiding, SPIE Multimedia Systems and Applications IV*, Denver, CO, August 20–24, 2001. (Volume 4518.) 275–280
23. Fridrich, J., Goljan, M., Soukal, D.: Perturbed quantization steganography using wet paper codes. In Dittman, J., Fridrich, J., eds.: *Proceedings ACM Multimedia and Security Workshop, Magdeburg, Germany, September 20–21, 2004*, (ACM Press, New York) 4–15
24. Wainwright, M.J., Maneva, E.: Lossy source encoding via message-passing and decimation over generalized codewords of LDGM codes. In: *Proceedings of the International Symposium on Information Theory, Adelaide, Australia. (2005)*