RC22671 (W0205-037) May 6, 2002
Other

IBM Research Report

Application of Blurring Filters to Improve Detection of Invisible Image Watermarks

Gordon W. Braudaway, Frederick C. Mintzer
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 218
Yorktown Heights, NY 10598

Application of blurring filters to improve detection of invisible image watermarks

Gordon W. Braudaway, Fred Mintzer

International Business Machines Corporation, Thomas J. Watson Research Center
Yorktown Heights, New York 10598

ABSTRACT

Earlier, we presented a highly robust invisible watermarking method having a payload of one bit - indicating the presence or absence of the watermark. Other invisible watermarking methods also possess this property. Our method, by design, places a major portion of the scant energy of the watermark into the low spatial frequencies of the watermarked image while leaving the higher spatial frequencies, contributed by the original image, largely unaltered. In this paper we will show that application of *blurring filters* to a presumed watermarked image before attempting watermark detection can improve the probability of detection. Blurring filters, because they suppress high spatial frequencies, are generally quite destructive of image quality. However, for a watermark which has dominant low frequency content, the application of a blurring filter can serve to improve the statistical environment for watermark detection and thereby improves detection probability. This is especially true for detections attempted from scans of printed images. The halftone screening process used in preparing images for printing adds significant high frequency noise to the printed image, and that noise is detrimental to watermark detection. Application of a rudimental blurring filter to a watermarked image before detection will be demonstrated showing that increased amounts of blurring improve the probability of watermark detection. However, the amount of blurring can not be increased without limit, and the amount of beneficial blurring will be shown to be related to the amount of low-pass filtering used in constructing the watermark.

Keywords: image and document security, invisible watermarking, image processing, fingerprinting

1. INTRODUCTION

In a previous paper [1] we presented a highly robust invisible watermarking method having a payload of one bit. The one-bit payload indicates the presence or absence of the watermark. Limiting the payload to one bit has the potential of maximizing the robustness of an embedded watermark. A few other invisible watermarking methods also possess this property. Our method, by design, places a major portion of the scant energy of the watermark into the low spatial frequencies of the watermarked image. The higher spatial frequencies are therefore contributed only by the original image and are largely unaltered. Because of this deliberate design choice, we will be able to show that application of *blurring filters* (also called low-pass filters) to a presumed watermarked image before attempting watermark detection can improve the probability of detection. Blurring filters, because they suppress high spatial frequencies, are generally quite destructive of image quality. However, for a watermark which has dominant low frequency content, we will show that the application of a blurring filter can serve to improve the statistical environment for watermark detection and thereby improves the probability of detection. This is especially true for detections attempted from scans of printed images. The halftone screening process used in preparing images for printing adds significant high frequency noise to the printed image, and that noise is generally detrimental to watermark detection. However, the amount of blurring used to improve detection can not be increased without limit, and the amount of beneficial blurring will be shown to be related to the amount of low-pass filtering used in constructing the watermark.

2. OVERVIEW OF A WATERMARK EMBEDDING TECHNIQUE

An overview of a the referenced robust invisible watermarking technique is presented here as foundation for the following discussion. The watermark that is to be embedded into a digitized image is represented as a rectangular array of numeric elements, called a *watermarking plane*. A watermarking plane either has the same number of rows and columns as the original image into which it is to be embedded or it is replicated by tiling to cover the entire image, with any excess trimmed at the right and bottom edges of the image. The invisible watermark is embedded as a random, but reproducible, small modulation of the luminance of every image pixel. This is done by multiplying the luminance of each pixel by its corresponding value in the watermarking plane. The watermark thus becomes a permanent and inseparable part of the watermarked image. *Small random luminance modulation without chroma alteration is the essence of this method of invisible watermarking.*

The techniques of constructing the watermarking plane that help insure its resistance to malicious removal will be discussed in more detail shortly. Let it suffice for now to say that the elements of the watermarking plane are random values that have been linearly mapped into the domain of 1 to $1\text{-}2\beta$, where β is called the *modulation strength*, and additionally, the elements, as an ensemble, are scaled and adjusted to have both a mean and median of $1\text{-}\beta$.

3. OVERVIEW OF WATERMARK DETECTION

Detection of an embedded watermark is a difficult task, especially if manipulation of the watermarked image has occurred. It requires detecting the presence of a particular known small modulation of a random two-dimensional carrier, where the carrier is composed of the pixel luminance values of the original image.

The first challenge in detecting a watermark is to determine in what ways a watermarked image may have been manipulated. It may have been cropped, reduced nonlinearly in size, and rotated through a small angle. A previous paper [2] details how a digitized image, suspected of being a watermarked image, can be realigned automatically by referencing its original image. Since the original image and the watermarking plane were geometrically aligned with each other at the time of watermark embedding, any realigned watermarked image will also be aligned with an expected watermarking plane, and detection of the expected watermark can then be attempted.

For our technique, [1] the process of watermark detection is designed to establish, by hypothesis testing, the probability that an expected watermark was embedded into the image. This is implemented by examining the luminance of every pixel in the watermarked image and comparing it to the average luminance of its neighboring pixels. A useful pixel neighborhood large enough to have reasonably well behaved statistics is an 11 by 11 square, giving each pixel 120 neighbors. The corresponding element from the reconstructed watermarking plane is also compared to the average of its neighboring elements in a neighborhood of corresponding location and size.

The correspondence of the suspected watermarked image with a candidate watermarking plane is established by counting coincidence events. If the luminance of a pixel is greater than the average luminance of its neighboring pixels, and the value of the corresponding element from the watermarking plane is greater than the average value of its neighboring elements, the agreement is considered to be a coincidence and a single *coincidence counter*, initially set to zero, is incremented by one. In the same manner, if both are less than their neighborhood averages the agreement is also considered to be a coincidence and the coincidence counter is also incremented by one. But if the luminance of a pixel is greater than the average luminance of its neighboring pixels and the value of the corresponding element from the watermarking plane is less than or equal to the average value of its neighboring elements, or vise-versa, the disagreement is considered to be a non-coincidence and the coincidence counter is decremented by one.

The count in the coincidence counter, in relation to the total number of coincidence events, is a very useful metric. The count is related directly to the probability of detection of a specific watermark. For example, if a watermarking plane is embedded into an image having uniform finite pixel values, and for the moment ignoring the effects of limited precision numerical truncation, it should be apparent that the expected value of the coincidence count will be equal to the number of pixels in the image. This expected result is a direct consequence of the careful statistical construction of the watermarking plane. Conversely, the expected value of the coincidence count should be zero if no watermarking plane has been embedded. This effect can be verified by embedding a watermark into an image having uniform gray pixel values and attempting detection by the technique described.

However, when using original images derived from natural scenes, which can have highly varying pixel values, the count in the coincidence counter is not as distinct. The inherent variability of the pixel values in the small neighborhood regions of natural scenes and the numerical truncation caused by use of short precision values are significant sources of noise. They cause the coincidence counter to have values other than those theoretically expected. But fortunately for most images of interest (other than artificial images composed of pure noise for pixel values, and, of course, the usual few pathological images) there is enough relative constancy in enough small neighborhoods to allow unequivocal detection of an embedded watermark with a probability that approaches unity (a significantly large positive count), and an unequivocal non-detection (a very small count, either positive or negative) if the expected watermark is not embedded into the image.

It must be emphasized that the count in the coincidence counter is image specific. It depends not only on the chosen modulation strength of the embedded watermark and the size of the detection neighborhood but also, and very strongly, on the variability of the pixels in the specific image. Nevertheless, for a given watermarked image it is a good quantitative metric for objectively comparing improvement in detection probabilities resulting from changes in the detection process, and it will be used here specifically for that purpose.

In undertaking a watermark detection, we are performing a hypothesis test. Under hypothesis $\mathbf{H}_0$, the image bears no watermark; under hypothesis $\mathbf{H}_1$, the image bears a watermark. In the following, $w(i,j)$ will refer to the value of an element of the watermarking plane lying in the row i and column j, and the $\underline{W}(i,j)$ will refer to the average of its 120 neighboring elements. The value $c_k(i,j)$ will refer to the luminance component value of an image pixel lying in row i, column j and color plane k, and $\underline{C}_k(i,j)$ the average of its 120 neighboring component values.

In the case that $\mathbf{H}_0$ is true, the "texture statistic" of the watermarking plane, $\textbf{\textit{sgn}}\,[w(i,j) - \underline{W}(i,j)]$, should be uncorrelated with the "texture statistic" of the k'th color plane of the candidate image, $\textbf{\textit{sgn}}\,[c_k(i,j) - \underline{C}_k(i,j)]$. In this case, the product of the two texture statistics, a coincidence event which we denote as $E(i,j)$, will be one, or minus one, with equal probability; this is a binomial distribution with a mean 0 and an event probability $p = 0.5$.

Suppose we sum N values of $E(i,j)$ for all values of i, j, and k, and call that sum S. If $\mathbf{H}_0$ is true, then the mean of S will be zero, the variance will be N, and the standard deviation, σ, will be $\sqrt{N}$. For large N, we can use the normal distribution to approximate the binomial and to relate multiples of standard deviation to defined probabilities. $\mathbf{H}_0$ will be judged false if S is more than ten σ from the mean, as the probability of this is extremely small. When $\mathbf{H}_0$ is false, $\mathbf{H}_1$ must be true.

4. IMPORTANT CONSIDERATIONS FOR CONSTRUCTING A WATERMARKING PLANE

The method of constructing the watermarking plane is fundamental to insuring *robustness* of the embedded watermark, meaning its ability to survive determined attacks or extensive manipulation, including halftone screening. To this end, the procedure by which the values of its elements are chosen must be carefully cast using techniques borrowed from cryptography and two-dimensional signal processing theory, as well as from mathematical statistics. Elements of the watermarking plane are based on sixteen-bit pseudo-random values taken from a *cryptographicly secure* sequence of bits. A practical interpretation of *cryptographicly secure* is that if n bits of the sequence are known, the probability of correctly predicting the value of the next bit in the sequence is 0.5, and this condition remains true for sufficiently large values of n. The secure sequence is a necessary component of the claim of robustness of the watermark, since if the sequence can be deduced by a malicious party, the watermarking plane can be reproduced and the watermark can be removed easily and effectively by inversion. The secure sequence must also be reproducible, at will, knowing only the details of its generating method, a specific private cryptographic key and a publicly known seed needed to initialize a particular sequence. The property of reproducibility allows discarding of the watermarking plane after it is embedded in an image. At a future time a reproduced sequence can be used to reconstruct the watermarking plane when watermark detection is to be attempted.

The value of each element in the watermarking plane is initially a uniformly distributed sixteen-bit pseudo-random number determined by linearly mapping successive sixteen-bit groups taken from a secure sequence. A plane having elements determined in this manner will contain a substantial amount of high spatial frequencies, some having wavelengths as short as two pixels, but no shorter. (The limiting short wavelength is known from the Nyquist sampling theorem.[3] It can be intuitively verified by imagining an image of alternating black and white pixels arranged in a checkerboard pattern. Any attempt to represent a wavelength shorter than two pixels is not possible.) Although high frequency content is beneficial in making an embedded watermark less visible, it also makes it vulnerable when a watermarked image is reduced in size or halftoned in preparation for printing. The deliberate suppression of high frequency content in the watermarking plane, thereby producing a noise pattern with larger features, is done to make the watermark less vulnerable to typical image manipulation (and deliberate attacks) at the cost of possibly making it more visible.

We have chosen to suppress high frequency content in the watermarking plane using relatively complex "frequency-domain" low-pass filtering. To do this, a square array filled with uniformly distributed sixteen-bit pseudo-random number is converted to frequency coordinates using a two-dimensional Discrete Fast Fourier Transform (DFFT). The corresponding horizontal and vertical frequency coordinates are referred to as υ and ω. The dimensions of the υ-ω space are enlarged by a factor that is a power of two. We have chosen the enlargement factor to be four, and the frequency coordinates of the enlarged space are referred to as s and t. The added surface area of the s- t space is filled with zero values. A circularly-symmetric low-pass filter is then applied to all frequency values in the s- t space. Figure 1 shows one quadrant of a low-pass filter for an initial square having of 256x256 elements. The filter has a small flat circular area near the origin and a smooth transition from an amplitude of 1 to 0. The transition is a displaced raised

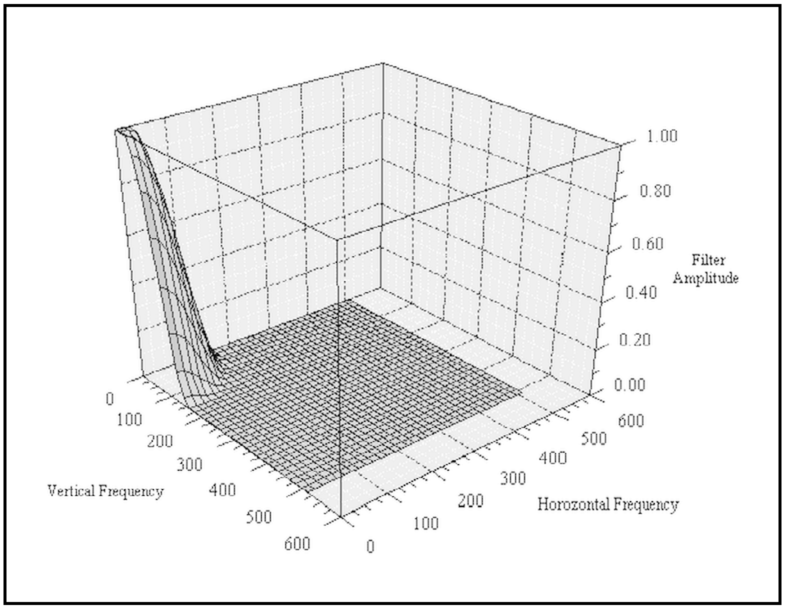

Figure 1. Single quadrant low-pass filter in frequency coordinates for a 1024 x1024 element watermark. Filter roll-off is a raised cosine function.

cosine function, and the smoothness of the transition, with its continuous first derivative, tends to minimize undesired consequences of this rather severe filter. The area in s- t space with the enlargement has spatial frequencies ranging from -512 to 512 cycles.

The final watermarking plane is produced using a reverse DFFT, and is a spatial array of real values with dimensions x and y that are 1024x1024, four times as large as the original 256x256 element plane. Elements of the new plane have no wavelength shorter than eight pixels. Elements of the enlarged filtered plane are then linearly mapped into the domain of 1 to $1\text{-}2\beta$, and additionally, as an ensemble, scaled and adjusted to have both a mean and median of $1\text{-}\beta$. Also importantly, the 1024x1024 watermarking plane created by this method can be abutted to itself on the any side using tiling, and the abutted planes will match seamlessly with one another. Thus, by tiling the watermarking plane can be made as large as desired. The watermarking plane can also be made larger or smaller by choosing an original plane larger or smaller than 256x256, such as 512x512.

5. EFFECT OF BLURRING FILTERS ON THE IMAGE AND ITS EMBEDDED WATERMARK

Meticulous construction of the watermarking plane, with careful attention given to the statistical properties of its elements, will pay off well in protecting the watermark from routine image manipulation as well as deliberate attacks. But the deliberate low-pass filtering used in its construction is equally important. To demonstrate this fact, a uniformly gray image is watermarked using an excessively large modulation strength of 50% (normal modulation strengths are 2% to 3%, and almost never greater than 10%.) A watermark with modulation strength of 50% embedded in a uniformly gray image is highly visible. An enlarged corner of the watermarked image is shown in Figure 2a. The image is then subjected

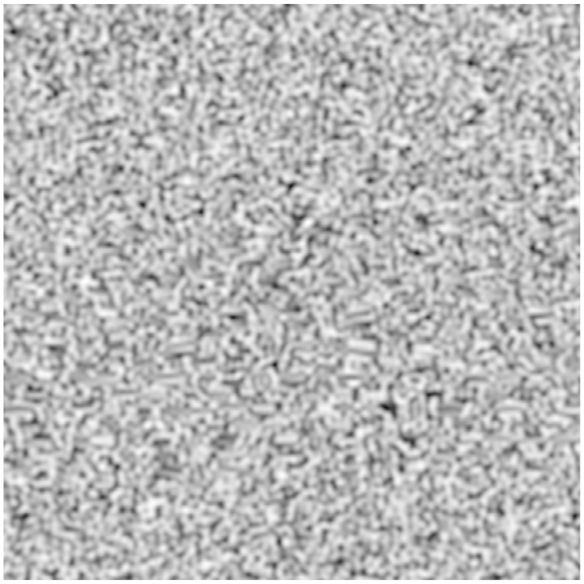

Figure 2a. Uniform gray plane watermarked with modulation strength β = 50%

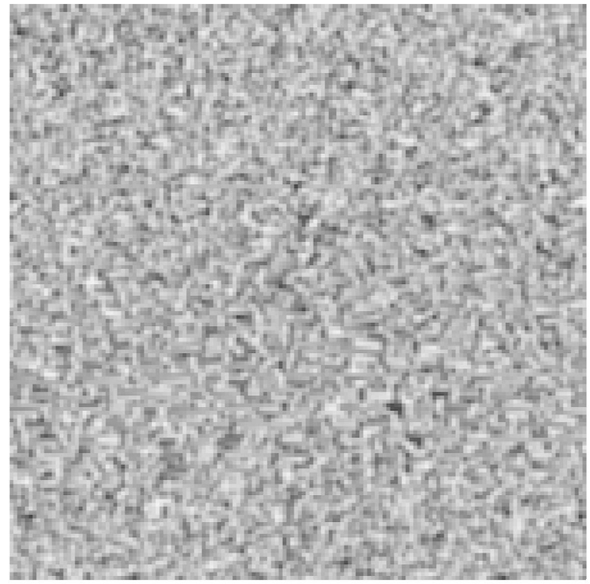

Figure 2b. Watermarked image of Figure 2a after primitive blurring (reduced x0.25, then enlarged x4). Watermark pattern is largely unaltered.

to a severe and primitive blurring filter. The primitive filter used first creates an intermediate image one sixteenth as large as the original with new pixel values that are the average of pixel values from conjoined square arrays of sixteen pixels from the original. The height and width of the intermediate image are each reduced by a factor of four. The intermediate image is then enlarged by pixel replication; that is, by copying the value of each pixel into a square array of sixteen pixels to produce a final filtered image the same size as the original. A corresponding section of the filtered image is shown in Figure 2b. This primitive blurring filter is a far less sophisticated than the low-pass filter used in creating the watermarking plane, but by observation, the damage done to the structure of the embedded watermark is relatively small, and a nearly perfect detection of the watermark is still possible.

The effect of the same severe blurring filter on image content (but with using more sophisticated reduction and enlargement filters) will be demonstrated and analyzed using two test images shown in Figures 3a and 3b. Figure 3a is a high resolution scan of real objects, originally in color, made using an IBM Pro3000[4] scanner mounted on a copy stand. Figure 3b is scanned from a photograph and is representative of a class of images that have a large content of high frequency features, making them challenging to watermark.

An enlarged section of the Figure 3a is shown before and after blurring in Figures 4a and 4b. The reduction and enlargement are, as before, 0.25 and 4. Image blurring is obvious when comparing Figure 4a with Figure 4b. It is particularly evident when comparing the warp and woof of the woven fabric in the lower left and upper right and in the pitting of the almond shell at the center. These images demonstrate that such severe blurring is obviously damaging to image quality, yet, as will be shown by objective measurement, are beneficial to embedded watermark detection.

The measurement plan is to embed a watermark having a particular modulation strength into a test image and then attempt to detect the embedded watermark while varying the degree of blurring applied to the watermarked image. We will use a rudimental but easily quantifiable blurring filter consisting of two stages, first an image reduction by a specified factor and second, an image enlargement by the reciprocal of that factor. For example, if the fixed reduction factor is ½, then the enlargement factor will be 2, and the filtered image will be the same height and width as the original watermarked image (this particular blurring will be referred to as a $2 \bullet 2$ filter). The objective measure of detection strength will be the ratio of positive coincidence counts to total coincidence counts, as described above.

The relationship between the ratio of positive to total coincidence counts, called the *detection ratio*, D, can be connected to the probability of a watermark detection using concepts of mathematical statistics stated above. For N observed

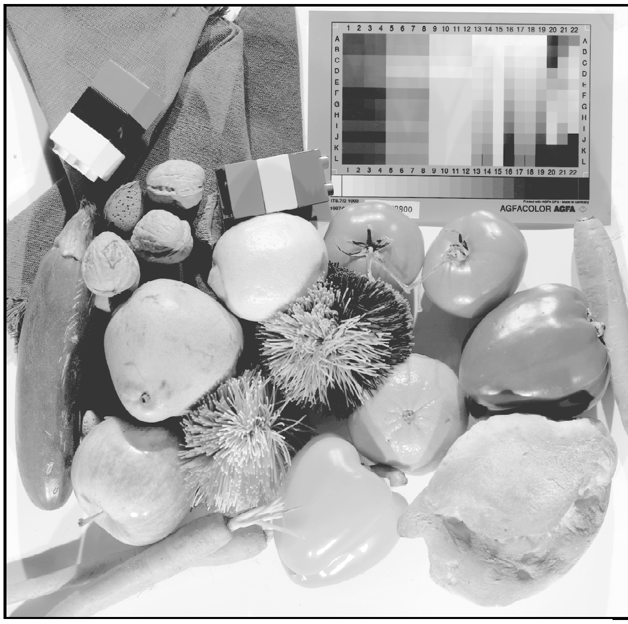

Figure 3a. A digitized test image composed of natural objects directly scanned using an IBM Pro3000 scanner (2237x2184 24-bit pixels)

Figure 3b. A digitized test image scanned from photographic negative using an IBM Pro3000 scanner. (2000x1568 24-bit pixels)

Figure 4a. A section of Figure 3a enlarged to show high frequency detail.

Figure 4b. A section of Figure 3a after severe low-pass filtering enlarged to show loss of high frequency detail.

coincidence events, the σ-*multiple* is $|N(D\text{-}0.5)|/\sigma$, and, if ten or greater, represents a near certainty of watermark detection.

Results of the measurements are shown in Table 1. Notice in columns 2 and 5 of the Table 1 the images were processed with a modulation strength, β, of 0%, meaning that they were not watermarked. This represents a required credibility test that every watermarking method must satisfy. No watermarking method can be considered credible if it can find watermarks in images that do not have them. The example method passes this test with no difficulty. The behavior

	Image 3a; β = 0%	Image 3a; β = 2%	Image 3a; β = 5%	Image 3b; β = 0%	Image 3b; β = 2%
Blurring Filter Factor	Detection Ratio: σ-Multiple	Detection Ratio: σ-Multiple	Detection Ratio: σ-Multiple	Detection Ratio: σ-Multiple	Detection Ratio: σ-Multiple
1 • 1	0.5004: 1.002	0.5470: 116.581	0.5604: 176.106	0.4999: 0.119	0.5270: 45.204
2 • 2	---	0.5705: 188.204	0.5977: 294.004	---	0.5329: 56.570
3 • 3	---	0.5805: 221.063	0.6103: 340.655	---	0.5364: 64.896
4 • 4	---	0.5835: 235.775	0.6139: 359.516	---	0.5413: 77.091
5 • 5	---	0.5679: 196.944	0.5925: 296.716	---	0.5386: 75.700
6 • 6	---	0.5569: 169.587	0.5806: 261.662	---	0.5281: 57.897
8 • 8	---	0.5274: 85.142	0.5405: 133.596	---	0.5181: 40.952
10 • 10	---	0.5098: 31.227	0.5168: 55.952	---	0.5075: 18.253

Table 1. *Detection ratios* and *σ-multiples* of watermarks detected in images 3a and 3b with increasing blurring filter factors. Any σ-multiple greater than 10.0 (ten times the standard deviation) is considered a very strong watermark detection.

exhibited in columns 3, 4, and 6 shows that the probability of detection increases up to a filter factor of 4 • 4, and then decreases relatively rapidly at 8 • 8 and greater. This behavior is entirely consistent and expected, and is caused by the amount of blurring used in the design of the watermarking plane. The frequency plane roll-off of the blurring filter reaches zero, as shown graphically in Figure 1, at about frequency 128, which is one fourth the Nyquist limit frequency of 512. The scant energy of the watermark is therefore concentrated in the low spatial frequencies (long wavelengths). Conversely, interfering carrier noise from the image, the pixel luminance variations, is distributed across all spatial frequencies. Blurring of the watermarked image has the effect of reducing high frequency noise contributed by the image without damaging the low-frequency dominant watermark, thereby improving the signal-to-noise ratio of the watermark and making it easier to detect.

In all cases for test images 3a and 3b the detection ratios are well below the theoretical limit 1.0, which can be obtained using a uniform gray plane as the test image. However, the probability of detection remains statistically extremely high. The number of coincidence events, N, ranged from 6.15×10^6 to 10.29×10^6 for Image 3a, and from 2.14×10^6 to 5.12×10^6 for Image 3b. These yield standard deviations in the range 2480 to 3210 for Image 3a and 1462 to 2265 for Image 3b. (These ranges are caused by "outlier" rejection screens that ignore coincidence events that are obviously improbable, such as when the luminance of a pixel is greater than (1 + 3β) times the average luminance of its neighbors.) Also note that increasing the modulation strength from 2% to 5% increases the probability of watermark detection, as expected, all be it at a possible cost of making the watermark texture visible. Any σ-multiple greater than ten can be considered, for practical purposes, a certainty. (Even DNA matching can not determine unique human identity beyond a statistical probability.)

Another interesting phenomenon observed is that strong watermark detection occurs at filter factors of 8 • 8, where image height and width have been reduced by factors of eight and image area by sixty-four. This result would argue that to best protect small images, the watermark should be applied before image reduction, and further, that reasonable assurance of watermark detection can occur using enlargements of small images watermarked in this manner.

6. BENEFICIAL EFFECT OF BLURRING FILTERS ON HALFTONED IMAGES

Watermarking has its greatest utility and value if it survives printing and is detectable in a scan of the printed image. This should be a minimum requirement for claims of robustness. Preparation for printing can be very damaging to image watermarks because the necessary halftone screening process introduces a large amount of high-energy high-frequency noise into the image. As evidenced above, the presence of high-frequency noise decreases the signal-to-noise ratio of the embedded watermark and makes it more difficult to detect. Blurring was shown to improve the signal-to-noise ratio that was beneficial to watermark detection, and that result should apply to scans of printed images as well.

To test this conjecture, Image 3a was enlarged by a factor of two, its pixels were converted from three-dimensional color to four dimensional color, (Cyan, Magenta, Yellow, and Black) the printer's primaries. The four-dimensional pixels were then halftoned into four-bit toner-density values ready for printing. In this example, the halftoned image was then "soft-proofed;" that is, it was converted back to the original three dimensional color space. An enlarged segment of the

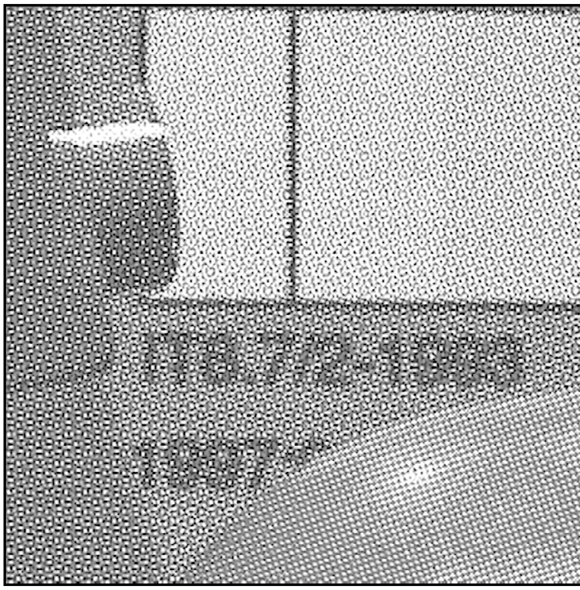

Figure 5a. A section of Figure 3a enlarged to show high-energy high-frequency artifacts caused by halftoning .

Figure 5b. Figure 5a after applying a blurring filter with a filter factor of 8 • 4.

halftoned image is shown in Figure 5a. Figure 5b shows the same enlarged image segment, but after a blurring filter having an 8 • 4 filter factor was applied. Note that the 8 • 4 filter factor is needed because the image was enlarged by a factor of two before being halftoned.

The "soft-proofed" halftoned image was reduced to its original dimensions and watermark detection was attempted. The results are shown in Table 2. It is apparent from results that the halftoning process damaged the watermark sufficiently

	Image 3a; $\beta = 2\%$
Low-pass Filter Factor	Detection Ratio: σ-Multiple
2 • 1	0.5010: 1.410*
4 • 2	0.5083: 16.300
6 • 3	0.5110: 25.241
8 • 4	0.5122: 29.679
10 • 5	0.5095: 24.380
12 • 6	0.5067: 17.844
16 • 8	0.5020: 5.599*
20 • 10	0.5005: 1.357*

Table 2. *Detection ratios* and *σ-multiples* of watermarks detected in image 3a after being enlarged, converted to CMYK, halftoned, reduced, and converted back to 3-D color.
***watermark not detected**

badly to make it undetectable (σ-multiple of 1.41) without the aid of a blurring filter, but with filtering, the watermark can be recovered from a scan of the printed image. Although the σ-multiple for an 8 • 4 filter factor is 29.68 instead of 235.78 for the non halftoned image, it is still abundantly large, and beyond any reasonable doubt the watermark was detected.

7. CONCLUDING REMARKS

In this paper we have shown that application of a blurring filter to a presumed watermarked image before attempting watermark detection can substantially improve the probability of detection. Blurring filters, because they suppress high spatial frequencies, are generally quite destructive of image quality. However, for watermarks that are designed to have dominant low frequency content, the application of a blurring filter can serve to improve the statistical environment for watermark detection. This is especially true for detections attempted from scans of printed images. Although the halftone screening process used in preparing images for printing adds significant high-energy high-frequency noise, often making watermark detection impossible, blurring of the scanned image prior to attempting watermark detection has demonstrated the ability to unequivocally detect the embedded watermark.

The watermarking method used as an example demonstrates that increased amounts of blurring can improve the probability of watermark detection. However the amount of blurring can not be increased without limit, since eventually a uniform monochrome intermediate image would be produced. What is surprising is the amount of beneficial blurring that can be done. Although the highest probability of detection occurs when the characteristics of the applied blurring filter match those of the low-pass filter used in constructing the watermarking plane, as signal processing theory would predict, useful watermark detection does occur with even greater amounts of blurring. When the optimum filter factor is $4 \bullet 4$, very high detection probabilities are still achieved using a filter factor of $8 \bullet 8$, and even as high as $10 \bullet 10$ for non-scanned images, with the latter producing an intermediate image 1% the size of the original. This means that considerable correlated information exists in the watermarking plane at frequencies well below the roll-off frequency of its constructing low-pass filter.

A useful consequence of this observation is that it leads to a dependable strategy for watermarking small images. The strategy is to watermark an image at high resolution (e.g., 3000x2000 pixels) before reducing its dimensions by a factor as large as ten, and its area by a factor as large as 100. A watermark can then be detected unequivocally in an enlargement of the small image. Small images are the standard medium of the Internet, and this strategy may prove very useful for protecting them. It has been applied to small images produced for the website of the Hermitage Museum, St. Petersburg, Russia. [5]

8. REFERENCES

1. Gordon W. Braudaway, "Protecting Publicly-Available Images with an Invisible Image Watermark," *Proc. of the IEEE ICIP'97*, Vol. I, pp. 524-527, October 1997.
2. Gordon W. Braudaway, Fred Mintzer, "Automatic recovery of invisible image watermarks from geometrically distorted images," *Journal of Electronic Imaging*, Vol. 9, No. 4, pp. 477-483, October 2000.
3. H. Nyquist, "Certain Topics in Telegraph Transmission Theory," *Trans. of the A.I.E.E.,* pp. 617-644, April 1928.
4. F. P. Giordano, et al., "Evolution of a high-quality digital imaging system," in *Sensors, Cameras, and Imaging Applications of Digital Photography,* Nitin Sampat, Thomas Yeh, Editors, Proc. of SPIE Vol. 3650, pp. 110-118, January 1999.
5. Fred Mintzer, Gordon Braudaway, et al., "Populating the *Hermitage Museum's* New Web Site," *Communications of the ACM*, pp. 52-60, August 2001.