

Data Embedding In Hardcopy Images Via Halftone-Dot Orientation Modulation

Orhan Bulan^a, Vishal Monga^b, Gaurav Sharma^a, and Basak Oztan^a

^aUniversity of Rochester, Rochester, NY, USA

^bXerox Corporation, Webster, NY, USA

ABSTRACT

The principal challenge in hardcopy data hiding is achieving robustness to the print-scan process. Conventional robust hiding schemes are not well-suited because they do not adapt to the print-scan distortion channel, and hence are fundamentally limited in a detection theoretic sense. We consider data embedding in images printed with clustered dot halftones. The input to the print-scan channel in this scenario is a binary halftone image, and hence the distortions are also intimately tied to the nature of the halftoning algorithm employed. We propose a new framework for hardcopy data hiding based on halftone dot orientation modulation. We develop analytic halftone threshold functions that generate elliptically shaped halftone dots in any desired orientation. Our hiding strategy then embeds a binary symbol as a particular choice of the orientation. The orientation is identified at the decoder via statistically motivated moments following appropriate global and local synchronization to address the geometric distortion introduced by the print scan channel. A probabilistic model of the print-scan process, which conditions received moments on input orientation, allows for Maximum Likelihood (ML) optimal decoding. Our method bears similarities to the paradigms of informed coding and QIM, but also makes departures from classical results in that constant and smooth image areas are better suited for embedding via our scheme as opposed to busy or “high entropy” regions. Data extraction is automatically done from a scanned hardcopy, and results indicate significantly higher embedding rate than existing methods, a majority of which rely on visual or manual detection.

Keywords: Hardcopy data embedding, halftone image watermarking, print watermark

1. INTRODUCTION

Multimedia data hiding has been an active area of research over the past decade,^{1 2} with several emerging applications in content authentication, anti-piracy data hiding, and digital rights management. Most of the techniques developed however, focus on hiding data in digital images, audio and video that are consumed in the digital space as well.

Hardcopy watermarking or data hiding in images that is required to survive the print-scan (digital-analog-digital) process has come to receive significant attention recently. The fact that printed media comprises a significant slice of official, legal and transactional data such as IDs, passports, contracts etc. makes the problem of considerable economic interest. Particular motivating applications lie in fingerprinting (identifying where a hardcopy document originated), integrity verification and copy protection (ensuring that originals/copies of an original hardcopy are distinguishable and identifiable as such), copyright enforcement (ready identification of copyright owner from printed content), and meta-data communication (auxiliary data embedding such as web URLs, product UPC codes, etc).

From a technical standpoint, hardcopy data hiding is distinguished from other multimedia hiding schemes by the presence of the print-scan process. The similarity between robust data hiding, and communications with side information is now well-known.¹ In that view, hardcopy data hiding schemes must specifically adapt to the

Send correspondence to O. Bulan: E-mail: bulan@ece.rochester.edu, Telephone: 1 585 275-8122, Address: Electrical and Computer Engineering Department, University of Rochester, Rochester, NY, 14627-0126, USA, WWW: www.ece.rochester.edu/projects/iplab

This work was supported in part by Xerox Corporation and by the a grant from New York State Office of Science, Technology and Academic Research (NYSTAR) through the Center for Electronic Imaging Systems (CEIS).

characteristics of the print-scan distortion channel. The key component of the printing process is a bit-depth reduction step called digital halftoning which reduces the original continuous tone (typically 8 bits/pixel) image to a 1 bit/pixel binary image. Halftoning aims to produce an illusion of continuous tone by cleverly trading off amplitude resolution for spatial resolution. Algorithms for digital halftoning can be classified into: 1.) point processes, or screening which involves comparison with a spatially varying threshold array, 2.) neighborhood processes or error diffusion, and 3.) iterative or search based methods. A detailed tutorial can be found in.³ Of the above three, screening (particularly clustered-dot) is dominant in xerographic printers while error diffusion is used extensively in inkjet printing. Once printed, the image transitions to the analog domain, and can be re-digitized by scanning. For the hardcopy data hiding problem, it should be noted that the input to the “physical print-scan channel” is a binary or halftone image, and hence the distortions are intimately tied to the nature of the halftoning algorithm employed.

The methods proposed for data-embedding in hardcopy images may be grouped into two main categories. The first corresponds to robust embedding methods that are intended to survive printing and scanning but do not directly exploit characteristics of the printing process in the embedding.^{4,5} The second corresponds to techniques that use the particular characteristics of the printing process (e.g. halftoning) for embedding.⁶⁻¹³ The exploitation of this specific knowledge typically offers greater potential for embedding and is well-suited to hardcopy applications since the embedding occurs just prior to printing.

The methods in the first category are fundamentally limited by the lack of knowledge of the halftone print process. From a communications viewpoint, this adversely affects detection accuracy because the embedding does not adapt to the characteristics of the (print-scan) channel. A majority of the methods in the latter category¹¹⁻¹⁴ embed text and other binary patterns that require a (manual) visual detection process, e.g. by overlaying a pre-designed binary pattern on the hardcopy print. While this is useful for some applications such as hardcopy authentication, a large class of applications such as meta-data embedding, document tracking in workflows, and secure hiding in adversarial scenarios, can only be enabled by automated data recovery. Many methods that do allow automated detection perform embedding by^{6,7} modulating binary halftone outputs with pre-decided binary message signals to produce the watermarked image. This leads to significant image distortions even for modest embedding rate. A joint halftoning and watermarking scheme was proposed in.⁸ However, their method employs a search over several candidate halftone patterns, in a way similar to direct binary search,¹⁵ and is therefore very slow for real-time printing applications. In¹⁰ Damera-Venkata *et al.* developed a real-time joint halftoning and embedding method for high-rate information embedding in error diffused halftone images used in inkjet printers. While practical, their scheme does not extend to the large xerographic family of printers which do not use error diffusion.

In this paper, we propose a high-rate hardcopy data hiding framework by halftone dot orientation modulation. The key novel aspect of our work is that we break the trade-off between being robust to the xerographic print-scan process, and still enabling automated data extraction. This is achieved by careful halftone dot design such that the shape (orientation) of the halftone dot maintains robustness under print-scan, and a probabilistic model of the print-scan channel which follows naturally.

In particular, we use analytic threshold functions to produce screening-based clustered-dot halftones, such that for a given gray level, we generate elliptic halftone dots. Further, our formulation allows control of the ellipse orientation which is then used for data embedding. Statistical moments that uniquely identify an ellipse orientation are used for detection purposes. We then characterize densities of the “received moments” conditioned on a given orientation, i.e. the probabilistic channel model. Maximum-likelihood (ML) decoding is then performed using the conditional densities.

Because our scheme uses the host signal, i.e. the image gray level, in the mapping from the message (binary symbol) to the mark (orientation) space, our scheme may be categorized as an informed coding scheme.¹ Further, it's a quantization based embedding where the quantization is done in the various possible elliptical orientations, e.g. horizontal (0 degree), vertical (90 degree), diagonal (+/- 45 degree) orientations etc. Despite the similarities, many characteristics of our proposed scheme make departures from classical results. As an example, in our proposed data hiding scheme, constant or low-entropy regions of the (contone) cover image are better suited for embedding and detection. As will become clear in the ensuing description, this is desirable both from the

viewpoint of introducing low-visual distortion, and also mandated by the print-scan channel which is more likely to preserve halftone dots in smooth or constant regions.

Our method of embedding has similarities with a hardcopy barcode scheme called DataGlyphs.^{16,17} It is a technique for interpreting printed binary patterns as digital data, which is commonly the case with barcodes. The similarity lies in that the binary codewords in DataGlyphs are also differently (orthogonally) oriented. These codewords have also been used to approximate binary halftone patterns,¹⁶ which translates to embedding in real images intended for printing. Here lies the first major difference of our scheme with DataGlyphs. In our scheme, we essentially develop and adapt a halftoning method for data embedding, which means that image fidelity is an explicit goal and not an outcome as with DataGlyphs. The second and more crucial difference is in the detection. DataGlyphs detect by first estimating the binary pattern from a scanned image and then correlating with the embedded binary codewords. We however detect from the scanned image by using statistical criterion, and print-scan channel modeling. Accurately estimating the binary image that were input to the print process is feasible for low-resolution barcodes, but not for high resolution image printing as is typical today. On a conceptual level, our method can in fact be understood as an extension of DataGlyph type image barcodes to a true image watermarking method.

The rest of the paper is organized as follows. Section 2 describes halftone print-scan channel characteristics and motivates our data hiding scheme. Embedding and detection by modulating the orientation of elliptical halftone dots is detailed in Section 3. Section 3 also describes the print-scan probabilistic channel model and corresponding maximum-likelihood (ML) decoding. Experimental results in the form of practically achievable embedding rates are presented in Section 4. Section 5 concludes the paper by summarizing the key aspects of our scheme and provides directions for future work.

2. HALFTONE PRINT CHANNEL CHARACTERISTICS

Figure 1 illustrates the data-embedding and recovery processes for a halftone printing channel. A contone original image $I(x, y)$ and message data \mathbf{m} constitute the inputs at the encoder. The halftoning and embedding process produces a halftoned image $I^h(x, y)$ that: (a) visually approximates the original image $I(x, y)$, and (b) carries the message data \mathbf{m} . The halftoned image is printed and the decoder attempts to recover the embedded data from the scan of the print. In practice, in the block diagram representation of Fig. 1, the process of data embedding may be performed prior to, concurrently with, or after the halftoning. If the embedding is performed prior to the halftoning, it operates on the contone image and the halftone print and scan channel appears as a rather severe distortion that the watermark must overcome. If on the other hand, the embedding is performed after the halftoning, it is necessary to ensure that the embedding process retains image quality which is challenging given that the image is already in a binarized representation. Methods that perform the embedding and the halftoning jointly are therefore typically preferable in applications where this is feasible.

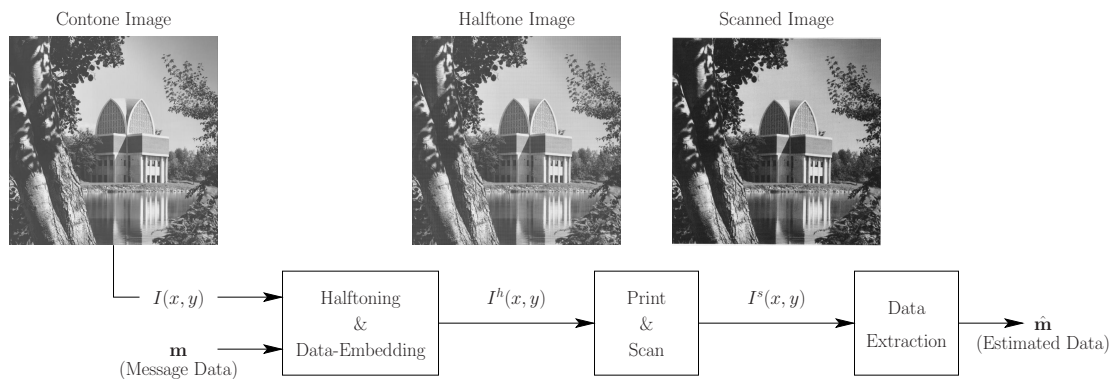


Figure 1. Data embedding for the halftone print channel.

Several characteristics of the halftone print-scan channel distinguish it from general multimedia communication channels:

1. Since the detection is performed based on a scan of printed image, print-scan geometric distortion (which formed the original inspiration for the StirMark¹⁸ rotation, scaling, translation and random-bending attack) is a natural part of the communication channel even in the absence of any malicious attacker. We note though that for screened halftones (both clustered and dispersed-dot), the inherent periodicity of the screening process provides a natural template for global synchronization. The more challenging problem of synchronizing in the presence of local geometric perturbations must, however, be addressed independently.
2. The discrete nature of the halftoning process ensures that there are a finite, albeit large, number of possible halftoned images. Of the set of possible halftoned images, a smaller subset will provide a perceptually acceptable representation of the original contone image. The cardinality of this set defines a limit on the capacity of the halftone channel, *even in the absence of any print-scan distortion and noise*.
3. Based on the preceding observation we note that the capacity of the data-embedding channel can be bounded above by the entropy of the available halftone configurations. While this characterization is intractable for complete images, some useful observations can be made with regard to the local embedding capacity for different image regions. Firstly observe that, in purely black and white regions of the cover image, the capacity is zero since there is only one configuration for the halftones in these regions. Furthermore, for regions with close to 50% halftone area coverage, again the capacity is rather small since the number of acceptable halftone configurations in these regions are small. For the specific case of clustered-dot halftones, typically only two differently phased checkerboard configurations are acceptable in the region of 50% area coverage and thus these regions have essentially a single bit capacity (over the relatively large regions that may have this area coverage).
4. For conventional multimedia, informed embedding attempts to shape the power distribution of the embedded mark to match the power distribution of the cover.¹⁹ Specifically, in the case of images, under this criterion, the embedding distortion tends to be concentrated in relatively busy regions of the image whereas relatively smooth regions carry lower watermark power. For the halftone channel, however, the situation is reversed. In relatively smooth regions that are not close to the critical gray levels indicated in the preceding part, there is considerable flexibility in the choice of the local halftone configuration yielding a higher upper bound on capacity as per the above criterion. On the other hand, in the busy image regions, typically, there is lesser flexibility in the choice of halftone configurations and concurrently a smaller capacity.
5. Mean-squared error metrics for the embedding distortion, that are often used for the multimedia embedding scenario, are less meaningful in the halftone channel, where it is well known that the pixel-wise thresholding provides the lowest mean squared error but rather poor visual quality. Alternate Human Visual System (HVS) based error metrics²⁰ must therefore be utilized to evaluate the embedding distortion in a meaningful fashion.


Motivated by these characteristics, we propose a halftone data embedding method, in the following section. Our experimental results in Sec. 4 corroborate several of the above observations.

3. EMBEDDING VIA HALFTONE ORIENTATION MODULATION


3.1. Encoder

We propose embedding data in clustered halftone dots by generating elliptically shaped dots, and modulating the dot orientation. Figure 2(b) shows the zoomed version of a constant gray level image with elliptical dots oriented along vertical and horizontal directions. This paper specifically focuses on binary embedding using two orthogonal orientations.

For notational simplicity, we base our description on the 0/90 degree case, though as we observe subsequently in this section, the +/- 45 degree case can provide beneficial symmetries in practical systems. Since host



(a) Gray level halftone image



(b) Zoomed version with dot orientation along vertical and horizontal directions

signal (gray level) information is incorporated in the mapping from the message (binary symbol) to the mark (orientation) space, our scheme is an instantiation of the paradigm of informed coding.¹

The most common technique to generate clustered halftone dots is comparing the contone image $I(x, y)$ against a periodic threshold function $T(x, y)$. One way of generating threshold function is to design discrete array of thresholds for a given screen frequency and dot shape. Alternatively, threshold function can be represented analytically which allows to control dot shape and other dot characteristics easily. Such a function is defined by Pellar^{21,22} for images taking values over the interval $[-1, 1]$ as:

$$T(x, y) = \cos(2\pi f_x x) \cos(2\pi f_y y) \quad (1)$$

where f_x and f_y represent the frequencies along the horizontal axis x and vertical axis y respectively. We modify the Pellar threshold function in such a way that it allows to modulate dot orientation along one of the orthogonal x and y axis. It is expressed as:

$$T(x, y) = \text{sgn}(\cos(2\pi f_x x) \cos(2\pi f_y y)) |\cos(2\pi f_x x)|^{\gamma_x} |\cos(2\pi f_y y)|^{\gamma_y} \quad (2)$$

where γ_x and γ_y are parameters that control the dot shape and $\text{sgn}(\cdot)$ denotes the signum function, i.e.,

$$\text{sgn}(t) = \begin{cases} 1 & \text{if } t \geq 0 \\ 0 & \text{if } t = 0 \\ -1 & \text{if } t < 0 \end{cases} \quad (3)$$

In the case γ_x equals to γ_y comparing a constant gray-level image with the threshold function above results in dot that is symmetric about the x and y axes and close to circular in the highlights and shadows. Setting one of them to larger value on the other hand, gives rise to elliptical shape oriented along one of vertical or horizontal directions (in highlights and shadows). If γ_x is greater than γ_y dot is oriented along vertical direction and when γ_y is larger it is oriented along the horizontal axis. Therefore, γ_x and γ_y are orientation parameters that offer flexibility in embedding. For our binary embedding scenario, we assume that γ_x and γ_y take on one of two positive values $\{\Gamma_1, \Gamma_2\}$ where $\Gamma_1 > \Gamma_2 > 0$, and the embedding rule is specified as

$$\gamma_x = \Gamma_1, \gamma_y = \Gamma_2, \quad \text{if } b_i = 1, \text{ and} \quad (4)$$

$$\gamma_y = \Gamma_1, \gamma_x = \Gamma_2, \quad \text{if } b_i = 0 \quad (5)$$

$$(6)$$

so that the “major axis” of the elliptical dot is oriented vertically for bit values of 1 and horizontally for bit values of 0.

The halftone image is then obtained as follows:

$$I^h(x, y) = \begin{cases} 1 & \text{if } I(x, y) \geq T(x, y), \\ 0 & \text{if } I(x, y) < T(x, y), \end{cases} \quad (7)$$

For typical printing resolutions, the proposed orientation based embedding causes a minimal degradation in halftone image quality. When the two orthogonal orientations used for the embedding are symmetrically located with respect to the two orthogonal directions that define the printer addressability, the symmetry in the printing system ensures that the tone response for halftones of the two orientations is identical. Thus the embedding does not change the average gray-level, which is the fundamental image fidelity criterion in the halftone reproduction. This happens for instance when the embedding utilizes a +/- 45 degree orientation in a printing system with the natural 0/90 degree axes. This characteristic is similar to that for GlossMarks⁹ that use orientation in order to encode a relatively coarse gloss pattern. Note also that in pure black or pure white regions of the cover image, the threshold function has no effect. Thus, when dictated by the cover image, our “power-law modulation” for changing the orientation sacrifices the orientation modulation in favor of maintaining image fidelity.

3.2. Decoder

At the decoder the goal is to estimate the embedded data from a scan of printed image. Since the data is embedded in dot orientation, statistical criteria that capture orientation must be developed for data extraction. In particular, we propose moments as a statistic for detection. The moments are computed as:

$$\sigma_x = \frac{\sum_{x,y \in C} I^s(x, y)(x - \bar{x})^2}{\sum_{x,y \in C} I^s(x, y)} \quad (8)$$

$$\sigma_y = \frac{\sum_{x,y \in C} I^s(x, y)(y - \bar{y})^2}{\sum_{x,y \in C} I^s(x, y)} \quad (9)$$

where $\bar{x} = \frac{\sum_{x,y \in C} I^s(x, y)x}{\sum_{x,y \in C} I^s(x, y)}$ and $\bar{y} = \frac{\sum_{x,y \in C} I^s(x, y)y}{\sum_{x,y \in C} I^s(x, y)}$ represent the center of mass of the halftone dot. σ_x ,

σ_y indicate moments along orthogonal x and y axes respectively. All pixels in the halftone cell contribute to moments along horizontal and vertical directions according to their value $I^s(x, y)$ and distance to center of mass in the horizontal or vertical direction. In essence, the moment vector $\sigma = (\sigma_x, \sigma_y)$ can be interpreted as a lower-dimensional (2-D) feature vector extracted from the high dimensional received signal, i.e. halftone dot that helps identify the orientation.

The print-scan channel introduces two kinds of distortion on the halftone image: 1.) geometric distortions due to imperfect dot registration, and 2.) noise dependent on the input gray level. Further, geometric distortions caused by the print-scan process comprise of both global rotation, scaling and also perturbations to local image geometry. In fact, the well known random being Stirmark attack was inspired by a modelling of the print-scan process.

It is well known⁵ that periodicity in the halftone dot structure may be exploited for estimating global rotation and scaling parameters. Therefore, we look for peaks in the Fourier Transform of the scanned image to estimate global rotation and scaling. Local geometric distortions are harder to model, and have not been explicitly addressed in previous work. We propose to use decision directed synchronization which is widely done in digital communications.²³ That is, we exploit the knowledge of the center of mass of a previously decoded halftone dot (using our decision statistic) to adjust local geometry. Fig. 2 summarizes our decoding process. Post synchronization, the detector makes a decision on the received orientation.

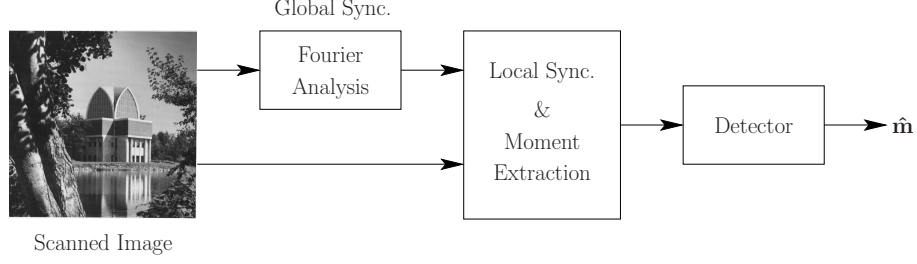


Figure 2. Decoder

3.2.1. Hard Decoding

Since the embedding orients the “major axis” of the elliptical dots along respectively the horizontal or vertical directions, hard decoding is performed by simply comparing the moments σ_x and σ_y estimated from the scanned image for a given cell, i.e., the hard detection decision rule is given by

$$\sigma_x \underset{0}{\overset{1}{\geq}} \sigma_y. \quad (10)$$

3.2.2. Soft Decoding and Statistical Channel Modeling

The hard decoding criterion in the previous subsection is intuitively meaningful but does not hold any optimality guarantees.

Next, we develop a detection criterion which is optimal in the Maximum Likelihood (ML) sense. The detector must make a decision on the received orientation based on the moment vector $\boldsymbol{\sigma} = (\sigma_x, \sigma_y)$. This can be accomplished via the knowledge of joint conditional density function $f_{\sigma_x, \sigma_y}(\sigma_x, \sigma_y | \Theta_i)$, where the conditioning is performed on the orientation Θ_i , $i = 1, 2$, i.e. Θ_1 and Θ_2 represent horizontal and vertical orientations for the elliptic halftone dots respectively. Describing multi-dimensional density functions analytically, esp. in the absence of a physically inspired model, constitutes a hard task. In order to make the problem tractable, we assume conditional independence given by

$$f_{\sigma_x, \sigma_y}(\sigma_x, \sigma_y | \Theta_i) = f_{\sigma_x}(\sigma_x | \Theta_i) \cdot f_{\sigma_y}(\sigma_y | \Theta_i) \quad (11)$$

We validated this assumption experimentally by observing the 2×2 covariance matrices of the received random vector $\boldsymbol{\sigma}$ conditioned on Θ_1, Θ_2 , and for several gray levels. Several empirical estimates of these covariance matrices obtained from scanned images corresponding to prints for different gray levels are included in the appendix. It is readily apparent from that most of these matrices are close to diagonal indicating that they may *approximately** be modeled as independent.

We therefore focus on performing a probabilistic modeling of the channel shown in Fig. 3. That is, we characterize the (marginal) conditional densities $f_{\sigma_x}(\sigma_x | \Theta_i)$, $f_{\sigma_y}(\sigma_y | \Theta_i)$ of the observed moments σ_x , σ_y .

Figs. 4 (a) - (d) shows the normalized histograms of the received moments σ_x , σ_y . In particular, for an encoded gray level of 32%, Fig. 4 (a) shows histograms of σ_x , σ_y for the (transmitted) horizontal orientation Θ_1 , while Fig. 4 (b) plots the same for the vertical orientation Θ_2 . Figs. 4 (c)-(d) repeat these plots for gray level 52%. Based on similar observations for several gray levels, we model the conditional densities by the exponential power density family²⁴ parameterized by the scale parameter a and shape parameter b . The density is given by

$$f_X(x) = \frac{1}{2a\Gamma(1 + \frac{1}{b})} e^{(-|\frac{x}{a}|^b)} \quad (12)$$

*Zero covariance across the off-diagonals guarantees uncorrelatedness and not independence. Hence we use the term approximately.


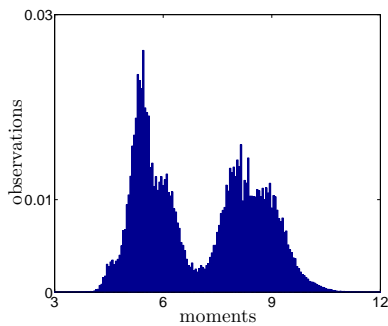


Figure 3. Probabilistic Model


For $b = 1$ this reduces to the Laplacian distribution, while for $b = 2$ this has the same form as the Gaussian. We estimated the parameters of the distribution by using Expectation-Maximization for exponential families.²⁵

A decision on the received orientation is then made using the maximum-likelihood criterion


$$\Theta^* = \arg \max_{i=1,2} f_{\sigma_x}(\sigma_x|\Theta_i) \cdot f_{\sigma_y}(\sigma_y|\Theta_i) \quad (13)$$



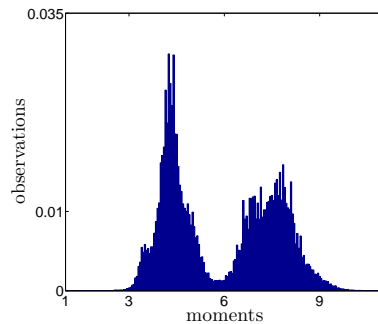
(a) Histogram of moments along horizontal direction given two orientations at gray level 32%



(b) Histogram of moments along vertical direction given two orientations at gray level 32%



(c) Histogram of moments along horizontal direction given two orientations at gray level 52%



(d) Histogram of moments along vertical direction given two orientations at gray level 52%

Figure 4.

4. EXPERIMENTAL RESULTS

First, we demonstrate performance over constant gray level images. We work with a 19200×19200 constant gray level image, and a 2400 dpi printer which means the image renders in a 8 inch by 8 inch area. We set orientation parameters γ_x to 1 and γ_y to 2 in order to achieve a horizontally oriented dot, and swap the parameters to orient the dot in vertical direction. The halftone images are generated with a screen frequency of 75 cells per inch along both vertical and horizontal directions. If each halftone cell is used for embedding, this amounts to 360,000 bits of embedding in the image.

Table 1 shows the bit error rate(BER) performances for different gray levels with hard decoding, and ML decoding. In the table, NS is used to indicate that symbol synchronization is not achieved for the corresponding gray level. As is evident from the table, data embedded in gray levels ranging from 20% to 40% and from 60% to 70% can be recovered with a small BER [†]. In highlights, shadows and mid-tone gray levels, however detection was not possible because symbol synchronization is completely lost in those regions.

Gray Level(percent)	10	20	30	40	50	60	70	80
Hard Decoding BER	NS	0.0109	0.0032	0.0238	NS	0.0045	0.0303	NS
ML Decoding BER	NS	0.0085	0.0012	0.0073	NS	0.0043	0.0339	NS

Table 1. Bit error rates of different gray level images for two orientation scenario with Hard decoding and ML decoding. NS denotes situations where no synchronization was achieved. A counter-intuitive observation from the plot is that at the 70% gray level, soft or ML decoding seems to perform marginally worse than hard decoding. This clearly is not possible in theory but we attribute this behavior due to deviations of the assumed channel model from the true one, and a fall in the ability to synchronize for gray levels in the 70% range.




Figure 5. BER performance of various gray levels for two orientation scenario with Hard decoding and ML decoding.

Figure 5 plots BER performance as in Table 1 but across a much larger number of gray levels. The plot reveals that the lowest BER is achieved for gray levels around 30% and 60%. This is intuitively clear as highlights, shadows, and midtones (approx. 50%) do not provide sufficiently robust dot configurations; best performance is hence expected between white and midtones, and black and midtones. A counter-intuitive observation from the plot is that at the 70% gray level, soft or ML decoding seems to perform marginally worse than hard decoding. This clearly is not possible in theory but we attribute this behavior due to deviations of the assumed channel model from the true one, and a fall in the ability to synchronize for gray levels in the 70% range.

Next, we present preliminary results for embedding in natural images. In Figure 6(a) we embed the bit sequence (1, 0, 1, 0, 1, 0, 1) in the dot orientation. In Figure 6(b), black regions depict the areas where alternating 0's and 1's are correctly detected in a sequence of seven halftone dots. The horizontal stripes in Fig. Figure 6(b)

[†]0% gray level corresponds to digital level 255 and 100% gray level corresponds to digital level 0.



(a) Halftoned Chapel image with bit sequence $\{1, 0, 1, 0, 1, 0, 1\}$ embedded along vertical and horizontal directions (b) Marked Chapel image for illustration. Black regions show that embedding and detection can be achieved successfully in those areas

Figure 6.

result from the fact that we perform a single global synchronization, and proceed in a left to right direction for detection. In principle, the Fourier transform based global synchronization method could be applied within local regions as well.

Observe crucially that in Fig. 6(b), the embedded data is recovered successfully with greater likelihood in the low entropy (constant or smooth) regions in comparison with the high entropy (busy) regions of the cover image. This corroborates our initial claims that low-entropy regions are better suited for data hiding (and detection), which is necessarily an outcome of the characteristics of the print-scan process. This is remarkably quite different from classical schemes that recommend embedding in high-entropy or busy image regions.

Finally, in the example presented here we successfully decode 93,935 bits from the printed and scanned Chapel image in Figure 6(a). The original size of the electronic image file was 1024×1024 .[‡] This is clearly much higher than the practical embedding rates achieved by both: a) visual detection methods,¹² that embed a small number of text characters, visual patterns and b) robust embedding methods designed to survive the print scan channel.^{4,5} In follow on work²⁶ we demonstrate a significant performance improvement by utilizing image adaptive decoding.

5. CONCLUSION AND DISCUSSION

Existing schemes for hardcopy data hiding either modulate binary halftone patterns with pre-decided message code words, and/or hinge on visual detection of text or other embedded patterns. In both cases, the rate of embedding is low. Additionally, many applications are enabled only by automated detection. In this paper, we develop a high-rate hardcopy data hiding scheme based on halftone dot orientation modulation. We define analytic threshold functions to produce screening-based clustered-dot halftones that produce elliptic halftone dots for a graylevel. Quantization based embedding is performed by associating message symbols with particular ellipse orientations. A statistical modeling of the print-scan channel is performed that enables maximum-likelihood (ML) identification of the orientation at the decoder. Experimental results show that our method achieves high embedding rates while simultaneously allowing for automated detection. Finally, a major conceptual insight brought forward by our scheme is that constant and smooth image regions are better suited for hardcopy data hiding vs. “busy” or high-entropy regions. We demonstrate that this is because of the unique characteristics of the print-scan channel which is relatively benign to constant regions. Further the characteristics of the halftone process imply that is desirable even from the viewpoint of introducing low visual distortion to the image.

[‡]The image was scaled prior to halftoning (jointly embedding) and subsequently printing.

APPENDIX A.

In Sec. 3.2.2 we construct a probabilistic model for the print-scan process which conditions received moments on input orientations by assuming conditional independence for the multi-dimensional density functions. We validated this assumption by observing the 2×2 covariance matrices of the received random moments σ_x, σ_y conditioned on Θ_1, Θ_2 for several gray levels. We listed here some covariance matrices for different gray levels.

$$C_1 = \begin{bmatrix} 0.36 & 0.08 \\ 0.08 & 0.17 \end{bmatrix}$$

$$C_2 = \begin{bmatrix} 0.54 & 0.02 \\ 0.02 & 0.26 \end{bmatrix}$$

$$C_3 = \begin{bmatrix} 0.56 & 0.07 \\ 0.07 & 0.29 \end{bmatrix}$$

$$C_4 = \begin{bmatrix} 0.80 & 0.07 \\ 0.07 & 0.53 \end{bmatrix}$$

$$C_5 = \begin{bmatrix} 1.02 & -0.20 \\ -0.20 & 1.21 \end{bmatrix}$$

$$C_6 = \begin{bmatrix} 0.20 & 0.07 \\ 0.07 & 0.40 \end{bmatrix}$$

$$C_7 = \begin{bmatrix} 0.22 & 0.07 \\ 0.07 & 0.40 \end{bmatrix}$$

$$C_8 = \begin{bmatrix} 0.26 & 0.01 \\ 0.01 & 0.56 \end{bmatrix}$$

$$C_9 = \begin{bmatrix} 0.39 & 0.06 \\ 0.06 & 0.70 \end{bmatrix}$$

$$C_{10} = \begin{bmatrix} 0.50 & -0.09 \\ -0.09 & 0.82 \end{bmatrix}$$

REFERENCES

1. I. J. Cox, M. L. Miller, and J. A. Bloom, *Digital Watermarking*, Morgan Kaufmann, 2001.
2. C. C. J. Kuo, T. Kalker, and W. Zhou, "Special issue on digital rights management," *IEEE Signal Processing Magazine*, Mar 2004.
3. R. Ulichney, *Digital Halftoning*, MIT Press, 1987.
4. D. Corp, "Digimarc mediabridge," 2000.
5. K. Solanki, U. Madhow, B. S. Manjunath, S. Chandrasekaran, and I. El-Khalil, "Print and scan resilient data hiding in images," *IEEE Trans. Info. Forensics and Security* **1**, pp. 464–478, Dec. 2006.
6. K. Tanaka, Y. Nakamura, and K. Matsui, "Embedding secret information into a dithered multi-level image," *Proc. IEEE Military Communications Conf.*, pp. 216–220, Sept. 1990.
7. Z. Baharav and D. Shaked, "Watermarking of dithered halftone images," **3657**, pp. 307–316, Jan. 1999.
8. D. Kacker and J. P. Allebach, "Joint halftoning and watermarking," *IEEE Trans. Sig. Proc.*, pp. 243 – 257, Apr. 2003.
9. C. Liu, S. Wang, and B. Xu, "Authenticate your digital prints with Glossmark images," in *Proceedings IS&T NIP20: International Conference on Digital Printing Technologies*, pp. 312–316, Oct. 2004.
10. N. Damera-Venkata, J. Yen, V. Monga, and B. L. Evans, "Hardcopy image barcodes via block error diffusion," *IEEE Trans. Image Proc.* **14**(12), pp. 1977–1989, 2005.
11. K. T. Knox and S.-G. Wang, "Digital watermarks using stochastic screens," in *Proc. SPIE: Color Imaging: Device Independent Color, Color Hardcopy, and Graphic Arts II*, G. B. Beretta and R. Eschbach, eds., **3018**, pp. 316–322, Feb. 1997.

12. G. Sharma and S.-G. Wang, "Show-through watermarking of duplex printed documents," in *Proc. SPIE: Security, Steganography, and Watermarking of Multimedia Contents VI*, E. J. Delp and P. W. Wong, eds., **5306**, Jan. 2004.
13. B. Oztan and G. Sharma, "Continuous phase modulated halftones and their application to halftone data embedding," in *Proc. IEEE Intl. Conf. Acoustics Speech and Sig. Proc.*, **II**, pp. 333–336, May 2006.
14. S. Wang, "Digital watermarking using phase-shifted stoclastic screens." US Pat. #6,252,971, 2001.
15. M. Analoui and J. P. Allebach, "Model based halftoning using direct binary search," *Proc. SPIE Annual Symp. on Electronic Imaging*, Feb. 1992.
16. D. L. Hecht, "Printed embedded data graphical user interfaces," *IEEE Computer*, pp. 47–55, Mar. 2001.
17. D. L. Hecht, "Embedded data glyph technology for hardcopy digital documents," in *Proc. SPIE: Color hard copy and graphic arts III*, J. Bares, ed., **2171**, pp. 341–352, Feb. 1994.
18. F. A. P. Petitcolas, R. J. Anderson, and M. G. Kuhn, "Attacks on copyright marking systems," in *Information Hiding, Second International Workshop, IH'98*, pp. 219–239, (Portland, OR, USA), Apr. 1998.
19. J. K. Su and B. Girod, "Power-spectrum condition for energy-efficient watermarking," *IEEE Trans. Multimedia* **4**(4), pp. 551–560, 2002.
20. A. U. Agar, F. A. Baqai, and J. P. Allebach, "Human visual model based color halftoning," in *Digital Color Imaging Handbook*, G. Sharma, ed., CRC Press, Boca Raton, FL, 2003. Chapter 7.
21. R. J. Pellar and L. D. Green, "Electronic halftone generator." United States Patent No. 4149183, 1979.
22. R. J. Pellar, "Electronic halftone generator." United States Patent No. 4196451, 1980.
23. J. Proakis, *Digital Communications*, McGraw-Hill, New York, fourth ed., 2001.
24. H. Stark and J. W. Woods, *Probability, Random Processes and Estimation Theory for Engineers*, Prentice-Hall, Englewood Cliffs, NJ, 1986.
25. J. Salojarvi and K. Puolamaki, "Expectation maximization algorithms for conditional likelihoods," *Proc. 22nd International Conference on Machine Learning*, Aug. 2005.
26. O. Bulan, G. Sharma, and V. Monga, "Image adaptive decoding for halftone orientation based data hiding," To be submitted to ICIP 2008.