

Affine-invariant image watermarking using the hyperbolic chirp

Peter Fletcher, Kieran Larkin, Stephen Hardy
Canon Information Systems Research Australia
Sydney, Australia
Email: peter.fletcher@cisra.canon.com.au

Abstract—Image watermarking is the robust, imperceptible embedding of a small quantity of data into a digital image, and the subsequent recovery of this data, perhaps after the watermarked image has been distorted. We present a new watermarking technique which is robust to many image distortions, in particular arbitrary affine transformations of the image. The method achieves its robustness through the use of one-dimensional chirp functions. An affine-invariant affine-invariant detection method exists for such functions using a Radon transform, yet they are not detected trivially by a malicious attacker. The method also provides a way to determine any affine transformation applied to the watermarked image by using an affine-invariant property of groups of intersecting lines.

I. INTRODUCTION

With the growing popularity of the Internet and consequent attempts to protect copyrighted media, some applications of watermarking techniques have been well publicised, as they potentially assist the owners of copyrighted works to both prove the provenance of a work, and also to automatically detect the presence of that work.

The application of watermarking for digital rights management may provide motivation to maliciously attack any such watermarks, thereby removing any consequences of their presence.

There are also many benign uses for watermarking, for which there is little motivation to remove the mark, such as the inclusion of image meta-data (camera type, date, GPS coordinates) which are easily lost through publication or a simple file format conversion.

To be successful, a watermarking system must be robust to common image editing operations; printing and scanning; and malicious attack.

Common editing operations are scaling, cropping, rotation, colour changes, de-noising, data compression and sharpening.

The effect of printing and scanning an image will generally reduce high frequencies, add spatial harmonics from non-linear effects, clip image details in very dark and very light regions, change aspect ratio slightly, and introduce strong frequency content due to half-toning screens or dot placement.

It is not always possible to predict what form a malicious attack might take, as it will be targeted at removing the watermark itself, but we can discount methods which adversely

affect the quality of the watermarked image. The best one can do in this case is to make the watermark difficult to detect through common image processing operations, and difficult to remove without adversely affecting image quality.

This paper presents a partial description of a practical watermarking system which allows a distorted image to be rectified back to its original size and position, while using a mark which is not trivially removable.

II. OVERVIEW OF WATERMARKING SYSTEM

There are many trade-offs involved in the creation of an effective watermarking system.

These trade-offs are often presented in the form of a triangle, with the three properties of robustness, imperceptibility and information capacity all working against each other. For example, a watermark which is both extremely robust and carries a large amount of information is likely to be very visible.

We have set ourselves a goal of producing a watermarking scheme which allows the embedding 64 bits of information, which is imperceptible under normal viewing conditions, and is robust enough to attacks or distortion that damaging the watermark will also damage the image quality.

One way to improve both the perceptibility and robustness trade-off in any watermarking scheme is through the use of a *perceptual mask*, which varies the local embedding strength of a watermark in an image according to a local estimate of the perceptibility of that watermark.

While more complicated systems exist, our embedder is quite simple in concept. The embedder calculates an image watermark of the same size as the original image to be watermarked, analyzes the original to compute a mask representing the “hide-ability” of the watermark in different image regions, attenuates the watermark image by the mask image, and adds this to the original image.

Our image watermark itself consists of two components, with two independent purposes. The first mark is an *alignment pattern*, and it exists to provide a spatial reference for the recovery of the second mark. The alignment pattern is very robust to spatial deformations, affine (linear) transformations and cropping in particular. The second mark is an *information pattern*, which carries the information to be embedded.

This paper is primarily concerned with describing the generation and detection of the alignment pattern, as this pattern allows a distorted image to be rectified back to its original orientation, making the analysis of the information pattern quite straightforward.

After the watermarked image is produced, it is subject to many kinds of distortion, which will generally damage the watermark to some extent.

The decoder will process this distorted image and attempt to recover the information embedded with the original watermark. This decoder should be robust to a wide range of image distortions, and attempts to ameliorate their effect by using methods which are, to some extent, invariant under many distortions. The decoder will first attempt to undo the effect of the perceptual mask, and to thus restore the watermark to a more-or-less constant level across the original image. It is also possible at this point to attempt to undo any blurring introduced by printing and scanning the image, or by heavy compression.

However, the decoder must now attempt to detect the mark in the distorted image. Common techniques used for image matching, such as correlation and least-squares matching, are likely to be ineffective if used directly, because correlation values drop extremely quickly with both rotation and scaling. For example, a rotation of 1° or a scaling by 2% of a 512×512 image results in a drop of correlation peak magnitude of 25 times.

III. PERCEPTUAL MASKING

In some areas of an image, such as constant regions, an embedded watermark with a relatively small amplitude will be visible. In other areas of an image, such as near edges, and in highly textured regions, only an embedded watermark with a relatively large amplitude will be visible.

To reduce the visibility of a mark, and to make its visibility relatively constant over the cover image, a *perceptual mask* is used. Before embedding a watermark, the input image is analyzed to estimate the local visibility of an embedded watermark. The analysis generates the perceptual mask, which can be multiplied by the generated watermarking pattern and then added to the original image.

We use a simple scheme to calculate a perceptual mask, which is just a gradient magnitude, coupled with a 10×10 separable smoothing filter. A small constant offset may be added to the perceptual mask to ensure that some signal is embedded even in regions of an image with constant value. An example perceptual mask is shown in Figure 1.

Perceptual masking substantially distorts the original watermark signal, so that it is necessary for the decoder to undo the perceptual mask before detection of the watermark. Because the watermark is embedded at a low level, computing a perceptual mask on the watermarked image will return a similar result to computing a perceptual mask on the original image. To return the embedded watermark




Figure 1. *lena* image with perceptual mask: lighter mask values result in lower embedding strength

to a constant level, it is therefore possible to divide the watermarked image by a newly computed perceptual mask.

This process also has the advantage of reducing the energy of the cover image in regions containing higher amplitude features of higher frequency, thus providing some equalization of the total energy in each region of the image.

Another way to equalize the energy across the whole image is by a process of local spectral whitening, in which the image is divided into a grid of lapped tiles, each tile is subjected to a spectral transform (such as an FFT), the magnitude of spectral values set to a constant value, the result inverse transformed, and all tiles are summed. This process both normalizes the energy in every region of the image, and also normalizes the spectral energy from every region of the image. Such a process has two beneficial effects. Firstly, it compensates for non-linear effects of the printing process which can greatly reduce watermark amplitude in very dark or very light areas of an image. Secondly, it greatly reduces features in the cover image which have a strong spectral signature, and may interfere with the detection process.[8]

IV. SOME RST AND AFFINE-INVARIANT WATERMARKING TECHNIQUES

To be effective, a watermark really must be robust to common geometric transforms, such as rotation, spatial scaling, cropping and inevitable translation. Therefore, any watermark we embed must retain its detectability through all of these transformations.

One method for embedding an affine-invariant watermark is to use a periodic noise pattern, in which an identical noise patch is repeated multiple times across a watermark. Such patterns are easily detectable using auto-correlation, which results in a grid of magnitude peaks due to the strong self-correlation of the periodic noise patterns. An affine transformation of the watermark will not destroy the periodic structure of the tiles, but will simply change the spacing of the grid cells, so that the auto-correlation grid structure will be preserved after an affine transform, but

with the grid spacing altered to reflect the affine transform. Examination of this grid structure allows the affine transform to be estimated, allowing the watermarked image to be rectified to its original orientation.[10]

Unfortunately, this mark is inevitably vulnerable to malicious attack, because the watermark is visible as a grid in the Fourier transform or auto-correlation of the watermarked image. Such structures can easily be attacked, either by averaging tiles from the image to reconstruct an approximation of the embedded noise pattern, or by reducing the magnitude of the corresponding spectral points in the Fourier transform. When performed carefully, this kind of watermark removal may result in an attacked image that is actually closer to the original image than the watermarked one was.

Other methods have been posited which use feature points in the cover image to provide a spatial reference for embedding a watermark, such as that by Bas, Chassart and Macq.[1] which triangulates image feature points and embeds a spatial watermark in each triangle. While this kind of technique is potentially resistant to non-linear warps, the results appear to indicate little resistance to image scaling.

Pun et. al. propose a method which relies upon peaks in the Fourier magnitude spectrum to define a spatial reference. This has very good resistance to most attacks. However, a method which relies upon the presence of magnitude peaks might be easy to attack, as such peaks are readily detectable, and hence easily removed or damaged.[4]

Several watermarking techniques have been based, in one way or another, upon the use of two-dimensional Fourier Mellin basis functions. These patterns, also known as *logarithmic radial harmonic functions* (LRHFs), have been used for RST-invariant matching for a longer time than watermarking has been postulated as a problem in itself: in 1976, Casasent et al introduced the idea of RST-invariant correlation based upon a Mellin transform.[3]

The essential idea of such matching is to transform the image into a form which is RST-invariant, or, alternatively, where rotation and scaling is transformed into a simple translation, and so can be detected by conventional correlation. This form of matching has been used for recognizing military targets.

The two dimensional Fourier-Mellin transform of a function $f(r, \theta)$ can be defined by

$$M(s, m) = \int_{-\pi}^{+\pi} \int_0^{\infty} f(r, \theta) r^s e^{im\theta} r dr d\theta \quad (1)$$

The variable s is assumed to be complex, and m is an integer representing the circular harmonic order. It can be shown that a complete orthogonal sequence of functions results from setting $\Re(s) = -1$. The orthogonal two dimensional Fourier Mellin basis functions $g_{\alpha, k}(r, \theta)$ are given by

$$g_{\alpha, k}(r, \theta) = \frac{r^{i\alpha}}{r} e^{ik\theta} \quad (2)$$

It can be readily confirmed that this family of functions exhibits a wonderful scale and rotation invariant property,

$$g_{\alpha, k}(ar, \theta + \theta_0) = Ag_{\alpha, k}(r, \theta) \quad (3)$$

where A is complex constant independent of the polar coordinates (r, θ) .

The RST invariant nature of the Fourier-Mellin basis functions can be exploited in several ways.

The Fourier-Mellin transform itself can be used to watermark an image: by adding a signal in a Fourier-Mellin domain, that signal can be detected as a translation in a rotated and scaled image.[16]




Figure 2. An example LRHF

Lin et. al. propose a method in which a signal is embedded in one dimension which is a projection along the radial direction of a log-polar transform of the Fourier transform of the image. However, while their method appears very secure, it seems to have only moderate resistance to scaling, and does not appear to provide a way to embed more than a single bit of information.[15]

We previously presented a watermarking system in which these LRHFs were directly added to an image to produce a watermark, using their relative spatial positions to encode information[7][6][14]. Direct addition avoids the severe problems associated with log-polar remapping and interpolation. Because these patterns are spread-spectrum and have wide spatial support, they are not readily visible in either the spatial or the Fourier transform domain when embedded at a low level relative to the cover image. An example LRHF is shown in Figure 2.

Towards the centre of any of these patterns, the frequency approaches infinity. When embedding a pattern in an image with a rectangular pixel grid, these frequencies are not representable as they are higher than the Nyquist limit. Consequently, we simply zero the patterns at the centre to avoid inserting aliased frequencies. This places some constraints on the selection of the value of α and k : if these values are too large, then the frequency of the embedded pattern will be too high, with the excluded hole occupying too large an area and reducing the amount of signal which can be embedded. If these values are too small, then the resulting pattern will not contain a substantial content in the highest frequencies, and the resulting correlation peak

will contain too little high-frequency energy to be detectable. However, there are a large range of values which are usable as effective watermarks, and the size of the pattern can be matched to some extent with the desired image size we wish to watermark.

The advantage of such direct embedding is that the detection mechanism is extremely simple, with only a single correlation with the appropriate complex basis function being necessary to detect embedded patterns. Because correlation is a linear process, rotation of the original function will result in its multiplication by a single complex constant, and therefore the result of the correlation will also be the same up to multiplication by a single complex constant. While correlation is perfect with a complex basis function, it is almost as good when only using the real part of the basis pattern, which can be represented as the sum of two conjugate basis patterns:

$$\Re\left(\frac{r^{i\alpha}}{r} e^{ik\theta}\right) = \frac{r^{i\alpha} e^{ik\theta} + r^{-i\alpha} e^{-ik\theta}}{2r} \quad (4)$$

While these patterns have ideal correlation properties, simple correlation is not usable directly for watermarking applications. In a natural image, low frequencies predominate and will swamp any correlation from a watermarking function. A simple way to recover a good quality peak and to reduce the magnitude of the cover signal is to use *phase correlation*, in which the magnitude of the correlation image in the Fourier domain is set to a constant value of 1.0. This has the effect of normalizing the contribution from all frequencies, removing any MTF effect from a possible print and scan, and also results in a correlation peak resembling a true impulse.

To provide even more sensitivity, both for measuring peak amplitude and peak position, we process candidate peaks in our phase correlation by scaling a neighbourhood of the peak using a Fourier method and identifying the maximum value and position using in a fitted complex polynomial. This can recover a substantial loss in magnitude where a peak lies between four pixels, and also provides a measurement of spatial position to an accuracy much better than a single pixel, and reliably better than a quarter of a pixel.[11]

When a real function is used as a watermark and detected by correlation with a complex mark, the resulting delta function is reduced in magnitude by a factor of two, and an extraneous, but highly dispersed, signal introduced due to the correlation with the conjugated basis function.

Without knowledge of the α and k parameters, these marks are difficult to detect, although they could be found with a two-dimensional search of the space of these basis patterns.

However, all of the techniques based upon RST invariants have one fundamental disadvantage: their correlation properties are destroyed very quickly by anamorphic scaling. A

change in aspect ratio of only a few percent is likely to destroy watermark detectability.

V. EMBEDDING AND DETECTING THE HYPERBOLIC CHIRP FUNCTION

There is a one-dimensional function which has similar properties to LRHFs:[13]

$$f(x) = e^{i\alpha \ln(x)} \quad (5)$$

This function has logarithmic phase, and therefore a hyperbolic frequency, and is in fact called a *hyperbolic chirp*.

A spatial scaling of this function is equivalent to multiplication by a complex constant:

$$f(ax) = e^{i\alpha \ln(ax)} = e^{i\alpha \ln(a)} e^{i\alpha \ln(x)} \quad (6)$$

This function is defined only for $x > 0$, but it may be extended for negative x in many ways without losing its scale invariant property, for example by symmetric extension:

$$f(x) = e^{i\alpha \ln(|x|)} \quad (7)$$

Again, the real part of this function can be expressed as the sum of two complex functions which are conjugated relative to each other, and the function has extremely good correlation properties.

This one-dimensional hyperbolic chirp function in x can be converted into a two-dimensional function by extending it, or back-projection, along the y axis.

$$g(x, y) = f(x) \quad (8)$$




Figure 3. An extended hyperbolic chirp

An example of such a pattern is shown in Figure 3.

Nevertheless, the resulting function cannot be detected directly by correlation as can the 2D LRHFs: the result of a correlation is not a point, but a line, and any rotation of the pattern will destroy the correlation.

Instead of detecting the pattern directly in two dimensions, we can convert the detection into a one-dimensional problem by projecting the two-dimensional function $f(x, y)$ onto the x axis, $p(x)$, which will restore the pattern to one dimension,

from where it can be detected directly by one-dimensional correlation:

$$p(x) = \int_{y=-\infty}^{y=+\infty} f(x, y) dy \quad (9)$$

Translation will not affect its detectability: a translation along the y axis will leave the pattern unchanged; a translation along the x axis will simply shift the function, resulting in a shift in the position of the correlation peak.

Scaling in either x or y will also not affect its detectability: scaling in y will, again, leave the function unchanged, and scaling in x will result in the correlation function being multiplied by a complex constant, and thus will not affect the magnitude of the correlation peak.

While rotation will destroy the coherency of the projection, resulting in the correlation peak disappearing very quickly, projection along the rotated axes will again result in a reconstruction of our desired one-dimensional watermarking pattern, and it will again be detectable.

Thus, to detect our extended one-dimensional pattern under rotation, we must project the two-dimensional pattern along all angles, and it will be detectable where the projection axes line up with the axis of extension.

So, under each of the operations of translation, anamorphic scaling and rotation, these functions remain detectable by projection along the correct direction and one-dimensional correlation with the appropriate basis function. Because any affine transformation can be decomposed by singular value decomposition into these operations, the detectability of these functions is not affected by an affine transform.


Inversion of the projective transform was described initially by Funk[9] on the sphere, and, more generally, by Radon, and is now called the *Radon Transform*[18]. While using a brute-force algorithm to calculate all of these projections would be very expensive, we are fortunate that an efficient way was found by Bracewell. By the projection-slice theorem,[2] each projection to an axis with angle θ can be calculated by sampling a two-dimensional Fourier transform along a line orthogonal to θ and calculating the inverse Fourier transform of the one-dimensional result.

Thus, a complete set of projections can be obtained by computing a two-dimensional Fourier transform of a function, sampling that function along lines (slices) through the origin to produce a two-dimensional function of r and θ , and computing an inverse one-dimensional transform at all values of θ .

While this theory has been developed over continuous functions on an infinite plane, it maps directly to application when used with sampled images with bounded area: so long as these patterns are embedded to contain a substantial portion of the highest frequencies which are near to their origin, they are quite detectable by a Radon transform and

one-dimensional correlation. As with the two-dimensional patterns, the frequency of the hyperbolic chirp approaches infinity at the origin, so we set the central aliased region to zero. As with the LRHFs, the parameter α must be chosen with some care.

Figure 4a shows an example watermarking pattern consisting of four super-imposed one-dimensional chirps, and 4b shows the Radon transform of this pattern, and demonstrates that the one-dimensional patterns are easily separated. The image as shown is quasi-polar, with $r=0$ passing vertically through the centre.



(a) Four marks

(b) Quasi-polar Radon transform

Figure 4. Alignment mark and its radon transform

From a security perspective, the hyperbolic chirp function is probably less secure than the LRHF, as only a one-dimensional space must be searched to find the any embedded patterns, and the patterns are not as well distributed in the Fourier domain as LRHFs, as each embedded chirp appears one a line, instead of spread over all frequencies.

VI. USING DETECTED CHIRPS AS A SPATIAL REFERENCE

The direct embedding method using LRHFs results in a correlation peak after detection which represent a specific (x, y) positions corresponding to the origin of each embedded pattern. By embedding multiple patterns at different (x, y) positions, the relative position of any detected peaks can be used directly to determine an RST transform, for example by using the RANSAC algorithm.[5]

However, the detection peak resulting from the one-dimensional patterns does *not* correspond directly to (x, y) positions in the image. A point in the radon domain corresponds to a line in the spatial domain, so that a (r, θ) detected watermark in the radon domain corresponds directly to a line in the original image: in fact, it corresponds to the line at the centre of the extended pattern, expressed in terms of its distance from the origin r and its orientation θ .

Where several chirp functions are present, the detected lines can be used to identify specific (x, y) positions in the original image by identifying the *intersection points* of these lines. For example, where four lines are present and none are parallel, there will be six intersection points. As these points correspond directly to (x, y) positions in the image,

they can be used to determine an affine transform, again using the RANSAC algorithm, between the original and a transformed image.

One of the first problems with finding an affine transform between two sets of points is identifying a correspondence between each of the points which fits a derived model. In RANSAC, this correspondence is determined by creating random correspondences from which a candidate model is derived, and then testing this model against the remaining points. After a number of iterations, the probability of finding the correct model can be made almost certain.

However, the RANSAC algorithm is iterative, and therefore slow. There is a far superior method for determining an affine transform using an identified set of four lines using some affine-invariant properties of intersecting lines.

In a group of four mutually intersecting lines, each line is intersected by each of the other three lines, creating three co-linear intersection points for each line. These three intersection points create two line segments, and the length of this pair of line segments forms a ratio, so that with four lines, there will be four associated ratios. This ratio is fortunately invariant under affine transform, so that a set of four lines will produce a set of four ratios, invariant under affine transformation.[12]




Figure 5. Invariant ratios of four intersecting watermarks

Figure 5 shows graphically this arrangement of lines. The plus signs indicate intersection points, each (a_n, b_n) is a pair of line segments associated with each line by the intersection points, and each ratio $a_n : b_n$ is invariant under affine transformation.

Because each ratio belongs to a specific line, the four ratios can be used to uniquely identify each of the four lines, and because each of the line segments provide correspondences to three intersection points, the four ratios may be used to determine correspondences between the six intersection points in the original watermark and the detected watermark. So, assuming that the four watermarking patterns are correctly identified, the method of ratios can be used to directly determine a correspondence between the six

intersection points, from which a least-squares method can be used to determine an affine transform relating the sets of points in the original watermarked image and the detected watermark.

Aside from an increase in speed, the method of ratios can also provide a great improvement in the robustness of the watermark by providing a method for distinguishing watermark peaks in the watermark detection image from noisy peaks caused by interference with the cover image. When detecting the watermark, rather than simply identifying the four strongest peaks in the detection image, and hence the four strongest lines in the Radon transform, a larger number of candidate peaks can be detected. By computing the four ratios corresponding to each permutation of lines, the correct set of lines can be distinguished from the other lines with great probability, as the four ratio values are invariant to affine transformations, and so can be readily determined.

If n candidate peaks are detected, then the number of line quartets to be tested is simply:

$${}^n C_4 = \frac{n!}{4!(n-4)!} \quad (10)$$

With 20 peaks, around 5,000 quartets must be tested, and with 50 peaks, around a quarter of a million candidates. If the search occurs using the peaks of the largest magnitude first, and the search curtailed when a candidate group of lines is first identified, then the search can be made very fast for all images with a relatively undamaged watermark.

VII. RECTIFYING AN IMAGE FOR FAST WATERMARK DETECTION

Many images these days are produced at very high resolutions, with images containing tens of millions of pixels. However, these images are often downsampled to a much smaller size for presentation on the web or in an email.

Ideally we would wish our watermarking scheme to work with both kinds of images, and we would also wish our watermark to survive a scaling process from megapixels down to hundreds of kilopixels. This can be achieved if we simply assume the worst case, and that the high resolution image *will* be scaled to a smaller size. This assumption is actually a very liberating one, as it allows us to simply downsample a large image before detection.

So, when detecting a watermark in an image, the first step can be to simply downsample it to a reasonable size, such as between 512×512 pixels and 1024×1024 pixels.

This “lowest-common-denominator” approach also allows us to generate the watermark at a low resolution, at around 1 megapixels for example, and use interpolation to scale the watermark up to any desired size.

While it gives great robustness of the watermark detection to scaling, this does cause some limitations. Because the watermark is distributed across the whole input image, large amounts of cropping can result in losing some of the

watermark “lines” outside the image. It can also fail where a small image is watermarked and subsequently padded with a large amount of extra material. Where the padding is with some constant value, this may be detected and the padding removed, but where the image is placed within another, much larger, image, this larger image will be substantially down-sized before detection commences, and the smaller image may be downsized past the point of detectability.

Once the watermark has been detected and the affine transform determined, the inverse affine transform may be applied to the image to restore it to its original size and dimensions. However, if an information watermark is to be detected from the rectified image, the affine transformation can be altered by scaling to ensure that it does not make the input image any larger. This serves two purposes: firstly, as there will be no information content in the high frequencies of a larger image, there will be no need to analyze a larger image; secondly, the absence of these high frequencies will be actively detrimental in detecting any subsequent watermark, so it is better to stick with a smaller image in which the watermark may be present at all frequencies.

An alternative approach is to leave the input image unrectified, and guide the information watermark detection by the derived affine transform.

VIII. RESULTS

While this watermarking technique appears to have many attractive properties, the proof of the method is in the pudding.

Sadly, it is not possible for us to use the standard benchmarking suite, *Stirmark*, but we have recreated many of the early *stirmark* tests using our own applications.[17]

We watermarked two common test images, *lena*, a 512×512 image, and *Woman*, a 2048×2560 image, as shown in Figure 7. The watermark was embedded using perceptual masking, and included both an alignment portion, using the hyperbolic chirp, and a data carrying component, including 64 bits of information and 20 bits of checksum. The strength of the embedded watermark had a mean absolute value of approximately 1.5 grey levels, and the standard deviation was 1.7 grey levels with *lena* (PSNR=43.5dB), and 2.8 grey levels with *Woman* (PSNR=39.2dB). The nominal embedding strength for both images was the same, but the watermark was embedded more strongly in the *Woman* image because this image is somewhat sharper than *lena*, so the perceptual mask is consequently of a larger value.

The table in Figure 6 shows all of the tests we performed in attempting to read the 64 bits embedded in an altered image. † is used to indicate the tests which failed with *lena*, and ‡ the tests which failed with *Woman*.

For the large image, *Woman*, only three tests failed. Cropping the image by 75% removed too much of the alignment mark, so no watermark could be recovered. Scaling the image to 5% resulted in an image of size 102×128 , which




Figure 7. Two test images: 512×512 *lena* and 2048×2560 *Woman*

was too small for the watermark to be recovered. The warp test combined a single horizontal cycle of a sinusoid to warp the x coordinate combined with a single vertical sinusoid in y , and was survived up to a magnitude of 5 pixels, but failed at 10.

For the smaller image, *lena*, more tests failed. As with *Woman*, scaling to a size of 128×128 failed, as did cropping by 75%. The median filters, and jpeg compression with a quality of 60% or below, probably failed because of the reduction in high frequency information in the watermark. These did not cause problems in the larger image because the damaged high frequencies were removed by the initial downsampling step. The warp test failed with a magnitude of 5 pixels.

IX. CONCLUSION

This paper has presented a new method for watermarking images with an alignment watermark that is inherently robust to affine transforms, usable over a very large range of image scales, and moderately difficult to attack. It is also robust to a wide range of common image processing operations, although we would require access to a more modern benchmarking suite to make definitive statements about its robustness.

We have not attempted to describe the method used to embed information alongside our alignment watermark, but hopefully this may form the basis of a later paper.

While the method requires use of computationally intensive Fourier transforms, Radon transforms, interpolation and correlation, these operations can always be performed on an downsampled image of a size which is easy to process using current hardware.

More work needs to be done to demonstrate the commercial value of this method: we require access to a more modern watermark benchmarking suite, more experiments need to be performed with human observers to assess the method’s imperceptibility, and optimized embedding and detection methods need to be developed.

Ratio x 0.80 y 1.00	Scale to 5%†‡	Crop 1%	Rotation -0.25°	Rotation+Scale -0.25°
Ratio x 0.90 y 1.00	Scale to 10%†	Crop 2%	Rotation -0.50°	Rotation+Scale -0.50°
Ratio x 1.00 y 0.80	Scale to 25%†	Crop 5%	Rotation -0.75°	Rotation+Scale -0.75°
Ratio x 1.00 y 0.90	Scale to 50%	Crop 10%	Rotation -1.00°	Rotation+Scale -1.00°
Ratio x 1.00 y 1.10	Scale to 75%	Crop 15%	Rotation -2.00°	Rotation+Scale -2.00°
Ratio x 1.00 y 1.20	Scale to 90%	Crop 20%	Rotation 0.25°	Rotation+Scale 0.25°
Ratio x 1.10 y 1.00	Scale to 110%	Crop 25%	Rotation 0.50°	Rotation+Scale 0.50°
Ratio x 1.20 y 1.00	Scale to 150%	Crop 50%	Rotation 0.75°	Rotation+Scale 0.75°
1 Row 1 Col removed	Scale to 200%	Crop 75%†‡	Rotation 1.00°	Rotation+Scale 1.00°
1 Row 5 Col removed	Scale to 500%	Jpeg Quality 10%†	Rotation 10.00°	Rotation+Scale 10.00°
5 Row 1 Col removed	Scale to 1000%	Jpeg Quality 20%†	Rotation 15.00°	Rotation+Scale 15.00°
5 Row 17 Col removed	Linear 1.007 0.010 0.010 1.012	Jpeg Quality 30%†	Rotation 2.00°	Rotation+Scale 2.00°
17 Row 5 Col removed	Linear 1.010 0.013 0.009 1.011	Jpeg Quality 40%†	Rotation 30.00°	Rotation+Scale 30.00°
Gaussian filtering 3×3	Linear 1.013 0.008 0.011 1.008	Jpeg Quality 50%†	Rotation 45.00°	Rotation+Scale 45.00°
Sharpening 3×3 kernel	Shear x 0% y 1%	Jpeg Quality 60%†	Rotation 5.00°	Rotation+Scale 5.00°
Median 2×2	Shear x 0% y 5%	Jpeg Quality 70%	Rotation 90.00°	Rotation+Scale 90.00°
Median 3×3 †	Shear x 1% y 0%	Jpeg Quality 80%	Warp 1 pixel	
Median 4×4 †	Shear x 1% y 1%	Jpeg Quality 90%	Warp 2 pixels	
Transpose	Shear x 5% y 0%		Warp 5 pixels†	
	Shear x 5% y 5%		Warp 10 pixels†‡	

Figure 6. Watermarking Tests. Most tests succeeded; † and ‡ respectively mark failed tests in *lena* and *Woman*

ACKNOWLEDGMENT

The authors would like to thank Canon Information Systems Australia for supporting this work and consenting to its publication.

REFERENCES

- [1] Patrick Bas, Jean-Marc Chassery, and Benoît Macq, *Geometrically Invariant Watermarking Using Feature Points*, IEEE Transactions on Image Processing, Vol. 11, No. 9, p. 1014, 2002
- [2] R. N. Bracewell, *Strip Integration in Radio Astronomy*, Australian Journal of Physics, Vol. 9, p. 198, 1956
- [3] David Casasent and Demetri Psaltis, *Position, rotation and scale invariant optical correlation*, Applied Optics, Vol. 15, No. 7, p. 1795, 1976
- [4] Frédéric Deguillaume, Sviatoslav Voloshynovskiy, and Thierry Pun, *A method for the estimation and recovering from general affine transforms in digital watermarking applications*, University of Geneva
- [5] M. A. Fischer and R. C. Bolles, *Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography*, Communications of the ACM, Vol 24, p. 381, 1981
- [6] Peter A. Fletcher and Kieran G. Larkin, *Method for generating and detecting marks*, Patent, US7031493, filed 2001
- [7] Peter A. Fletcher and Kieran G. Larkin, *Direct Embedding and Detection of RST Invariant Watermarks*, 5th International Workshop on Information Hiding, Noordwijkerhout, The Netherlands, 2002
- [8] Peter A. Fletcher, *Local phase filter to assist correlation*, Patent, US7221774, filed 2003
- [9] Paul Funk, *Über eine geometrische anwendung der abelschen integralgleichung*, Math. Ann., Vol. 77, p. 128, 1916
- [10] Chris Honsinger and Majid Rabbani, *Data Embedding Using Phase Dispersion*, Eastman Kodak, 2000
- [11] Kieran G. Larkin and Peter A. Fletcher, *Method for the enhancement of complex peaks*, Patent, US7430301B2, filed 2002
- [12] Kieran G. Larkin and Peter A. Fletcher, *Method of estimating an affine relation between images*, Patent, US7532768B2, filed 2004
- [13] Kieran G. Larkin, Peter A. Fletcher and Stephen J. Hardy, *Mark embedding and detection using projective transforms*, Patent, US7313249, filed 2003
- [14] Kieran G. Larkin and Peter A. Fletcher, *Encoding information in a watermark*, Patent, US7158653, filed 2002
- [15] C-Y. Lin, M. Wu, J. A. Bloom, I. J. Cox and Y. M. Lui, *Rotation, Scale, and Translation Resilient Watermarking for Images*, IEEE Transactions on Image Processing, Vol. 10, No. 5, p.767, 2001
- [16] O'Ruanaidh, Jjko and Pun, T., *Rotation, Scale and Translation Invariant Spread Spectrum Digital Image Watermarking*, Signal Processing, Vol. 66, No. 3, p. 303, 1998
- [17] Fabien A. P. Petitcolas, Ross J. Anderson, Markus G. Kuhn. *Attacks on copyright marking systems*, in David Aucsmith (Ed), Information Hiding, Second International Workshop, Portland, Oregon, U.S.A., 1998
- [18] J. Radon, *Über die bestimmung von funktionen durch ihre intergralwerte langsgewisser mannigfaltigkeiten (on the determination of functions from their integrals along certain manifolds)*, Berichte Saechsische Akademie der Wissenschaften, Vol. 29, p. 262, 1917