



(19) **United States**

(12) **Patent Application Publication**  
**Rhoads**

(10) **Pub. No.: US 2004/0001608 A1**

(43) **Pub. Date: Jan. 1, 2004**

(54) **IMAGE PROCESSOR AND IMAGE PROCESSING METHOD**

(76) Inventor: **Geoffrey B. Rhoads**, West Linn, OR (US)

Correspondence Address:  
**DIGIMARC CORPORATION**  
19801 SW 72ND AVENUE  
SUITE 100  
TUALATIN, OR 97062 (US)

(21) Appl. No.: **10/336,650**

(22) Filed: **Jan. 2, 2003**

**Related U.S. Application Data**

(63) Continuation-in-part of application No. 09/626,984, filed on Jul. 27, 2000, now Pat. No. 6,542,620, which is a continuation of application No. 08/967,693, filed on Nov. 12, 1997, now Pat. No. 6,122,392, which is a continuation of application No. 08/614,521, filed on Mar. 15, 1996, now Pat. No. 5,745,604, which is a continuation of application No. 08/215,289, filed on Mar. 17, 1994, now abandoned, which is a continuation-in-part of application No. 08/154,866, filed on Nov. 18, 1993, now abandoned. Continuation-in-part of application No. 09/526,982, filed on Mar. 15, 2000, now Pat. No. 6,516,079, which is a continuation-in-part of application No.

09/186,962, filed on Nov. 5, 1998, which is a continuation of application No. 08/649,419, filed on May 16, 1996, now Pat. No. 5,862,260, which is a continuation-in-part of application No. PCT/US96/06618, filed on May 7, 1996, and which is a continuation of application No. 08/635,531, filed on Apr. 22, 1996, now abandoned, and which is a continuation of application No. 08/534,005, filed on Sep. 25, 1995, now Pat. No. 5,832,119, and which is a continuation of application No. 08/436,102, filed on May 8, 1995, now Pat. No. 5,748,783.

Continuation-in-part of application No. 10/202,367, filed on Jul. 22, 2002, which is a continuation of application No. 09/566,533, filed on May 8, 2000, now Pat. No. 6,424,725, which is a continuation-in-part of application No. 09/452,023, filed on Nov. 30, 1999, now Pat. No. 6,408,082.

**Publication Classification**

(51) **Int. Cl.<sup>7</sup>** ..... **G06K 9/00**  
(52) **U.S. Cl.** ..... **382/100**

(57) **ABSTRACT**

A pattern signal is steganographically encoded in a content object to aid in identifying distortion (e.g., image rotation) to which the object may be subjected. The signal may include a feature that can be used to unambiguously identify correct orientation.

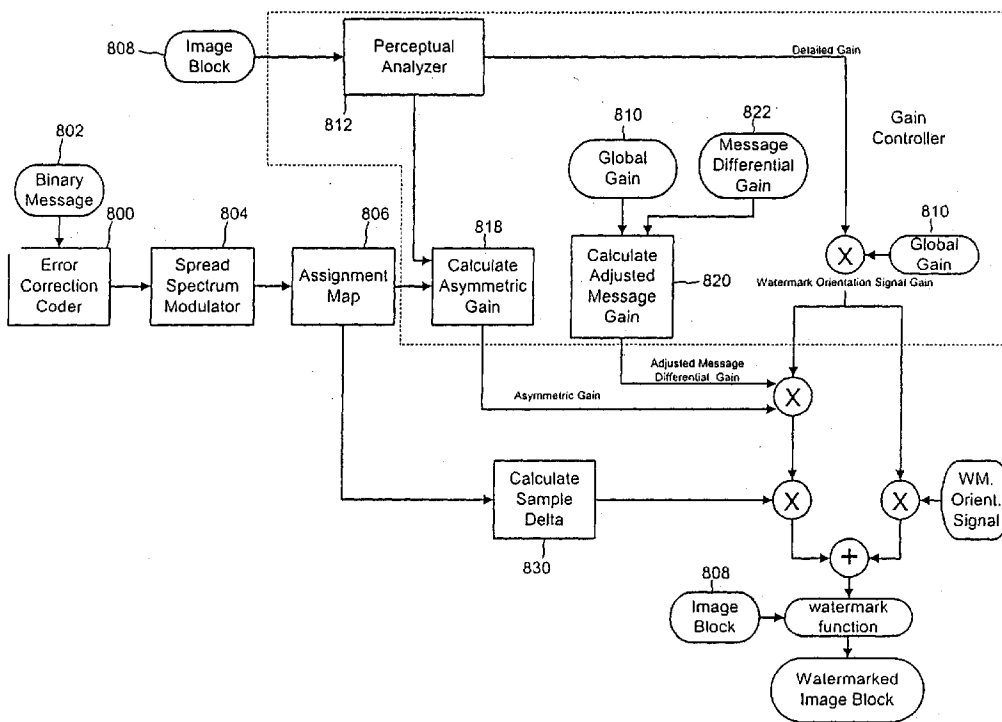


Fig. 1

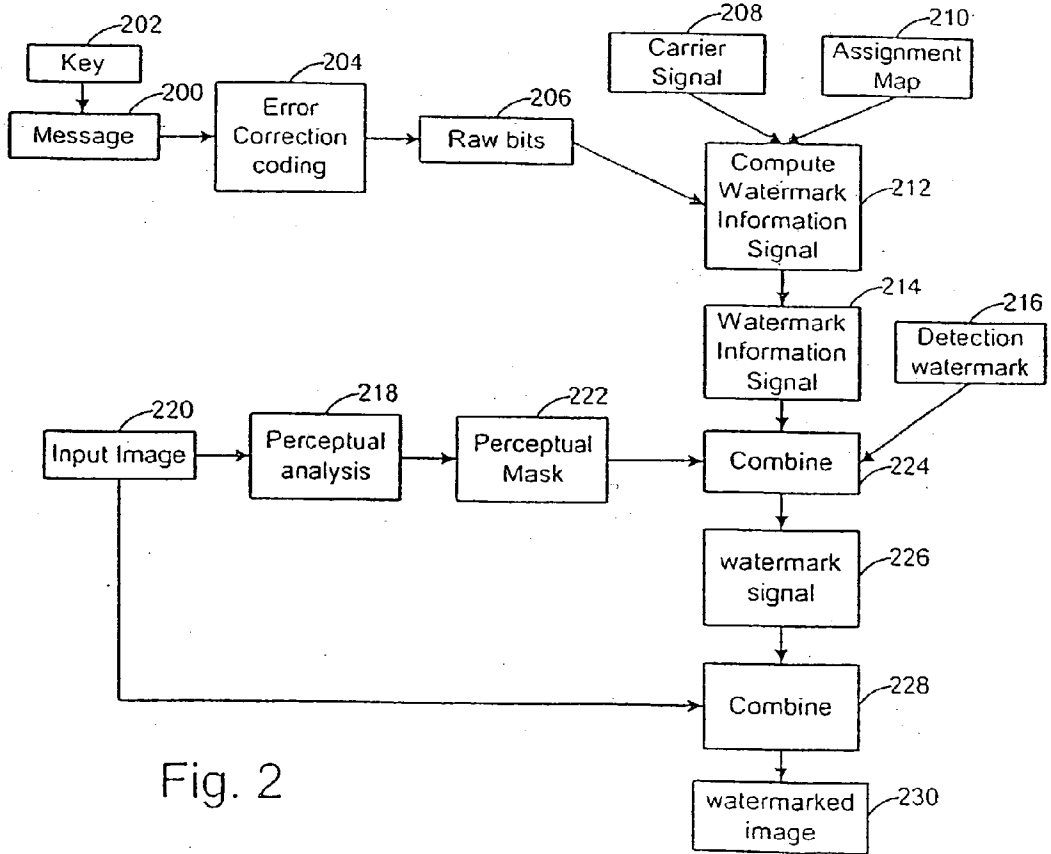



Fig. 2

One Quadrant Spatial  
Transform Domain

Fig. 3




Fig. 4


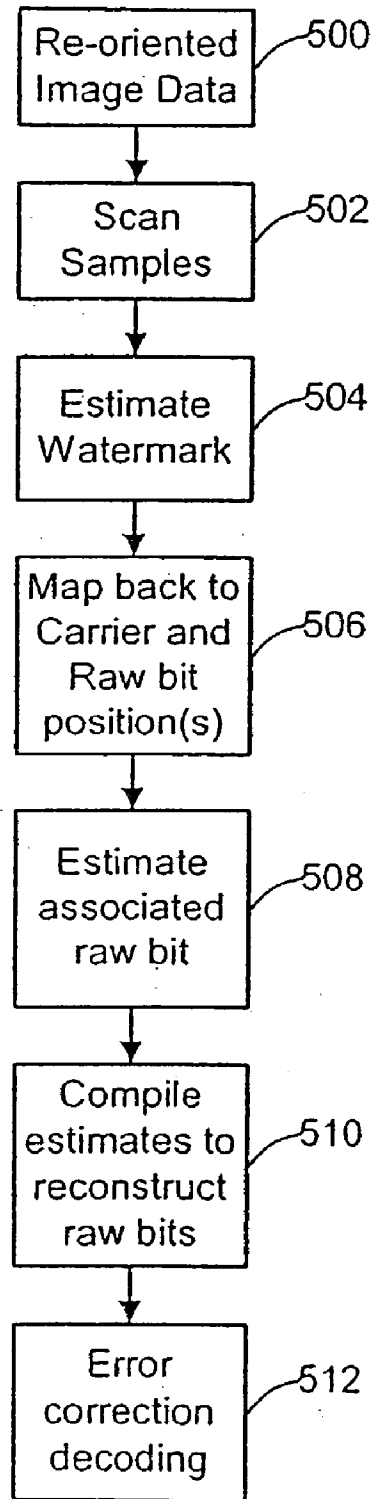


Fig. 5







FIG. 6

Fig. 7




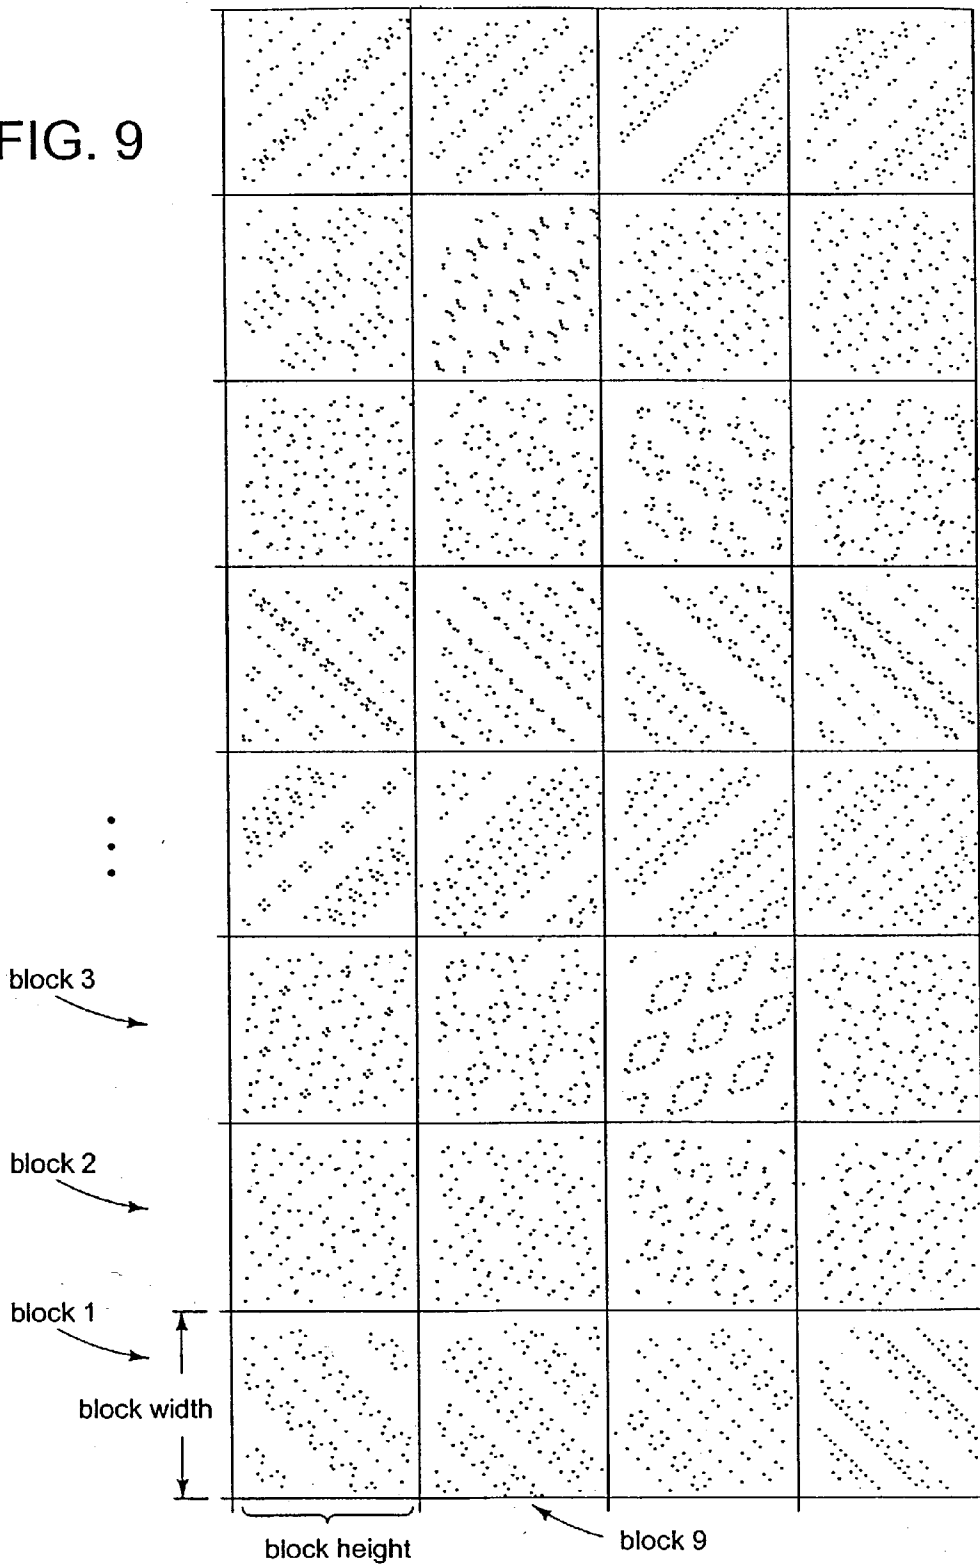


Fig. 8

FIG. 9





MAGNITUDE OF WATERMARK ORIENTATION SIGNAL  
IN FREQUENCY DOMAIN




FIG. 10

REAL VALUE OF WATERMARK ORIENTATION SIGNAL  
IN SPATIAL DOMAIN




FIG. 11

Fig. 12

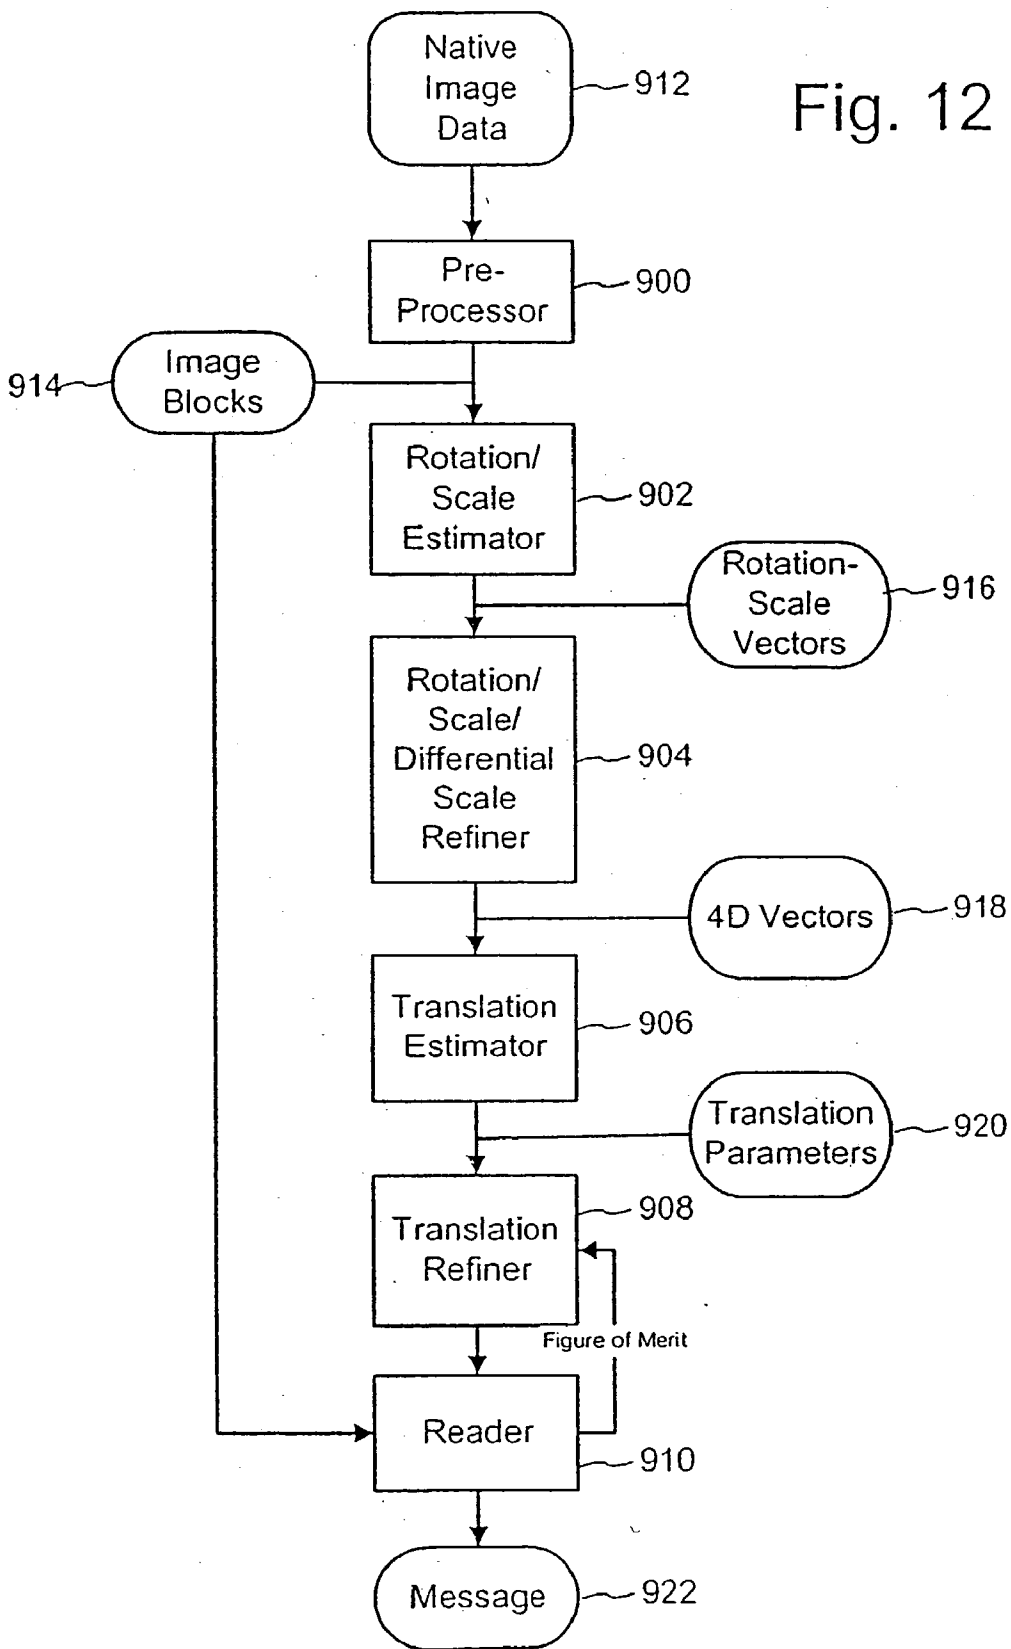


Fig. 13

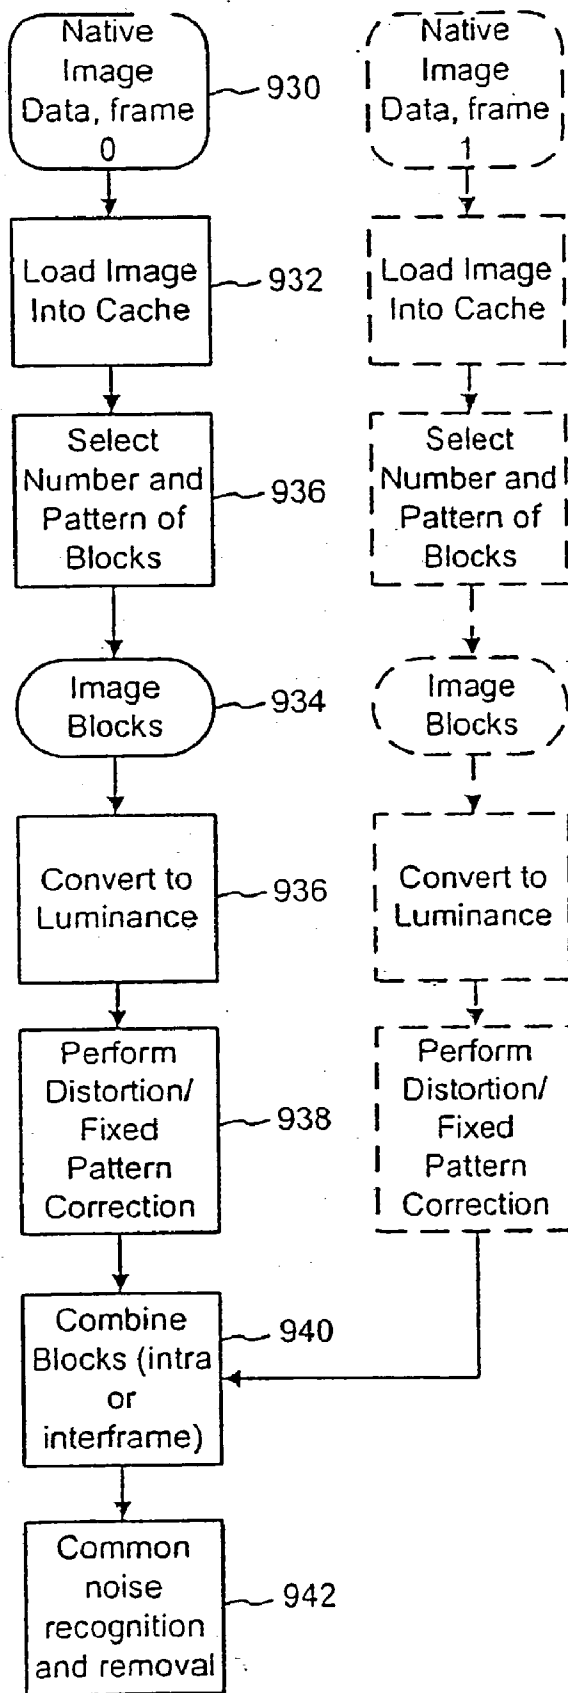



Fig. 14



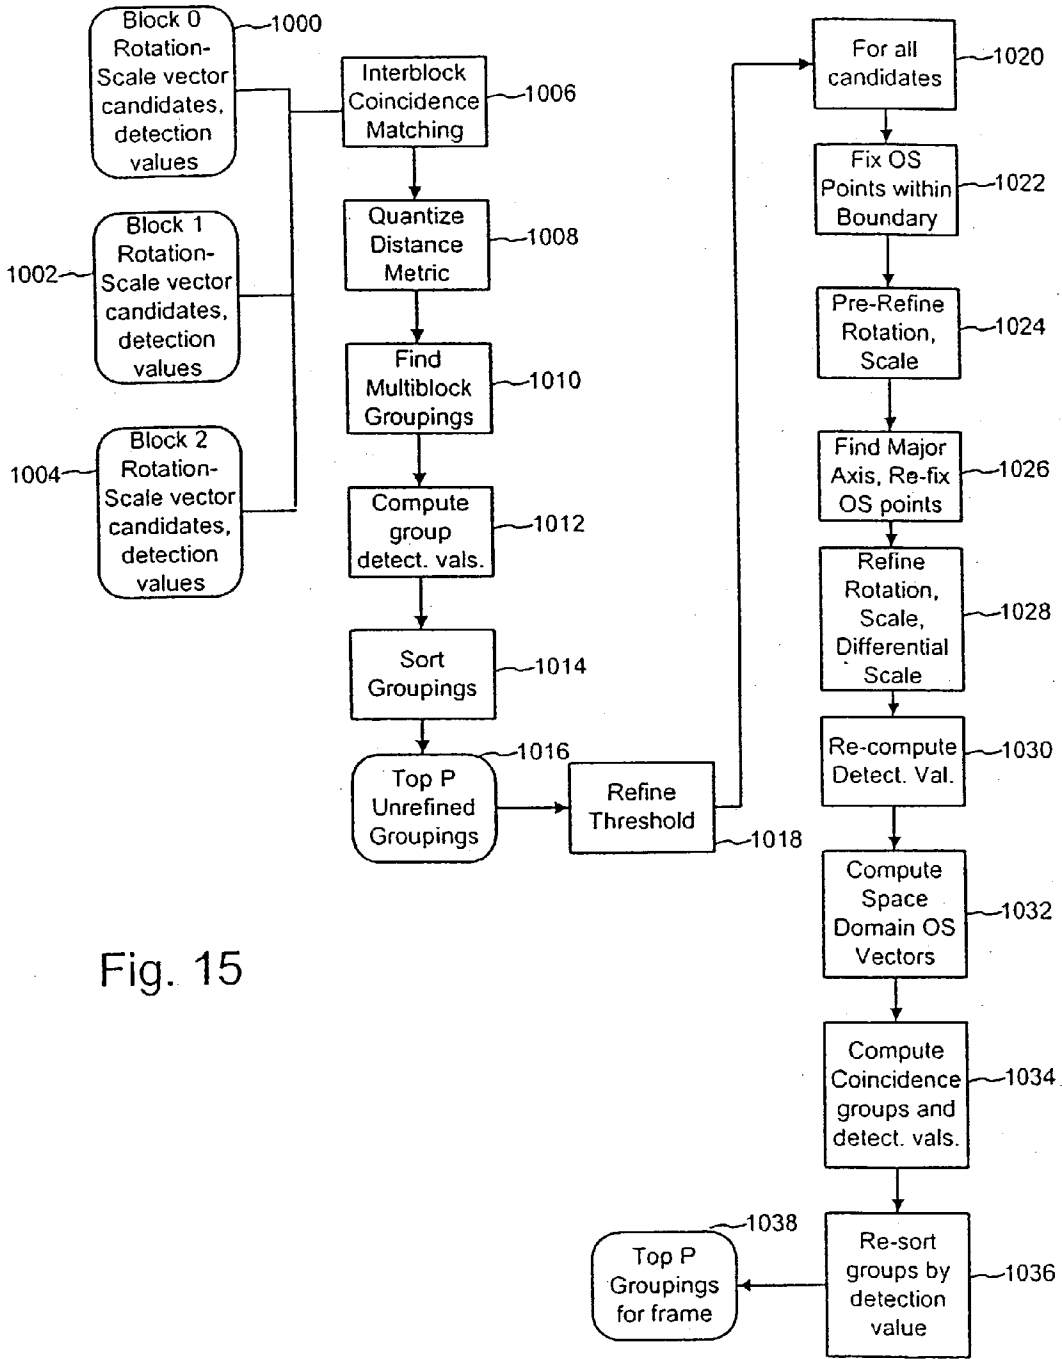
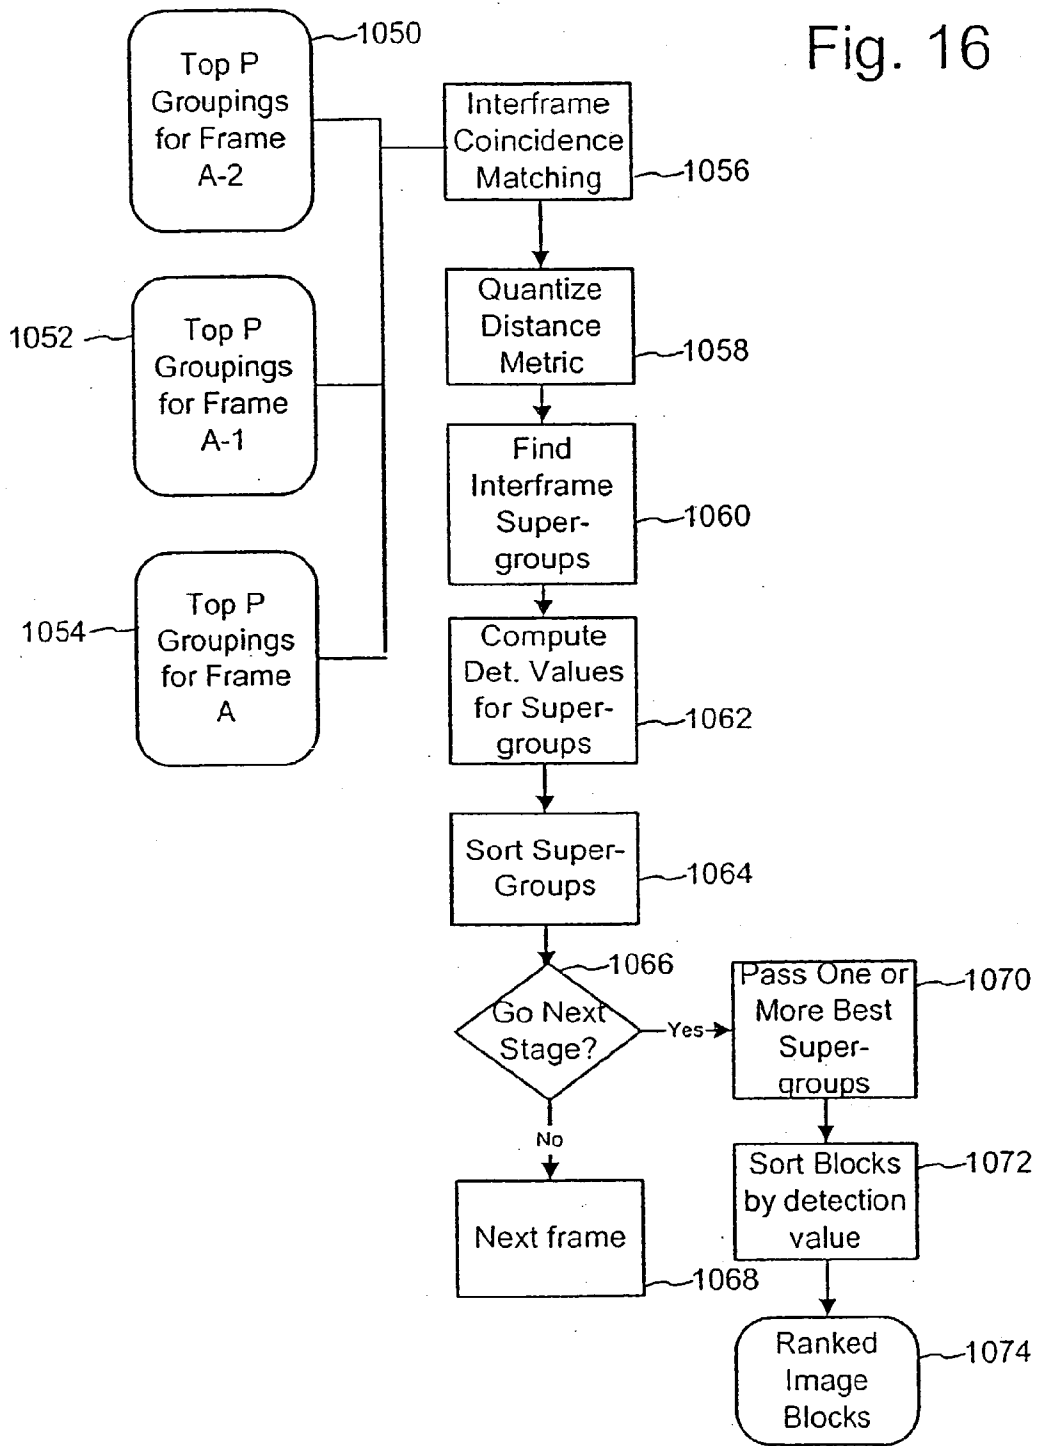


Fig. 15

Fig. 16



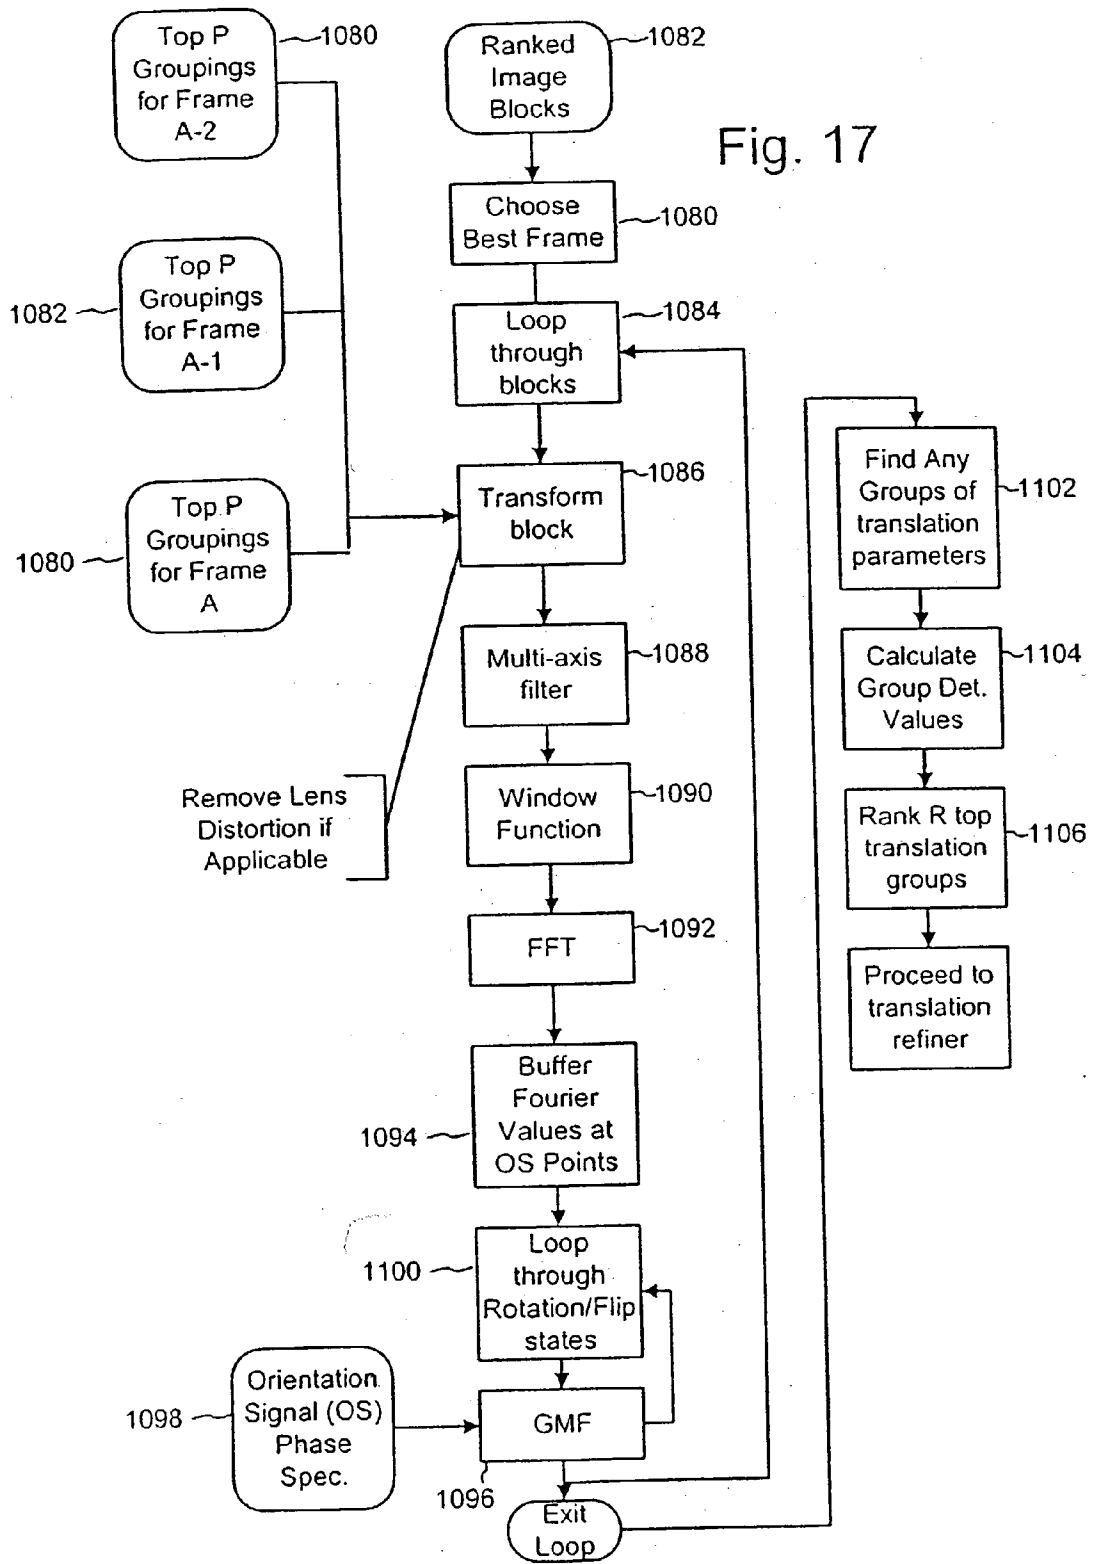


Fig. 18

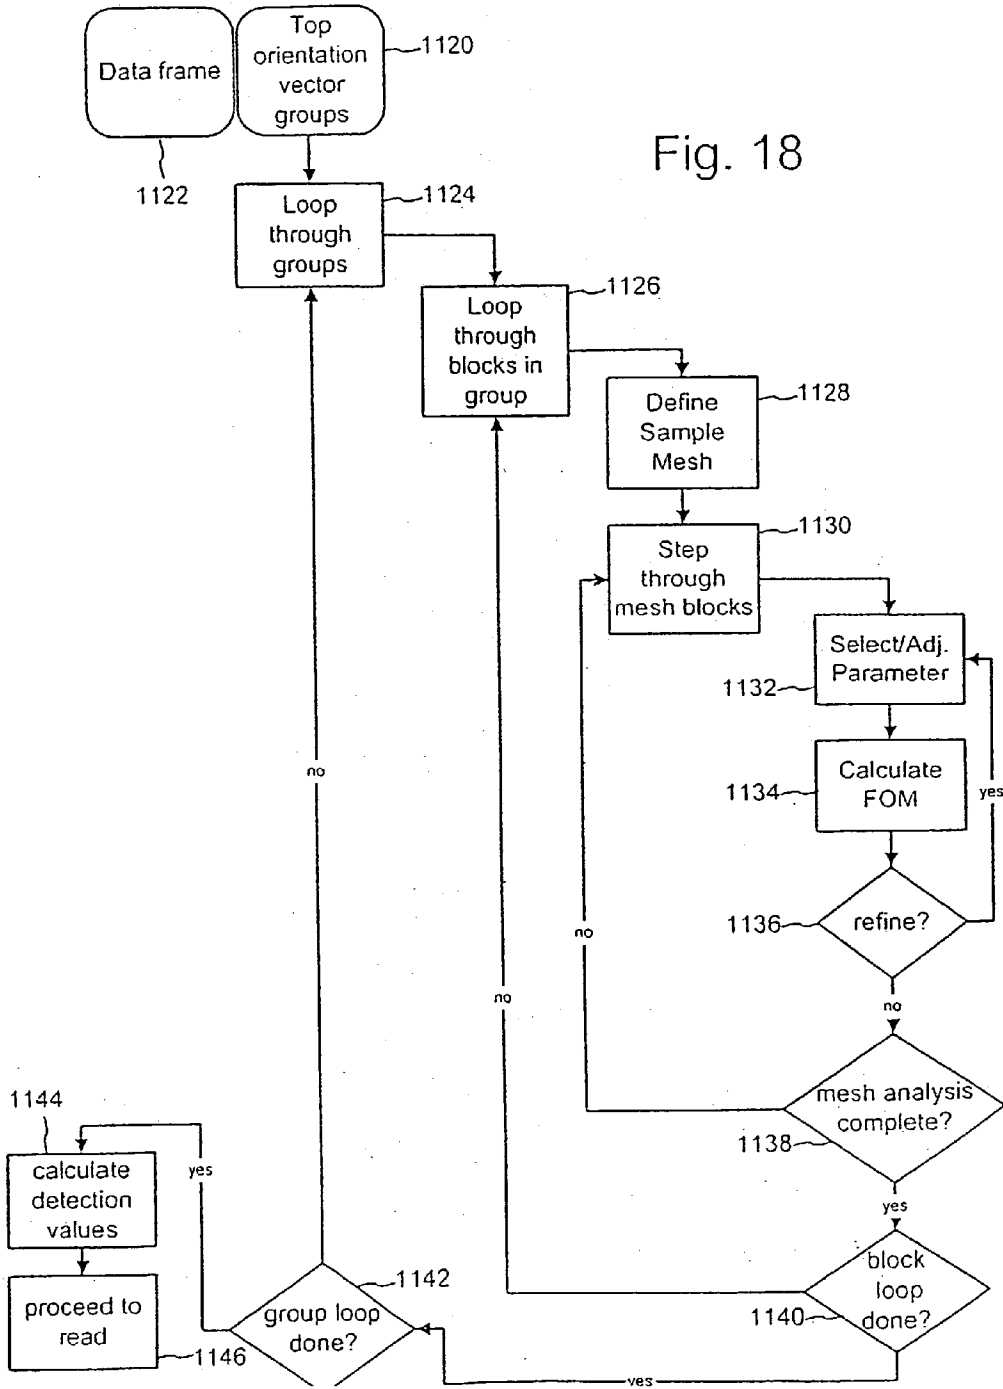




Fig. 19

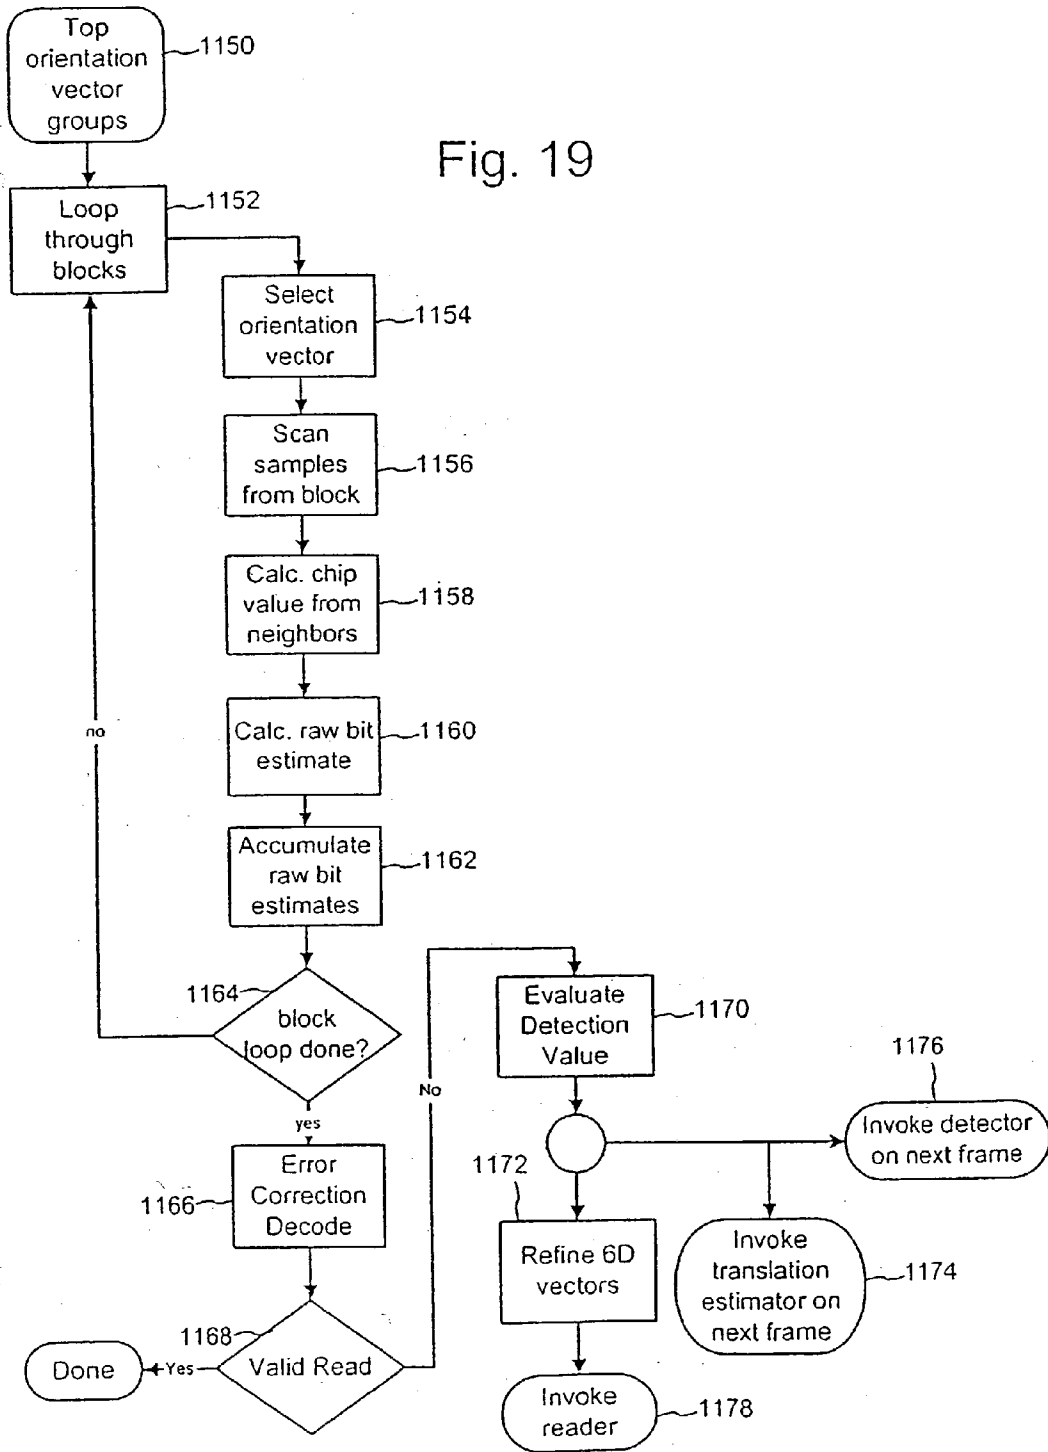
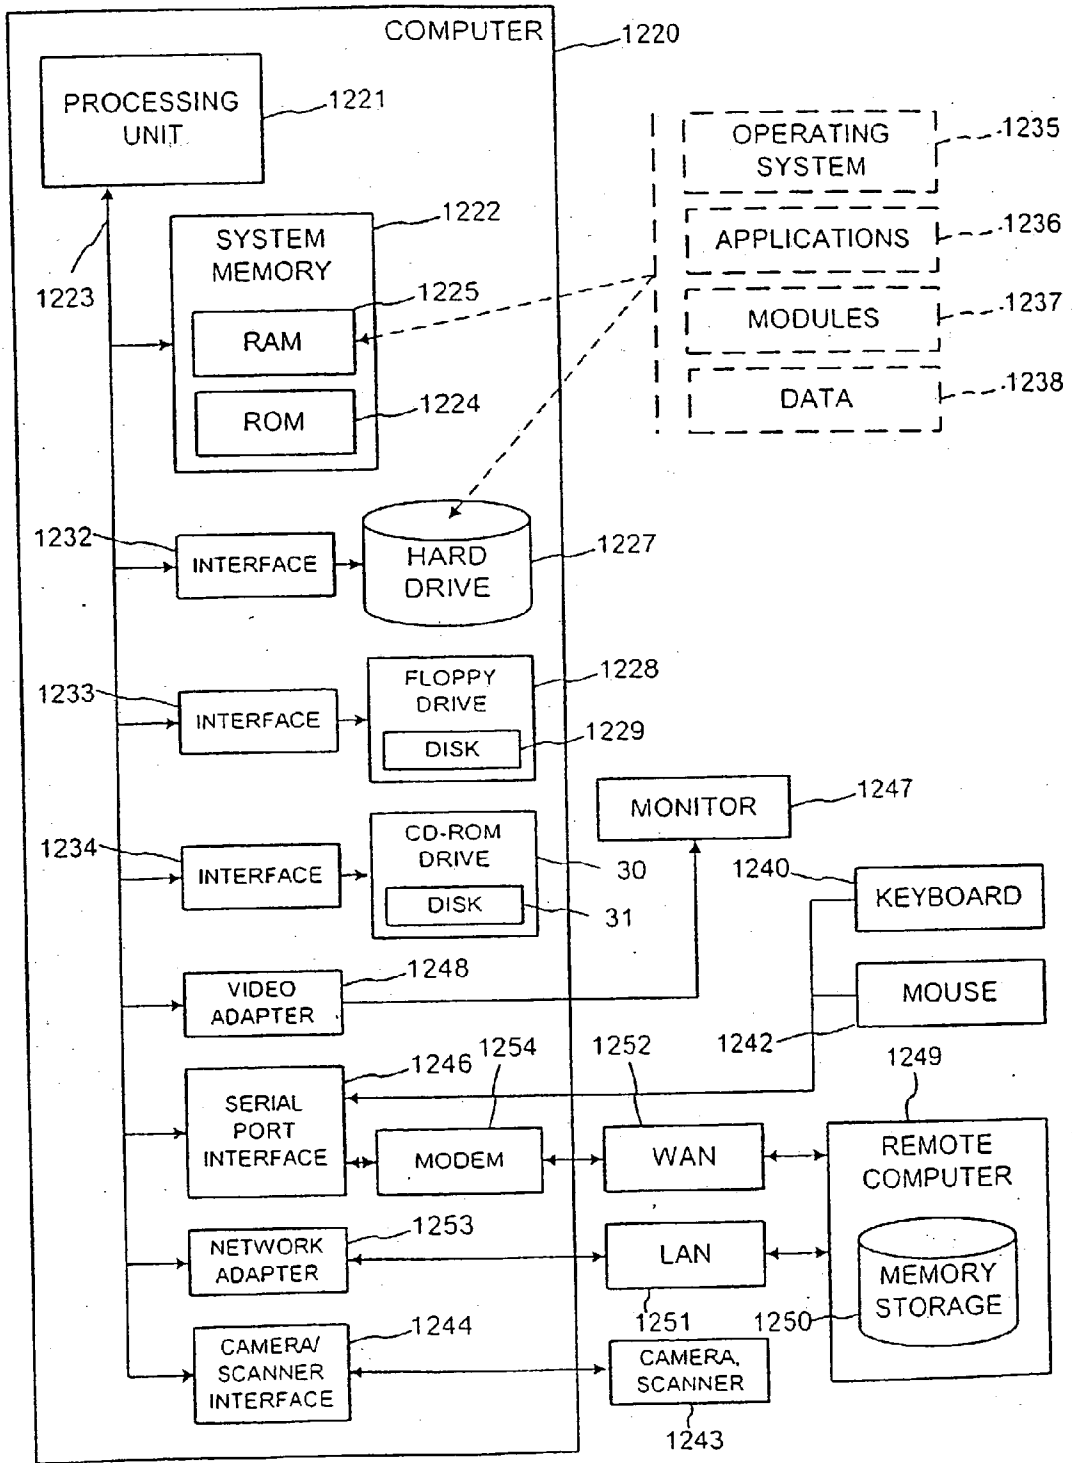


FIG. 20







FIG. 21



QUEST FOR MOSAICED KNOT PATTERNS WHICH "COVER" AND  
ARE COEXTENSIVE WITH ORIGINAL IMAGE; ALL ELEMENTAL  
KNOT PATTERNS CAN CONVEY THE SAME INFORMATION, SUCH  
AS A SIGNATURE, OR EACH CAN CONVEY A NEW MESSAGE IN A  
STEGANOGRAPHIC SENSE

FIG. 22






FIG. 23



NON-HARMONIC SPATIAL FREQUENCIES ALONG THE 45 DEGREE AXES, GIVING RISE TO A WEAVE-LIKE CROSS-HATCHING PATTERN IN THE SPATIAL DOMAIN

FIG. 24



NON-HARMONIC CONCENTRIC CIRCLES IN UV PLANE, WHERE PHASE HOPS QUASI-RANDOMLY ALONG EACH CIRCLE, GIVING RISE TO PSEUDO RANDOM LOOKING PATTERNS IN THE SPATIAL DOMAIN

FIG. 25






FIG. 24A




PHASE OF SPATIAL FREQUENCIES ALONG FORWARD 45 DEGREE AXES, 1008

FIG. 26A




PHASE OF SPATIAL FREQUENCIES ALONG BACKWARD 45 DEGREE AXES, 1010

FIG. 26B



PHASE OF SPATIAL FREQUENCIES ALONG FIRST CONCENTRIC RING, 1012

FIG. 27A



PHASE OF SPATIAL FREQUENCIES ALONG SECOND CONCENTRIC RING, 1014

FIG. 27B



PHASE OF SPATIAL FREQUENCIES ALONG THIRD CONCENTRIC RING, 1016

FIG. 27C



FIG. 28A




FIG. 28B






FIG. 28C



POWER PROFILE ALONG ANGLE A, AS NORMALIZED BY ITS OWN MOVING AVERAGE; ONLY A MINIMAL AMOUNT EXCEEDS THRESHOLD, GIVING A SMALL INTEGRATED VALUE

FIG. 28D



POWER PROFILE ALONG ANGLE B, AS NORMALIZED BY ITS OWN MOVING AVERAGE; THIS FINDS STRONG ENERGY ABOVE THE THRESHOLD

FIG. 28E




FIG. 29A




FIG. 29B




FIG. 29C




FIG. 29D




FIG. 29E






FIG. 30A




SCALE = A: ADD LL POWER VALUES AT THE "KNOWN" FREQUENCIES", 1042

FIG. 30B




SCALE = B: ADD LL POWER VALUES AT THE "KNOWN" FREQUENCIES", 1044

FIG. 30C




"SCALED-KERNEL" BASED MATCHED FILTER;  
PEAK IS WHERE THE SCALE OF THE SUBLIMINAL  
GRID WAS FOUND, 1046

FIG. 31A



ARBITRARY ORIGINAL IMAGE  
IN WHICH SUBLIMINAL GRATICULES  
MAY HAVE BEEN PLACED

FIG. 31B



"COLUMN SCAN" IS APPLIED  
ALONG A GIVEN ANGLE THROUGH  
THE CENTER OF THE IMAGE

FIG. 31C





FIG. 31D



PROCESS STEPS

1.    SCAN IN PHOTOGRAPH
2.    2D FFT
3.    GENERATE 2D POWER SPECTRUM, FILTER WITH E.G. 3X3 BLURRING  
      KERNEL
4.    STEP ANGLES FROM 0 DEGREES THROUGH 90 (1/2 DEG)
5.    GENERATE NORMALIZED VECTOR, WITH POWER VALUE AS  
      NUMERATOR, AND MOVING AVERAGED POWER VALUE AS  
      DENOMINATOR
6.    INTEGRATE VALUES AS SOME THRESHOLD, GIVING A SINGLE  
      INTEGRATED VALUE FOR THIS ANGLE
7.    END STEP ON ANGLES
8.    FIND TOP ONE OR TWO OR THREE "PEAKS" FROM THE ANGLES IN  
      LOOP 4, THEN FOR EACH PEAK...
9.    STEP SCALE FROM 25% TO 400%, STEP ~ 1.01
10.    ADD THE NORMALIZED POWER VALUES CORRESPONDING TO THE  
      'N' SCALED FREQUENCIES OF STANDARD
11.    KEEP TRACK OF HIGHEST VALUE IN LOOP
12.    END LOOP 9 AND 8, DETERMINE HIGHEST VALUE
13.    ROTATION AND SCALE NOW FOUND
14.    PERFORM TRADITIONAL MATCHED FILTER TO FIND EXACT SPATIAL  
      OFFSET
15.    PERFORM ANY "FINE TUNNING" TO PRECISELY DETERMINE  
      ROTATION, SCALE, OFFSET

FIG. 32

## IMAGE PROCESSOR AND IMAGE PROCESSING METHOD

### FIELD OF THE INVENTION

[0001] The present invention relates to digital watermarking, and more particularly relates to identifying rotation of watermarked content to aid in decoding.

### BACKGROUND AND SUMMARY OF THE INVENTION

[0002] Digital watermarking is a process for modifying media content to embed a machine-readable code into the data content. The data may be modified such that the embedded code is imperceptible or nearly imperceptible to the user, yet may be detected through an automated detection process. Most commonly, digital watermarking is applied to media such as images, audio signals, and video signals. However, it may also be applied to other types of data, including documents (e.g., through line, word or character shifting), software, multi-dimensional graphics models, and surface textures of objects.

[0003] Digital watermarking systems have two primary components: an embedding component that embeds the watermark in the media content, and a reading component that detects and reads the embedded watermark. The embedding component embeds a watermark pattern by altering data samples of the media content. The reading component analyzes content to detect whether a watermark pattern is present. In applications where the watermark encodes information, the reader extracts this information from the detected watermark.

[0004] One challenge to the developers of watermark embedding and reading systems is to ensure that the watermark is detectable even if the watermarked media content is transformed in some fashion. The watermark may be corrupted intentionally, so as to bypass its copy protection or anti-counterfeiting functions, or unintentionally through various transformations that result from routine manipulation of the content. In the case of watermarked images, such manipulation of the image may distort the watermark pattern embedded in the image.

[0005] The invention provides watermark structures, and related embedders, detectors, and readers for processing the watermark structures. In addition, it provides a variety of methods and applications associated with the watermark structures, embedders, detectors and readers. While adapted for images, the watermark system applies to other electronic and physical media. For example, it can be applied to electronic objects, including image, audio and video signals. It can be applied to mark blank paper, film and other substrates, and it can be applied by texturing object surfaces for a variety of applications, such as identification, authentication, etc. The detector and reader can operate on a signal captured from a physical object, even if that captured signal is distorted.

[0006] The watermark structure can have multiple components, each having different attributes. To name a few, these attributes include function, signal intensity, transform domain of watermark definition (e.g., temporal, spatial, frequency, etc.), location or orientation in host signal, redundancy, level of security (e.g., encrypted or scrambled). When

describing a watermark signal in the context of this document, intensity refers to an embedding level while strength describes reading level (though the terms are sometimes used interchangeably). The components of the watermark structure may perform the same or different functions. For example, one component may carry a message, while another component may serve to identify the location or orientation of the watermark in a combined signal. Moreover, different messages may be encoded in different temporal or spatial portions of the host signal, such as different locations in an image or different time frames of audio or video.

[0007] Watermark components may have different signal intensities. For example, one component may carry a longer message, yet have smaller signal intensity than another component, or vice-versa. The embedder may adjust the signal intensity by encoding one component more redundantly than others, or by applying a different gain to the components. Additionally, watermark components may be defined in different transform domains. One may be defined in a frequency domain, while another may be defined in a spatial or temporal domain.

[0008] The watermark components may be located in different spatial or temporal locations in the host signal. In images, for example, different components may be located in different parts of the image. Each component may carry a different message or perform a different function. In audio or video, different components may be located in different time frames of the signal.

[0009] The watermark components may be defined, embedded and extracted in different domains. Examples of domains include spatial, temporal and frequency domains. A watermark may be defined in a domain by specifying how it alters the host signal in that domain to effect the encoding of the watermark component. A frequency domain component alters the signal in the frequency domain, while a spatial domain component alters the signal in the spatial domain. Of course, such alterations may have an impact that extends across many transform domains.

[0010] While described here as watermark components, one can also construe the components to be different watermarks. This enables the watermark technology described throughout this document to be used in applications using two or more watermarks. For example, some copy protection applications of the watermark structure may use two or more watermarks, each performing similar or different function. One mark may be more fragile than another, and thus, disappear when the combined signal is corrupted or transformed in some fashion. The presence or lack of a watermark or watermark component conveys information to the detector to initiate or prohibit some action, such as playback, copying or recording of the marked signal.

[0011] A watermark system may include an embedder, detector, and reader. The watermark embedder encodes a watermark signal in a host signal to create a combined signal. The detector looks for the watermark signal in a potentially corrupted version of the combined signal, and computes its orientation. Finally, a reader extracts a message in the watermark signal from the combined signal using the orientation to approximate the original state of the combined signal.

[0012] There are a variety of alternative embodiments of the embedder and detector. One embodiment of the embed-

der performs error correction coding of a binary message, and then combines the binary message with a carrier signal to create a component of a watermark signal. It then combines the watermark signal with a host signal. To facilitate detection, it may also add a detection component to form a composite watermark signal having a message and detection component. The message component includes known or signature bits to facilitate detection, and thus, serves a dual function of identifying the mark and conveying a message. The detection component is designed to identify the orientation of the watermark in the combined signal, but may carry an information signal as well. For example, the signal values at selected locations in the detection component can be altered to encode a message.

[0013] One embodiment of the detector estimates an initial orientation of a watermark signal in the multidimensional signal, and refines the initial orientation to compute a refined orientation. As part of the process of refining the orientation, this detector computes at least one orientation parameter that increases correlation between the watermark signal and the multidimensional signal when the watermark or multidimensional signal is adjusted with the refined orientation.

[0014] Another detector embodiment computes orientation parameter candidates of a watermark signal in different portions of the target signal, and compares the similarity of orientation parameter candidates from the different portions. Based on this comparison, it determines which candidates are more likely to correspond to a valid watermark signal.

[0015] Yet another detector embodiment estimates orientation of the watermark in a target signal suspected of having a watermark. The detector then uses the orientation to extract a measure of the watermark in the target. It uses the measure of the watermark to assess merits of the estimated orientation. In one implementation, the measure of the watermark is the extent to which message bits read from the target signal match with expected bits. Another measure is the extent to which values of the target signal are consistent with the watermark signal. The measure of the watermark signal provides information about the merits of a given orientation that can be used to find a better estimate of the orientation.

[0016] In various embodiments, a feature of the embedded signal can be used to unambiguously identify correct orientation.

[0017] Further advantages and features of the invention will become apparent with reference to the following detailed description and accompanying drawings.

#### BRIEF DESCRIPTION OF THE DRAWINGS

[0018] FIG. 1 is a block diagram illustrating an image watermark system.

[0019] FIG. 2 is a block diagram illustrating an image watermark embedder.

[0020] FIG. 3 is a spatial frequency domain plot of a detection watermark signal.

[0021] FIG. 4 is a flow diagram of a process for detecting a watermark signal in an image and computing its orientation within the image.

[0022] FIG. 5 is a flow diagram of a process reading a message encoded in a watermark.

[0023] FIG. 6 is a diagram depicting an example of a watermark detection process.

[0024] FIG. 7 is a diagram depicting the orientation of a transformed image superimposed over the original orientation of the image at the time of watermark encoding.

[0025] FIG. 8 is a diagram illustrating an implementation of a watermark embedder.

[0026] FIG. 9 is a diagram depicting an assignment map used to map raw bits in a message to locations within a host image.

[0027] FIG. 10 illustrates an example of a watermark orientation signal in a spatial frequency domain.

[0028] FIG. 11 illustrates the orientation signal shown in FIG. 10 in the spatial domain.

[0029] FIG. 12 is a diagram illustrating an overview of a watermark detector implementation.

[0030] FIG. 13 is a diagram illustrating an implementation of the detector pre-processor depicted generally in FIG. 12.

[0031] FIG. 14 is a diagram illustrating a process for estimating rotation and scale vectors of a detection watermark signal.

[0032] FIG. 15 is a diagram illustrating a process for refining the rotation and scale vectors, and for estimating differential scale parameters of the detection watermark signal.

[0033] FIG. 16 is a diagram illustrating a process for aggregating evidence of the orientation signal and orientation parameter candidates from two or more frames.

[0034] FIG. 17 is a diagram illustrating a process for estimating translation parameters of the detection watermark signal.

[0035] FIG. 18 is a diagram illustrating a process for refining orientation parameters using known message bits in the watermark message.

[0036] FIG. 19 is a diagram illustrating a process for reading a watermark message from an image, after re-orienting the image data using an orientation vector.

[0037] FIG. 20 is a diagram of a computer system that serves as an operating environment for software implementations of a watermark embedder, detector and reader.

[0038] FIGS. 21-23 show details relating to embodiments of the present invention using ring-based patterns.

[0039] FIGS. 24 and 25 illustrate two embodiments by which subliminal digital graticules can be realized.

[0040] FIG. 24A shows a variation on the FIG. 24 embodiment.

[0041] FIGS. 26A and 26B show the phase of spatial frequencies along two inclined axes.

[0042] FIGS. 27A-27C show the phase of spatial frequencies along first, second and third concentric rings.

[0043] FIGS. 28A-28E show steps in the registration process for a subliminal graticule using inclined axes.

[0044] FIGS. 29A-29E show steps in the registration process for a subliminal graticule using concentric rings.

[0045] FIGS. 30A-30C shows further steps in the registration process for a subliminal graticule using inclined axes.

[0046] FIGS. 31A-31 D show another registration process that does not require a 2D FFT.

[0047] FIG. 32 is a flow chart summarizing a registration process for subliminal graticules.

#### DETAILED DESCRIPTION

[0048] 1.0 Introduction

[0049] A watermark can be viewed as an information signal that is embedded in a host signal, such as an image, audio, or some other media content. Watermarking systems based on the following detailed description may include the following components: 1) An embedder that inserts a watermark signal in the host signal to form a combined signal; 2) A detector that determines the presence and orientation of a watermark in a potentially corrupted version of the combined signal; and 3) A reader that extracts a watermark message from the combined signal. In some implementations, the detector and reader are combined.

[0050] The structure and complexity of the watermark signal can vary significantly, depending on the application. For example, the watermark may be comprised of one or more signal components, each defined in the same or different domains. Each component may perform one or more functions. Two primary functions include acting as an identifier to facilitate detection and acting as an information carrier to convey a message. In addition, components may be located in different spatial or temporal portions of the host signal, and may carry the same or different messages.

[0051] The host signal can vary as well. The host is typically some form of multi-dimensional media signal, such as an image, audio sequence or video sequence. In the digital domain, each of these media types is represented as a multi-dimensional array of discrete samples. For example, a color image has spatial dimensions (e.g., its horizontal and vertical components), and color space dimensions (e.g., YUV or RGB). Some signals, like video, have spatial and temporal dimensions. Depending on the needs of a particular application, the embedder may insert a watermark signal that exists in one or more of these dimensions.

[0052] In the design of the watermark and its components, developers are faced with several design issues such as: the extent to which the mark is impervious to jamming and manipulation (either intentional or unintentional); the extent of imperceptibility; the quantity of information content; the extent to which the mark facilitates detection and recovery, and the extent to which the information content can be recovered accurately.

[0053] For certain applications, such as copy protection or authentication, the watermark should be difficult to tamper with or remove by those seeking to circumvent it. To be robust, the watermark must withstand routine manipulation, such as data compression, copying, linear transformation, flipping, inversion, etc., and intentional manipulation intended to remove the mark or make it undetectable. Some applications require the watermark signal to remain robust through digital to analog conversion (e.g., printing an image

or playing music), and analog to digital conversion (e.g., scanning the image or digitally sampling the music). In some cases, it is beneficial for the watermarking technique to withstand repeated watermarking.

[0054] A variety of signal processing techniques may be applied to address some or all of these design considerations. One such technique is referred to as spreading. Sometimes categorized as a spread spectrum technique, spreading is a way to distribute a message into a number of components (chips), which together make up the entire message. Spreading makes the mark more impervious to jamming and manipulation, and makes it less perceptible.

[0055] Another category of signal processing technique is error correction and detection coding. Error correction coding is useful to reconstruct the message accurately from the watermark signal. Error detection coding enables the decoder to determine when the extracted message has an error.

[0056] Another signal processing technique that is useful in watermark coding is called scattering. Scattering is a method of distributing the message or its components among an array of locations in a particular transform domain, such as a spatial domain or a spatial frequency domain. Like spreading, scattering makes the watermark less perceptible and more impervious to manipulation.

[0057] Yet another signal processing technique is gain control. Gain control is used to adjust the intensity of the watermark signal. The intensity of the signal impacts a number of aspects of watermark coding, including its perceptibility to the ordinary observer, and the ability to detect the mark and accurately recover the message from it.

[0058] Gain control can impact the various functions and components of the watermark differently. Thus, in some cases, it is useful to control the gain while taking into account its impact on the message and orientation functions of the watermark or its components. For example, in a watermark system described below, the embedder calculates a different gain for orientation and message components of an image watermark.

[0059] Another useful tool in watermark embedding and reading is perceptual analysis. Perceptual analysis refers generally to techniques for evaluating signal properties based on the extent to which those properties are (or are likely to be) perceptible to humans (e.g., listeners or viewers of the media content). A watermark embedder can take advantage of a Human Visual System (HVS) model to determine where to place a watermark and how to control the intensity of the watermark so that chances of accurately recovering the watermark are enhanced, resistance to tampering is increased, and perceptibility of the watermark is reduced. Such perceptual analysis can play an integral role in gain control because it helps indicate how the gain can be adjusted relative to the impact on the perceptibility of the mark. Perceptual analysis can also play an integral role in locating the watermark in a host signal. For example, one might design the embedder to hide a watermark in portions of a host signal that are more likely to mask the mark from human perception.

[0060] Various forms of statistical analyses may be performed on a signal to identify places to locate the watermark, and to identify places where to extract the watermark.



[0061] For example, a statistical analysis can identify portions of a host image that have noise-like properties that are likely to make recovery of the watermark signal difficult. Similarly, statistical analyses may be used to characterize the host signal to determine where to locate the watermark.

[0062] Each of the techniques may be used alone, in various combinations, and in combination with other signal processing techniques.

[0063] In addition to selecting the appropriate signal processing techniques, the developer is faced with other design considerations. One consideration is the nature and format of the media content. In the case of digital images, for example, the image data is typically represented as an array of image samples. Color images are represented as a set of color vectors in a color space, such as RGB or YUV. The watermark may be embedded in one or more of the color components of an image. In some implementations, the embedder may transform the input image into a target color space, and then proceed with the embedding process in that color space.

[0064] 2.0 Digital Watermark Embedder and Reader Overview

[0065] The following sections describe implementations of a watermark embedder and reader that operate on digital signals. The embedder encodes a message into a digital signal by modifying its sample values such that the message is imperceptible to the ordinary observer in output form. To extract the message, the reader captures a representation of the signal suspected of containing a watermark and then processes it to detect the watermark and decode the message.

[0066] FIG. 1 is a block diagram summarizing signal processing operations involved in embedding and reading a watermark. There are three primary inputs to the embedding process: the original, digitized signal 100, the message 102, and a series of control parameters 104. The control parameters may include one or more keys. One key or set of keys may be used to encrypt the message. Another key or set of keys may be used to control the generation of a watermark carrier signal or a mapping of information bits in the message to positions in a watermark information signal.

[0067] The carrier signal or mapping of the message to the host signal may be encrypted as well. Such encryption may increase security by varying the carrier or mapping for different components of the watermark or watermark message. Similarly, if the watermark or watermark message is redundantly encoded throughout the host signal, one or more encryption keys can be used to scramble the carrier or signal mapping for each instance of the redundantly encoded watermark. This use of encryption provides one way to vary the encoding of each instance of the redundantly encoded message in the host signal. Other parameters may include control bits added to the message, and watermark signal attributes (e.g., orientation or other detection patterns) used to assist in the detection of the watermark.

[0068] Apart from encrypting or scrambling the carrier and mapping information, the embedder may apply different, and possibly unique carrier or mapping for different components of a message, for different messages, or from different watermarks or watermark components to be embedded in the host signal. For example, one watermark may be encoded in a block of samples with one carrier, while

another, possibly different watermark, is encoded in a different block with a different carrier. A similar approach is to use different mappings in different blocks of the host signal.

[0069] The watermark embedding process 106 converts the message to a watermark information signal. It then combines this signal with the input signal and possibly another signal (e.g., an orientation pattern) to create a watermarked signal 108. The process of combining the watermark with the input signal may be a linear or non-linear function. Examples of watermarking functions include:  $S^* = S + gX$ ;  $S^* = S(1 + gX)$ ; and  $S^* = S e^{gX}$ , where  $S^*$  is the watermarked signal vector,  $S$  is the input signal vector, and  $g$  is a function controlling watermark intensity. The watermark may be applied by modulating signal samples  $S$  in the spatial, temporal or some other transform domain.

[0070] To encode a message, the watermark encoder analyzes and selectively adjusts the host signal to give it attributes that correspond to the desired message symbol or symbols to be encoded. There are many signal attributes that may encode a message symbol, such as a positive or negative polarity of signal samples or a set of samples, a given parity (odd or even), a given difference value or polarity of the difference between signal samples (e.g., a difference between selected spatial intensity values or transform coefficients), a given distance value between watermarks, a given phase or phase offset between different watermark components, a modulation of the phase of the host signal, a modulation of frequency coefficients of the host signal, a given frequency pattern, a given quantizer (e.g., in Quantization Index Modulation) etc.

[0071] Some processes for combining the watermark with the input signal are termed non-linear, such as processes that employ dither modulation, modify least significant bits, or apply quantization index modulation. One type of non-linear modulation is where the embedder sets signal values so that they have some desired value or characteristic corresponding to a message symbol. For example, the embedder may designate that a portion of the host signal is to encode a given bit value. It then evaluates a signal value or set of values in that portion to determine whether they have the attribute corresponding to the message bit to be encoded. Some examples of attributes include a positive or negative polarity, a value that is odd or even, a checksum, etc. For example, a bit value may be encoded as a one or zero by quantizing the value of a selected sample to be even or odd. As another example, the embedder might compute a checksum or parity of an  $N$  bit pixel value or transform coefficient and then set the least significant bit to the value of the checksum or parity. Of course, if the signal already corresponds to the desired message bit value, it need not be altered. The same approach can be extended to a set of signal samples where some attribute of the set is adjusted as necessary to encode a desired message symbol. These techniques can be applied to signal samples in a transform domain (e.g., transform coefficients) or samples in the temporal or spatial domains.

[0072] Quantization index modulation techniques employ a set of quantizers. In these techniques, the message to be transmitted is used as an index for quantizer selection. In the decoding process, a distance metric is evaluated for all quantizers and the index with the smallest distance identifies the message value.

[0073] The watermark detector **110** operates on a digitized signal suspected of containing a watermark. As depicted generally in **FIG. 1**, the suspect signal may undergo various transformations **112**, such as conversion to and from an analog domain, cropping, copying, editing, compression/decompression, transmission etc. Using parameters **114** from the embedder (e.g., orientation pattern, control bits, key(s)), it performs a series of correlation or other operations on the captured image to detect the presence of a watermark. If it finds a watermark, it determines its orientation within the suspect signal.

[0074] Using the orientation, if necessary, the reader **116** extracts the message. Some implementations do not perform correlation, but instead, use some other detection process or proceed directly to extract the watermark signal. For instance in some applications, a reader may be invoked one or more times at various temporal or spatial locations in an attempt to read the watermark, without a separate pre-processing stage to detect the watermark's orientation.

[0075] Some implementations require the original, un-watermarked signal to decode a watermark message, while others do not. In those approaches where the original signal is not necessary, the original un-watermarked signal can still be used to improve the accuracy of message recovery. For example, the original signal can be removed, leaving a residual signal from which the watermark message is recovered. If the decoder does not have the original signal, it can still attempt to remove portions of it (e.g., by filtering) that are expected not to contain the watermark signal.

[0076] Watermark decoder implementations use known relationships between a watermark signal and a message symbol to extract estimates of message symbol values from a signal suspected of containing a watermark. The decoder has knowledge of the properties of message symbols and how and where they are encoded into the host signal to encode a message. For example, it knows how message bit values of one and a zero are encoded and it knows where these message bits are originally encoded. Based on this information, it can look for the message properties in the watermarked signal. For example, it can test the watermarked signal to see if it has attributes of each message symbol (e.g., a one or zero) at a particular location and generate a probability measure as an indicator of the likelihood that a message symbol has been encoded. Knowing the approximate location of the watermark in the watermarked signal, the reader implementation may compare known message properties with the properties of the watermarked signal to estimate message values, even if the original signal is unavailable. Distortions to the watermarked signal and the host signal itself make the watermark difficult to recover, but accurate recovery of the message can be enhanced using a variety of techniques, such as error correction coding, watermark signal prediction, redundant message encoding, etc.

[0077] One way to recover a message value from a watermarked signal is to perform correlation between the known message property of each message symbol and the watermarked signal. If the amount of correlation exceeds a threshold, for example, then the watermarked signal may be assumed to contain the message symbol. The same process can be repeated for different symbols at various locations to

extract a message. A symbol (e.g., a binary value of one or zero) or set of symbols may be encoded redundantly to enhance message recovery.

[0078] In some cases, it is useful to filter the watermarked signal to remove aspects of the signal that are unlikely to be helpful in recovering the message and/or are likely to interfere with the watermark message. For example, the decoder can filter out portions of the original signal and another watermark signal or signals. In addition, when the original signal is unavailable, the reader can estimate or predict the original signal based on properties of the watermarked signal. The original or predicted version of the original signal can then be used to recover an estimate of the watermark message. One way to use the predicted version to recover the watermark is to remove the predicted version before reading the desired watermark. Similarly, the decoder can predict and remove un-wanted watermarks or watermark components before reading the desired watermark in a signal having two or more watermarks.

#### [0079] 2.1 Image Watermark Embedder

[0080] **FIG. 2** is a block diagram illustrating an implementation of an exemplary embedder in more detail. The embedding process begins with the message **200**. As noted above, the message is binary number suitable for conversion to a watermark signal. For additional security, the message, its carrier, and the mapping of the watermark to the host signal may be encrypted with an encryption key **202**. In addition to the information conveyed in the message, the embedder may also add control bit values ("signature bits") to the message to assist in verifying the accuracy of a read operation. These control bits, along with the bits representing the message, are input to an error correction coding process **204** designed to increase the likelihood that the message can be recovered accurately in the reader.

[0081] There are several alternative error correction coding schemes that may be employed. Some examples include BCH, convolution, Reed Solomon and turbo codes. These forms of error correction coding are sometimes used in communication applications where data is encoded in a carrier signal that transfers the encoded data from one place to another. In the digital watermarking application discussed here, the raw bit data is encoded in a fundamental carrier signal.

[0082] In addition to the error correction coding schemes mentioned above, the embedder and reader may also use a Cyclic Redundancy Check (CRC) to facilitate detection of errors in the decoded message data.

[0083] The error correction coding function **204** produces a string of bits, termed raw bits **206**, that are embedded into a watermark information signal. Using a carrier signal **208** and an assignment map **210**, the illustrated embedder encodes the raw bits in a watermark information signal **212**, **214**. In some applications, the embedder may encode a different message in different locations of the signal. The carrier signal may be a noise image. For each raw bit, the assignment map specifies the corresponding image sample or samples that will be modified to encode that bit.

[0084] The embedder depicted in **FIG. 2** operates on blocks of image data (referred to as 'tiles') and replicates a watermark in each of these blocks. As such, the carrier signal and assignment map both correspond to an image block of

a pre-determined size, namely, the size of the tile. To encode each bit, the embedder applies the assignment map to determine the corresponding image samples in the block to be modified to encode that bit. Using the map, it finds the corresponding image samples in the carrier signal. For each bit, the embedder computes the value of image samples in the watermark information signal as a function of the raw bit value and the value(s) of the corresponding samples in the carrier signal.

[0085] To illustrate the embedding process further, it is helpful to consider an example. First, consider the following background. Digital watermarking processes are sometimes described in terms of the transform domain in which the watermark signal is defined. The watermark may be defined in the spatial or temporal domain, or some other transform domain such as a wavelet transform, Discrete Cosine Transform (DCT), Discrete Fourier Transform (DFT), Hadamard transform, Hartley transform, Karhunen-Loeve transform (KLT) domain, etc.

[0086] Consider an example where the watermark is defined in a transform domain (e.g., a frequency domain such as DCT, wavelet or DFT). The embedder segments the image in the spatial domain into rectangular tiles and transforms the image samples in each tile into the transform domain. For example in the DCT domain, the embedder segments the image into N by N blocks and transforms each block into an N by N block of DCT coefficients. In this example, the assignment map specifies the corresponding sample location or locations in the frequency domain of the tile that correspond to a bit position in the raw bits. In the frequency domain, the carrier signal looks like a noise pattern. Each image sample in the frequency domain of the carrier signal is used together with a selected raw bit value to compute the value of the image sample at the location in the watermark information signal.

[0087] Now consider an example where the watermark is defined in the spatial domain. The embedder segments the image in the spatial domain into rectangular tiles of image samples (i.e. pixels). In this example, the assignment map specifies the corresponding sample location or locations in the tile that correspond to each bit position in the raw bits. In the spatial domain, the carrier signal looks like a noise pattern extending throughout the tile. Each image sample in the spatial domain of the carrier signal is used together with a selected raw bit value to compute the value of the image sample at the same location in the watermark information signal.

[0088] With this background, the embedder proceeds to encode each raw bit in the selected transform domain as follows. It uses the assignment map to look up the position of the corresponding image sample (or samples) in the carrier signal. The image sample value at that position in the carrier controls the value of the corresponding position in the watermark information signal. In particular, the carrier sample value indicates whether to invert the corresponding watermark sample value. The raw bit value is either a one or zero. Disregarding for a moment the impact of the carrier signal, the embedder adjusts the corresponding watermark sample upward to represent a one, or downward to represent a zero. Now, if the carrier signal indicates that the corresponding sample should be inverted, the embedder adjusts the watermark sample downward to represent a one, and

upward to represent a zero. In this manner, the embedder computes the value of the watermark samples for a raw bit using the assignment map to find the spatial location of those samples within the block.

[0089] From this example, a number of points can be made. First, the embedder may perform a similar approach in any other transform domain. Second, for each raw bit, the corresponding watermark sample or samples are some function of the raw bit value and the carrier signal value. The specific mathematical relationship between the watermark sample, on one hand, and the raw bit value and carrier signal, on the other, may vary with the implementation. For example, the message may be convolved with the carrier, multiplied with the carrier, added to the carrier, or applied based on another non-linear function. Third, the carrier signal may remain constant for a particular application, or it may vary from one message to another. For example, a secret key may be used to generate the carrier signal. For each raw bit, the assignment map may define a pattern of watermark samples in the transform domain in which the watermark is defined. An assignment map that maps a raw bit to a sample location or set of locations (i.e. a map to locations in a frequency or spatial domain) is just one special case of an assignment map for a transform domain. Fourth, the assignment map may remain constant, or it may vary from one message to another. In addition, the carrier signal and map may vary depending on the nature of the underlying image. In sum, there many possible design choices within the implementation framework described above.

[0090] The embedder depicted in **FIG. 2** combines another watermark component, shown as the detection watermark **216**, with the watermark information signal to compute the final watermark signal. The detection watermark is specifically chosen to assist in identifying the watermark and computing its orientation in a detection operation.

[0091] **FIG. 3** is a spatial frequency plot illustrating one quadrant of a detection watermark. The points in the plot represent impulse functions indicating signal content of the detection watermark signal. The pattern of impulse functions for the illustrated quadrant is replicated in all four quadrants. There are a number of properties of the detection pattern that impact its effectiveness for a particular application. The selection of these properties is highly dependent on the application. One property is the extent to which the pattern is symmetric about one or more axes. For example, if the detection pattern is symmetrical about the horizontal and vertical axes, it is referred to as being quad symmetric. If it is further symmetrical about diagonal axes at an angle of 45 degrees, it is referred to as being octally symmetric (repeated in a symmetric pattern 8 times about the origin). Such symmetry aids in identifying the watermark in an image, and aids in extracting the rotation angle. However, in the case of an octally symmetric pattern, the detector includes an additional step of testing which of the four quadrants the orientation angle falls into.

[0092] Another criterion is the position of the impulse functions and the frequency range that they reside in. Preferably, the impulse functions fall in a mid frequency range. If they are located in a low frequency range, they may be noticeable in the watermarked image. If they are located in the high frequency range, they are more difficult to

recover. Also, they should be selected so that scaling, rotation, and other manipulations of the watermarked signal do not push the impulse functions outside the range of the detector. Finally, the impulse functions should preferably not fall on the vertical or horizontal axes, and each impulse function should have a unique horizontal and vertical location. While the example depicted in **FIG. 3** shows that some of the impulse functions fall on the same horizontal axis, it is trivial to alter the position of the impulse functions such that each has a unique vertical or horizontal coordinate.

[0093] Returning to **FIG. 2**, the embedder makes a perceptual analysis **218** of the input image **220** to identify portions of the image that can withstand more watermark signal content without substantially impacting image fidelity. Generally, the perceptual analysis employs a HVS model to identify signal frequency bands and/or spatial areas to increase or decrease watermark signal intensity to make the watermark imperceptible to an ordinary observer. One type of model is to increase watermark intensity in frequency bands and spatial areas where there is more image activity. In these areas, the sample values are changing more than other areas and have more signal strength. The output of the perceptual analysis is a perceptual mask **222**. The mask may be implemented as an array of functions, which selectively increase the signal strength of the watermark signal based on a HVS model analysis of the input image. The mask may selectively increase or decrease the signal strength of the watermark signal in areas of greater signal activity.

[0094] The embedder combines (**224**) the watermark information, the detection signal and the perceptual mask to yield the watermark signal **226**. Finally, it combines (**228**) the input image **220** and the watermark signal **226** to create the watermarked image **230**. In the frequency domain watermark example above, the embedder combines the transform domain coefficients in the watermark signal to the corresponding coefficients in the input image to create a frequency domain representation of the watermarked image. It then transforms the image into the spatial domain. As an alternative, the embedder may be designed to convert the watermark into the spatial domain, and then add it to the image.

[0095] In the spatial watermark example above, the embedder combines the image samples in the watermark signal to the corresponding samples in the input image to create the watermarked image **230**.

[0096] The embedder may employ an invertible or non-invertible, and linear or non-linear function to combine the watermark signal and the input image (e.g., linear functions such as  $S^* = S + gX$ ; or  $S^* = S(1 + gX)$ , convolution, quantization index modulation). The net effect is that some image samples in the input image are adjusted upward, while others are adjusted downward. The extent of the adjustment is greater in areas or subbands of the image having greater signal activity.

[0097] 2.2. Overview of a Detector and Reader

[0098] **FIG. 4** is a flow diagram illustrating an overview of a watermark detection process. This process analyzes image data **400** to search for an orientation pattern of a watermark in an image suspected of containing the watermark (the target image). First, the detector transforms the image data to another domain **402**, namely the spatial

frequency domain, and then performs a series of correlation or other detection operations **404**. The correlation operations match the orientation pattern with the target image data to detect the presence of the watermark and its orientation parameters **406** (e.g., translation, scale, rotation, and differential scale relative to its original orientation).

[0099] Finally, it re-oriens the image data based on one or more of the orientation parameters **408**.

[0100] If the orientation of the watermark is recovered, the reader extracts the watermark information signal from the image data (optionally by first re-orienting the data based on the orientation parameters). **FIG. 5** is flow diagram illustrating a process of extracting a message from re-oriented image data **500**. The reader scans the image samples (e.g., pixels or transform domain coefficients) of the re-oriented image (**502**), and uses known attributes of the watermark signal to estimate watermark signal values **504**. Recall that in one example implementation described above, the embedder adjusted sample values (e.g., frequency coefficients, color values, etc.) up or down to embed a watermark information signal. The reader uses this attribute of the watermark information signal to estimate its value from the target image. Prior to making these estimates, the reader may filter the image to remove portions of the image signal that may interfere with the estimating process. For example, if the watermark signal is expected to reside in low or medium frequency bands, then high frequencies may be filtered out.

[0101] In addition, the reader may predict the value of the original un-watermarked image to enhance message recovery. One form of prediction uses temporal or spatial neighbors to estimate a sample value in the original image. In the frequency domain, frequency coefficients of the original signal can be predicted from neighboring frequency coefficients in the same frequency subband. In video applications for example, a frequency coefficient in a frame can be predicted from spatially neighboring coefficients within the same frame, or temporally neighboring coefficients in adjacent frames or fields. In the spatial domain, intensity values of a pixel can be estimated from intensity values of neighboring pixels. Having predicted the value of a signal in the original, un-watermarked image, the reader then estimates the watermark signal by calculating an inverse of the watermarking function used to combine the watermark signal with the original signal.

[0102] For such watermark signal estimates, the reader uses the assignment map to find the corresponding raw bit position and image sample in the carrier signal (**506**). The value of the raw bit is a function of the watermark signal estimate, and the carrier signal at the corresponding location in the carrier. To estimate the raw bit value, the reader solves for its value based on the carrier signal and the watermark signal estimate. As reflected generally in **FIG. 5** (**508**), the result of this computation represents only one estimate to be analyzed along with other estimates impacting the value of the corresponding raw bit. Some estimates may indicate that the raw bit is likely to be a one, while others may indicate that it is a zero. After the reader completes its scan, it compiles the estimates for each bit position in the raw bit string, and makes a determination of the value of each bit at that position (**510**). Finally, it performs the inverse of the error correction coding scheme to construct the message (**512**). In some implementations, probabilistic models may be

employed to determine the likelihood that a particular pattern of raw bits is just a random occurrence rather than a watermark.

#### [0103] 2.2.1 Example Illustrating Detector Process

[0104] FIG. 6 is a diagram depicting an example of a watermark detection process. The detector segments the target image into blocks (e.g., 600, 602) and then performs a 2-dimensional fast fourier transform (2D FFT) on several blocks. This process yields 2D transforms of the magnitudes of the image contents of the blocks in the spatial frequency domain as depicted in the plot 604 shown in FIG. 6.

[0105] Next, the detector process performs a log polar remapping of each transformed block. The detector may add some of the blocks together to increase the watermark signal to noise ratio. The type of remapping in this implementation is referred to as a Fourier Mellin transform. The Fourier Mellin transform is a geometric transform that warps the image data from a frequency domain to a log polar coordinate system. As depicted in the plot 606 shown in FIG. 6, this transform sweeps through the transformed image data along a line at angle  $\theta$ , mapping the data to a log polar coordinate system shown in the next plot 608. The log polar coordinate system has a rotation axis, representing the angle  $\theta$ , and a scale axis. Inspecting the transformed data at this stage, one can see the orientation pattern of the watermark begin to be distinguishable from the noise component (i.e., the image signal).

[0106] Next, the detector performs a correlation 610 between the transformed image block and the transformed orientation pattern 612. At a high level, the correlation process slides the orientation pattern over the transformed image (in a selected transform domain, such as a spatial frequency domain) and measures the correlation at an array of discrete positions. Each such position has a corresponding scale and rotation parameter associated with it. Ideally, there is a position that clearly has the highest correlation relative to all of the others. In practice, there may be several candidates with a promising measure of correlation. As explained further below, these candidates may be subjected to one or more additional correlation stages to select the one that provides the best match.

[0107] There are a variety of ways to implement the correlation process. Any number of generalized matching filters may be implemented for this purpose. One such filter performs an FFT on the target and the orientation pattern, and multiplies the resulting arrays together to yield a multiplied FFT. Finally, it performs an inverse FFT on the multiplied FFT to return the data into its original log-polar domain. The position or positions within this resulting array with the highest magnitude represent the candidates with the highest correlation.

[0108] When there are several viable candidates, the detector can select a set of the top candidates and apply an additional correlation stage. Each candidate has a corresponding rotation and scale parameter. The correlation stage rotates and scales the FFT of the orientation pattern and performs a matching operation with the rotated and scaled pattern on the FFT of the target image. The matching operation multiplies the values of the transformed pattern with sample values at corresponding positions in the target image and accumulates the result to yield a measure of the

correlation. The detector repeats this process for each of the candidates and picks the one with the highest measure of correlation. As shown in FIG. 6, the rotation and scale parameters (614) of the selected candidate are then used to find additional parameters that describe the orientation of the watermark in the target image.

[0109] The detector applies the scale and rotation to the target data block 616 and then performs another correlation process between the orientation pattern 618 and the scaled and rotated data block 616. The correlation process 620 is a generalized matching filter operation. It provides a measure of correlation for an array of positions that each has an associated translation parameter (e.g., an x, y position). Again, the detector may repeat the process of identifying promising candidates (i.e. those that reflect better correlation relative to others) and using those in an additional search for a parameter or set of orientation parameters that provide a better measure of correlation.

[0110] At this point, the detector has recovered the following orientation parameters: rotation, scale and translation. For many applications, these parameters may be sufficient to enable accurate reading of the watermark. In the read operation, the reader applies the orientation parameters to re-orient the target image and then proceeds to extract the watermark signal.

[0111] In some applications, the watermarked image may be stretched more in one spatial dimension than another. This type of distortion is sometimes referred to as differential scale or shear. Consider that the original image blocks are square. As a result of differential scale, each square may be warped into a parallelogram with unequal sides. Differential scale parameters define the nature and extent of this stretching.

[0112] There are several alternative ways to recover the differential scale parameters. One general class of techniques is to use the known parameters (e.g., the computed scale, rotation, and translation) as a starting point to find the differential scale parameters. Assuming the known parameters to be valid, this approach warps either the orientation pattern or the target image with selected amounts of differential scale and picks the differential scale parameters that yield the best correlation.

[0113] Another approach to determination of differential scale is set forth in application Ser. No. 09/452,022 (filed Nov. 30, 1999, and entitled Method and System for Determining Image Transformation, attorney docket 60057).

#### [0114] 2.2.2 Example Illustrating Reader Process

[0115] FIG. 7 is a diagram illustrating a re-oriented image 700 superimposed onto the original watermarked image 702. The difference in orientation and scale shows how the image was transformed and edited after the embedding process. The original watermarked image is sub-divided into tiles (e.g., pixel blocks 704, 706, etc.). When superimposed on the coordinate system of the original image 702 shown in FIG. 7, the target image blocks typically do not match the orientation of the original blocks.

[0116] The reader scans samples of the re-oriented image data, estimating the watermark information signal. It estimates the watermark information signal, in part, by predicting original sample values of the un-watermarked image.

The reader then uses an inverted form of the watermarking function to estimate the watermark information signal from the watermarked signal and the predicted signal. This inverted watermarking function expresses the estimate of the watermark signal as a function of the predicted signal and the watermarked signal. Having an estimate of the watermark signal, it then uses the known relationship among the carrier signal, the watermark signal, and the raw bit to compute an estimate of the raw bit. Recall that samples in the watermark information signal are a function of the carrier signal and the raw bit value. Thus, the reader may invert this function to solve for an estimate of the raw bit value.

[0117] Recall that the embedder implementation discussed in connection with FIG. 2 redundantly encodes the watermark information signal in blocks of the input signal. Each raw bit may map to several samples within a block. In addition, the embedder repeats a mapping process for each of the blocks. As such, the reader generates several estimates of the raw bit value as it scans the watermarked image.

[0118] The information encoded in the raw bit string can be used to increase the accuracy of read operations. For instance, some of the raw bits act as signature bits that perform a validity checking function. Unlike unknown message bits, the reader knows the expected values of these signature bits. The reader can assess the validity of a read operation based on the extent to which the extracted signature bit values match the expected signature bit values. The estimates for a given raw bit value can then be given a higher weight depending on whether they are derived from a tile with a greater measure of validity.

[0119] 3.0 Embedder Implementation:

[0120] The following sections describe an implementation of the digital image watermark embedder depicted in FIG. 8. The embedder inserts two watermark components into the host image: a message component and a detection component (called the orientation pattern). The message component is defined in a spatial domain or other transform domain, while the orientation pattern is defined in a frequency domain. As explained later, the message component serves a dual function of conveying a message and helping to identify the watermark location in the image.

[0121] The embedder inserts the watermark message and orientation pattern in blocks of a selected color plane or planes (e.g., luminance or chrominance plane) of the host image. The message payload varies from one application to another, and can range from a single bit to the number of image samples in the domain in which it is embedded. The blocks may be blocks of samples in a spatial domain or some other transform domain.

[0122] 3.1 Encoding the Message

[0123] The embedder converts binary message bits into a series of binary raw bits that it hides in the host image. As part of this process, a message encoder 800 appends certain known bits to the message bits 802. It performs an error detection process (e.g., parity, Cyclic Redundancy Check (CRC), etc.) to generate error detection bits and adds the error detection bits to the message. An error correction coding operation then generates raw bits from the combined known and message bit string.

[0124] For the error correction operation, the embedder may employ any of a variety of error correction codes such as Reed Solomon, BCH, convolution or turbo codes. The encoder may perform an M-ary modulation process on the message bits that maps groups of message bits to a message signal based on an M-ary symbol alphabet.

[0125] In one application of the embedder, the component of the message representing the known bits is encoded more redundantly than the other message bits. This is an example of a shorter message component having greater signal strength than a longer, weaker message component. The embedder gives priority to the known bits in this scheme because the reader uses them to verify that it has found the watermark in a potentially corrupted image, rather than a signal masquerading as the watermark.

[0126] 3.2 Spread Spectrum Modulation

[0127] The embedder uses spread spectrum modulation as part of the process of creating a watermark signal from the raw bits. A spread spectrum modulator 804 spreads each raw bit into a number of "chips." The embedder generates a pseudo random number that acts as the carrier signal of the message. To spread each raw bit, the modulator performs an exclusive OR (XOR) operation between the raw bit and each bit of a pseudo random binary number of a pre-determined length. The length of the pseudo random number depends, in part, on the size of the message and the image. Preferably, the pseudo random number should contain roughly the same number of zeros and ones, so that the net effect of the raw bit on the host image block is zero. If a bit value in the pseudo random number is a one, the value of the raw bit is inverted. Conversely, if the bit value is a zero, then the value of the raw bit remains the same.

[0128] The length of the pseudorandom number may vary from one message bit or symbol to another. By varying the length of the number, some message bits can be spread more than others.

[0129] 3.3 Scattering the Watermark Message

[0130] The embedder scatters each of the chips corresponding to a raw bit throughout an image block. An assignment map 806 assigns locations in the block to the chips of each raw bit. Each raw bit is spread over several chips. As noted above, an image block may represent a block of transform domain coefficients or samples in a spatial domain. The assignment map may be used to encode some message bits or symbols (e.g., groups of bits) more redundantly than others by mapping selected bits to more locations in the host signal than other message bits. In addition, it may be used to map different messages, or different components of the same message, to different locations in the host signal.

[0131] FIG. 9 depicts an example of the assignment map. Each of the blocks in FIG. 9 correspond to an image block and depict a pattern of chips corresponding to a single raw bit. FIG. 9 depicts a total of 32 example blocks. The pattern within a block is represented as black dots on a white background. Each of the patterns is mutually exclusive such that each raw bit maps to a pattern of unique locations relative to the patterns of every other raw bit. Though not a requirement, the combined patterns, when overlapped, cover every location within the image block.

**[0132]** 3.4 Gain Control and Perceptual Analysis

**[0133]** To insert the information carried in a chip to the host image, the embedder alters the corresponding sample value in the host image. In particular, for a chip having a value of one, it adds to the corresponding sample value, and for a chip having a value of zero, it subtracts from the corresponding sample value. A gain controller in the embedder adjusts the extent to which each chip adds or subtracts from the corresponding sample value.

**[0134]** The gain controller takes into account the orientation pattern when determining the gain. It applies a different gain to the orientation pattern than to the message component of the watermark. After applying the gain, the embedder combines the orientation pattern and message components together to form the composite watermark signal, and combines the composite watermark with the image block. One way to combine these signal components is to add them, but other linear or non-linear functions may be used as well.

**[0135]** The orientation pattern is comprised of a pattern of quad symmetric impulse functions in the spatial frequency domain. In the spatial domain, these impulse functions look like cosine waves. An example of the orientation pattern is depicted in **FIGS. 10 and 11**. **FIG. 10** shows the impulse functions as points in the spatial frequency domain, while **FIG. 11** shows the orientation pattern in the spatial domain. Before adding the orientation pattern component to the message component, the embedder may transform the watermark components to a common domain. For example, if the message component is in a spatial domain and the orientation component is in a frequency domain, the embedder transforms the orientation component to a common spatial domain before combining them together.

**[0136]** **FIG. 8** depicts the gain controller used in the embedder. Note that the gain controller operates on the blocks of image samples **808**, the message watermark signal, and a global gain input **810**, which may be specified by the user. A perceptual analyzer component **812** of the gain controller performs a perceptual analysis on the block to identify samples that can tolerate a stronger watermark signal without substantially impacting visibility. In places where the naked eye is less likely to notice the watermark, the perceptual analyzer increases the strength of the watermark. Conversely, it decreases the watermark strength where the eye is more likely to notice the watermark.

**[0137]** The perceptual analyzer shown in **FIG. 8** performs a series of filtering operations on the image block to compute an array of gain values. There are a variety of filters suitable for this task. These filters include an edge detector filter that identifies edges of objects in the image, a non-linear filter to map gain values into a desired range, and averaging or median filters to smooth the gain values. Each of these filters may be implemented as a series of one-dimensional filters (one operating on rows and the other on columns) or two-dimensional filters. The size of the filters (i.e. the number of samples processed to compute a value for a given location) may vary (e.g., 3 by 3, 5 by 5, etc.). The shape of the filters may vary as well (e.g., square, cross-shaped, etc.). The perceptual analyzer process produces a detailed gain multiplier. The multiplier is a vector with elements corresponding to samples in a block.

**[0138]** Another component **818** of the gain controller computes an asymmetric gain based on the output of the

image sample values and message watermark signal. This component analyzes the samples of the block to determine whether they are consistent with the message signal. The embedder reduces the gain for samples whose values relative to neighboring values are consistent with the message signal.

**[0139]** The embedder applies the asymmetric gain to increase the chances of an accurate read in the watermark reader. To understand the effect of the asymmetric gain, it is helpful to explain the operation of the reader. The reader extracts the watermark message signal from the watermarked signal using a predicted version of the original signal. It estimates the watermark message signal value based on values of the predicted signal and the watermarked signal at locations of the watermarked signal suspected of containing a watermark signal. There are several ways to predict the original signal. One way is to compute a local average of samples around the sample of interest. The average may be computed by taking the average of vertically adjacent samples, horizontally adjacent samples, an average of samples in a cross-shaped filter (both vertical and horizontal neighbors), an average of samples in a square-shaped filter, etc. The estimate may be computed one time based on a single predicted value from one of these averaging computations. Alternatively, several estimates may be computed based on two or more of these averaging computations (e.g., one estimate for vertically adjacent samples and another for horizontally adjacent samples). In the latter case, the reader may keep estimates if they satisfy a similarity metric. In other words, the estimates are deemed valid if they within a predetermined value or have the same polarity.

**[0140]** Knowing this behavior of the reader, the embedder computes the asymmetric gain as follows. For samples that have values relative to their neighbors that are consistent with the watermark signal, the embedder reduces the asymmetric gain. Conversely, for samples that are inconsistent with the watermark signal, the embedder increases the asymmetric gain. For example, if the chip value is a one, then the sample is consistent with the watermark signal if its value is greater than its neighbors. Alternatively, if the chip value is a zero, then the sample is consistent with the watermark signal if its value is less than its neighbors.

**[0141]** Another component **820** of the gain controller computes a differential gain, which represents an adjustment in the message vs. orientation pattern gains. As the global gain increases, the embedder emphasizes the message gain over the orientation pattern gain by adjusting the global gain by an adjustment factor. The inputs to this process **820** include the global gain **810** and a message differential gain **822**. When the global gain is below a lower threshold, the adjustment factor is one. When the global gain is above an upper threshold, the adjustment factor is set to an upper limit greater than one. For global gains falling within the two thresholds, the adjustment factor increases linearly between one and the upper limit. The message differential gain is the product of the adjustment factor and the global gain.

**[0142]** At this point, there are four sources of gain: the detailed gain, the global gain, the asymmetric gain, and the message dependent gain. The embedder applies the first two gain quantities to both the message and orientation watermark signals. It only applies the latter two to the message watermark signal. **FIG. 8** depicts how the embedder applies

the gain to the two watermark components. First, it multiplies the detailed gain with the global gain to compute the orientation pattern gain. It then multiplies the orientation pattern gain with the adjusted message differential gain and asymmetric gain to form the composite message gain.

[0143] Finally, the embedder forms the composite watermark signal. It multiplies the composite message gain with the message signal, and multiplies the orientation pattern gain with the orientation pattern signal. It then combines the result in a common transform domain to get the composite watermark. The embedder applies a watermarking function to combine the composite watermark to the block to create a watermarked image block. The message and orientation components of the watermark may be combined by mapping the message bits to samples of the orientation signal, and modulating the samples of the orientation signal to encode the message.

[0144] The embedder computes the watermark message signal by converting the output of the assignment map **806** to delta values, indicating the extent to which the watermark signal changes the host signal. As noted above, a chip value of one corresponds to an upward adjustment of the corresponding sample, while a chip value of zero corresponds to a downward adjustment. The embedder specifies the specific amount of adjustment by assigning a delta value to each of the watermark message samples (**830**).

[0145] 4.0 Detector Implementation

[0146] **FIG. 12** illustrates an overview of a watermark detector that detects the presence of a detection watermark in a host image and its orientation. Using the orientation pattern and the known bits inserted in the watermark message, the detector determines whether a potentially corrupted image contains a watermark, and if so, its orientation in the image.

[0147] Recall that the composite watermark is replicated in blocks of the original image. After an embedder places the watermark in the original digital image, the watermarked image is likely to undergo several transformations, either from routine processing or from intentional tampering. Some of these transformations include: compression, decompression, color space conversion, digital to analog conversion, printing, scanning, analog to digital conversion, scaling, rotation, inversion, flipping differential scale, and lens distortion. In addition to these transformations, various noise sources can corrupt the watermark signal, such as fixed pattern noise, thermal noise, etc.

[0148] When building a detector implementation for a particular application, the developer may implement counter-measures to mitigate the impact of the types of transformations, distortions and noise expected for that application. Some applications may require more counter-measures than others. The detector described below is designed to recover a watermark from a watermarked image after the image has been printed, and scanned. The following sections describe the counter-measures to mitigate the impact of various forms of corruption. The developer can select from among these counter-measures when implementing a detector for a particular application.

[0149] For some applications, the detector will operate in a system that provides multiple image frames of a watermarked object. One typical example of such a system is a

computer equipped with a digital camera. In such a configuration, the digital camera can capture a temporal sequence of images as the user or some device presents the watermarked image to the camera.

[0150] As shown in **FIG. 12**, the principal components of the detector are: 1) preprocessor **900**; 2) rotation and scale estimator **902**; 3) orientation parameter refiner **904**; 4) translation estimator **906**; 5) translation refiner **908**; and reader **910**.

[0151] The preprocessor **900** takes one or more frames of image data **912** and produces a set of image blocks **914** prepared for further analysis. The rotation-scale estimator **902** computes rotation-scale vectors **916** that estimate the orientation of the orientation signal in the image blocks. The parameter refiner **904** collects additional evidence of the orientation signal and further refines the rotation scale vector candidates by estimating differential scale parameters. The result of this refining stage is a set of 4D vectors candidates **918** (rotation, scale, and two differential scale parameters). The translation estimator **906** uses the 4D vector candidates to re-orient image blocks with promising evidence of the orientation signal. It then finds estimates of translation parameters **920**. The translation refiner **908** invokes the reader **910** to assess the merits of an orientation vector. When invoked by the detector, the reader uses the orientation vector to approximate the original orientation of the host image and then extracts values for the known bits in the watermark message. The detector uses this information to assess the merits of and refine orientation vector candidates.

[0152] By comparing the extracted values of the known bits with the expected values, the reader provides a figure of merit for an orientation vector candidate. The translation refiner then picks a 6D vector, including rotation, scale, differential scale and translation, that appears likely produce a valid read of the watermark message **922**. The following sections describe implementations of these components in more detail.

[0153] 4.1 Detector Pre-Processing

[0154] **FIG. 13** is a flow diagram illustrating preprocessing operations in the detector shown in **FIG. 12**. The detector performs a series of pre-processing operations on the native image **930** to prepare the image data for further analysis. It begins by filling memory with one or more frames of native image data (**932**), and selecting sets of pixel blocks **934** from the native image data for further analysis (**936**). While the detector can detect a watermark using a single image frame, it also has support for detecting the watermark using additional image frames. As explained below, the use of multiple frames has the potential for increasing the chances of an accurate detection and read.

[0155] In applications where a camera captures an input image of a watermarked object, the detector may be optimized to address problems resulting from movement of the object. Typical PC cameras, for example, are capable of capturing images at a rate of at least 10 frames a second. A frustrated user might attempt to move the object in an attempt to improve detection. Rather than improving the chances of detection, the movement of the object changes the orientation of the watermark from one frame to the next, potentially making the watermark more difficult to detect.



One way to address this problem is to buffer one or more frames, and then screen the frame or frames to determine if they are likely to contain a valid watermark signal. If such screening indicates that a frame is not likely to contain a valid signal, the detector can discard it and proceed to the next frame in the buffer, or buffer a new frame. Another enhancement is to isolate portions of a frame that are most likely to have a valid watermark signal, and then perform more detailed detection of the isolated portions.

[0156] After loading the image into the memory, the detector selects image blocks 934 for further analysis. It is not necessary to load or examine each block in a frame because it is possible to extract the watermark using only a portion of an image. The detector looks at only a subset of the samples in an image, and preferably analyzes samples that are more likely to have a recoverable watermark signal.

[0157] The detector identifies portions of the image that are likely to have the highest watermark signal to noise ratio. It then attempts to detect the watermark signal in the identified portions. In the context of watermark detection, the host image is considered to be a source of noise along with conventional noise sources. While it is typically not practical to compute the signal to noise ratio, the detector can evaluate attributes of the signal that are likely to evince a promising watermark signal to noise ratio. These properties include the signal activity (as measured by sample variance, for example), and a measure of the edges (abrupt changes in image sample values) in an image block. Preferably, the signal activity of a candidate block should fall within an acceptable range, and the block should not have a high concentration of strong edges. One way to quantify the edges in the block is to use an edge detection filter (e.g., a LaPlacian, Sobel, etc.).

[0158] In one implementation, the detector divides the input image into blocks, and analyzes each block based on pre-determined metrics. It then ranks the blocks according to these metrics. The detector then operates on the blocks in the order of the ranking. The metrics include sample variance in a candidate block and a measure of the edges in the block. The detector combines these metrics for each candidate block to compute a rank representing the probability that it contains a recoverable watermark signal.

[0159] In another implementation, the detector selects a pattern of blocks and evaluates each one to try to make the most accurate read from the available data. In either implementation, the block pattern and size may vary. This particular implementation selects a pattern of overlapping blocks (e.g., a row of horizontally aligned, overlapping blocks). One optimization of this approach is to adaptively select a block pattern that increases the signal to noise ratio of the watermark signal. While shown as one of the initial operations in the preparation, the selection of blocks can be postponed until later in the pre-processing stage.

[0160] Next, the detector performs a color space conversion on native image data to compute an array of image samples in a selected color space for each block (936). In the following description, the color space is luminance, but the watermark may be encoded in one or more different color spaces. The objective is to get a block of image samples with lowest noise practical for the application. While the implementation currently performs a row by row conversion of the native image data into 8 bit integer luminance values, it

may be preferable to convert to floating-point values for some applications. One optimization is to select a luminance converter that is adapted for the sensor used to capture the digital input image. For example, one might experimentally derive the lowest noise luminance conversion for commercially available sensors, e.g., CCD cameras or scanners, CMOS cameras, etc. Then, the detector could be programmed to select either a default luminance converter, or one tuned to a specific type of sensor.

[0161] At one or more stages of the detector, it may be useful to perform operations to mitigate the impact of noise and distortion. In the pre-processing phase, for example, it may be useful to evaluate fixed pattern noise and mitigate its effect (938). The detector may look for fixed pattern noise in the native input data or the luminance data, and then mitigate it.

[0162] One way to mitigate certain types of noise is to combine data from different blocks in the same frame, or corresponding blocks in different frames 940. This process helps augment the watermark signal present in the blocks, while reducing the noise common to the blocks. For example, merely adding blocks together may mitigate the effects of common noise.

[0163] In addition to common noise, other forms of noise may appear in each of the blocks such as noise introduced in the printing or scanning processes. Depending on the nature of the application, it may be advantageous to perform common noise recognition and removal at this stage 942. The developer may select a filter or series of filters to target certain types of noise that appear during experimentation with images. Certain types of median filters may be effective in mitigating the impact of spectral peaks (e.g., speckles) introduced in printing or scanning operations.

[0164] In addition to introducing noise, the printing and image capture processes may transform the color or orientation of the original, watermarked image. As described above, the embedder typically operates on a digital image in a particular color space and at a desired resolution. The watermark embedders normally operate on digital images represented in an RGB or CYMK color space at a desired resolution (e.g., 100 dpi or 300 dpi, the resolution at which the image is printed). The images are then printed on paper with a screen printing process that uses the CYMK subtractive color space at a line per inch (LPI) ranging from 65-200. 133 lines/in is typical for quality magazines and 73 lines/in is typical for newspapers. In order to produce a quality image and avoid pixelization, the rule of thumb is to use digital images with a resolution that is at least twice the press resolution. This is due to the half tone printing for color production. Also, different presses use screens with different patterns and line orientations and have different precision for color registration.

[0165] One way to counteract the transforms introduced through the printing process is to develop a model that characterizes these transforms and optimize watermark embedding and detecting based on this characterization. Such a model may be developed by passing watermarked and unwatermarked images through the printing process and observing the changes that occur to these images. The resulting model characterizes the changes introduced due to the printing process. The model may represent a transfer function that approximates the transforms due to the printing

process. The detector then implements a pre-processing stage that reverses or at least mitigates the effect of the printing process on watermarked images. The detector may implement a pre-processing stage that performs the inverse of the transfer function for the printing process.

[0166] A related challenge is the variety in paper attributes used in different printing processes. Papers of various qualities, thickness and stiffness, absorb ink in various ways. Some papers absorb ink evenly, while others absorb ink at rates that vary with the changes in the paper's texture and thickness. These variations may degrade the embedded watermark signal when a digitally watermarked image is printed. The watermark process can counteract these effects by classifying and characterizing paper so that the embedder and reader can compensate for this printing-related degradation.

[0167] Variations in image capture processes also pose a challenge. In some applications, it is necessary to address problems introduced due to interlaced image data. Some video camera produce interlaced fields representing the odd or even scan lines of a frame. Problems arise when the interlaced image data consists of fields from two consecutive frames. To construct an entire frame, the preprocessor may combine the fields from consecutive frames while dealing with the distortion due to motion that occurs from one frame to the next. For example, it may be necessary to shift one field before interleaving it with another field to counteract inter-frame motion. A de-blurring function may be used to mitigate the blurring effect due to the motion between frames.

[0168] Another problem associated with cameras in some applications is blurring due to the lack of focus. The preprocessor can mitigate this effect by estimating parameters of a blurring function and applying a de-blurring function to the input image.

[0169] Yet another problem associated with cameras is that they tend to have color sensors that utilize different color pattern implementations. As such, a sensor may produce colors slightly different than those represented in the object being captured. Most CCD and CMOS cameras use an array of sensors to produce colored images. The sensors in the array are arranged in clusters of sensitive to three primary colors red, green, and blue according to a specific pattern. Sensors designated for a particular color are dyed with that color to increase their sensitivity to the designated color. Many camera manufacturers use a Bayer color pattern GR/BG. While this pattern produces good image quality, it causes color mis-registration that degrades the watermark signal. Moreover, the color space converter, which maps the signal from the sensors to another color space such as YUV or RGB, may vary from one manufacturer to another. One way to counteract the mis-registration of the camera's color pattern is to account for the distortion due to the pattern in a color transformation process, implemented either within the camera itself, or as a pre-processing function in the detector.

[0170] Another challenge in counteracting the effects of the image capture process is dealing with the different types of distortion introduced from various image capture devices. For example, cameras have different sensitivities to light. In addition, their lenses have different spherical distortion, and noise characteristics. Some scanners have poor color repro-

duction or introduce distortion in the image aspect ratio. Some scanners introduce aliasing and employ interpolation to increase resolution. The detector can counteract these effects in the pre-processor by using an appropriate inverse transfer function. An off-line process first characterizes the distortion of several different image capture devices (e.g., by passing test images through the scanner and deriving a transfer function modeling the scanner distortion). Some detectors may be equipped with a library of such inverse transfer functions from which they select one that corresponds to the particular image capture device

[0171] Yet another challenge in applications where the image is printed on paper and later scanned is that the paper deteriorates over time and degrades the watermark. Also, varying lighting conditions make the watermark difficult to detect. Thus, the watermark may be selected so as to be more impervious to expected deterioration, and recoverable over a wider range of lighting conditions.

[0172] At the close of the pre-processing stage, the detector has selected a set of blocks for further processing. It then proceeds to gather evidence of the orientation signal in these blocks, and estimate the orientation parameters of promising orientation signal candidates. Since the image may have suffered various forms of corruption, the detector may identify several parts of the image that appear to have attributes similar to the orientation signal. As such, the detector may have to resolve potentially conflicting and ambiguous evidence of the orientation signal. To address this challenge, the detector estimates orientation parameters, and then refines these estimates to extract the orientation parameters that are more likely to evince a valid signal than other parameter candidates.

[0173] 4.2 Estimating Initial Orientation Parameters

[0174] FIG. 14 is a flow diagram illustrating a process for estimating rotation-scale vectors. The detector loops over each image block (950), calculating rotation-scale vectors with the best detection values in each block. First, the detector filters the block in a manner that tends to amplify the orientation signal while suppressing noise, including noise from the host image itself (952). Implemented as a multi-axis Laplacian filter, the filter highlights edges (e.g., high frequency components of the image) and then suppresses them. The term, "multi-axis," means that the filter includes a series of stages that each operates on particular axis. First, the filter operates on the rows of luminance samples, then operates on the columns, and adds the results. The filter may be applied along other axes as well. Each pass of the filter produces values at discrete levels. The final result is an array of samples, each having one of five values:  $\{-2, -1, 0, 1, 2\}$ .

[0175] Next, the detector performs a windowing operation on the block data to prepare it for an FFT transform (954). This windowing operation provides signal continuity at the block edges. The detector then performs an FFT (956) on the block, and retains only the magnitude component (958).

[0176] In an alternative implementation, the detector may use the phase signal produced by the FFT to estimate the translation parameter of the orientation signal. For example, the detector could use the rotation and scale parameters extracted in the process described below, and then compute the phase that provided the highest measure of correlation with the orientation signal using the phase component of the FFT process.

[0177] After computing the FFT, the detector applies a Fourier magnitude filter (960) on the magnitude components. The filter in the implementation slides over each sample in the Fourier magnitude array and filters the sample's eight neighbors in a square neighborhood centered at the sample. The filter boosts values representing a sharp peak with a rapid fall-off, and suppresses the fall-off portion. It also performs a threshold operation to clip peaks to an upper threshold.

[0178] Next, the detector performs a log-polar re-sample (962) of the filtered Fourier magnitude array to produce a log-polar array 964. This type of operation is sometimes referred to as a Fourier Mellin transform. The detector, or some off-line pre-processor, performs a similar operation on the orientation signal to map it to the log-polar coordinate system. Using matching filters, the detector implementation searches for an orientation signal in a specified window of the log-polar coordinate system. For example, consider that the log-polar coordinate system is a two dimensional space with the scale being the vertical axis and the angle being the horizontal axis. The window ranges from 0 to 90 degrees on the horizontal axis and from approximately 50 to 2400 dpi on the vertical axis. Note that the orientation pattern should be selected so that routine scaling does not push the orientation pattern out of this window. The orientation pattern can be designed to mitigate this problem, as noted above, and as explained in co-pending patent application Ser. No. 60/136, 572, filed May 28, 1999, by Ammon Gustafson, entitled Watermarking System With Improved Technique for Detecting Scaling and Rotation, filed May 28, 1999.

[0179] The detector proceeds to correlate the orientation and the target signal in the log polar coordinate system. As shown in FIG. 14, the detector uses a generalized matched filter GMF (966). The GMF performs an FFT on the orientation and target signal, multiplies the resulting Fourier domain entities, and performs an inverse FFT. This process yields a rectangular array of values in log-polar coordinates, each representing a measure of correlation and having a corresponding rotation angle and scale vector. As an optimization, the detector may also perform the same correlation operations for distorted versions (968, 970, 972) of the orientation signal to see if any of the distorted orientation patterns results in a higher measure of correlation. For example, the detector may repeat the correlation operation with some pre-determined amount of horizontal and vertical differential distortion (970, 972). The result of this correlation process is an array of correlation values 974 specifying the amount of correlation that each corresponding rotation-scale vector provides.

[0180] The detector processes this array to find the top M peaks and their location in the log-polar space 976. To extract the location more accurately, the detector uses interpolation to provide the inter-sample location of each of the top peaks 978. The interpolator computes the 2D median of the samples around a peak and provides the location of the peak center to an accuracy of 0.1 sample.

[0181] The detector proceeds to rank the top rotation-scale vectors based on yet another correlation process 980. In particular, the detector performs a correlation between a Fourier magnitude representation for each rotation-scale vector candidate and a Fourier magnitude specification of the orientation signal 982. Each Fourier magnitude repre-

sentation is scaled and rotated by an amount reflected by the corresponding rotation-scale vector. This correlation operation sums a point-wise multiplication of the orientation pattern impulse functions in the frequency domain with the Fourier magnitude values of the image at corresponding frequencies to compute a measure of correlation for each peak 984. The detector then sorts correlation values for the peaks (986).

[0182] Finally, the detector computes a detection value for each peak (988). It computes the detection value by quantizing the correlation values. Specifically, it computes a ratio of the peak's correlation value and the correlation value of the next largest peak. Alternatively, the detector may compute the ratio of the peak's correlation value and a sum or average of the correlation values of the next n highest peaks, where n is some predetermined number. Then, the detector maps this ratio to a detection value based on a statistical analysis of unmarked images.

[0183] The statistical analysis plots a distribution of peak ratio values found in unmarked images. The ratio values are mapped to a detection value based on the probability that the value came from an unmarked image. For example, 90% of the ratio values in unmarked images fall below a first threshold T1, and thus, the detection value mapping for a ratio of T1 is set to 1. Similarly, 99% of the ratio values in unmarked images fall below T2, and therefore, the detection value is set to 2. 99.9% of the ratio values in unmarked images fall below T3, and the corresponding detection value is set to 3. The threshold values, T1, T2 and T3, may be determined by performing a statistical analysis of several images. The mapping of ratios to detection values based on the statistical distribution may be implemented in a look up table.

[0184] The statistical analysis may also include a maximum likelihood analysis. In such an analysis, an off-line detector generates detection value statistics for both marked and unmarked images. Based on the probability distributions of marked and unmarked images, it determines the likelihood that a given detection value for an input image originates from a marked and unmarked image.

[0185] At the end of these correlation stages, the detector has computed a ranked set of rotation-scale vectors 990, each with a quantized measure of correlation associated with it. At this point, the detector could simply choose the rotation and scale vectors with the highest rank and proceed to compute other orientation parameters, such as differential scale and translation. Instead, the detector gathers more evidence to refine the rotation-scale vector estimates. FIG. 15 is a flow diagram illustrating a process for refining the orientation parameters using evidence of the orientation signal collected from blocks in the current frame.

[0186] Continuing in the current frame, the detector proceeds to compare the rotation and scale parameters from different blocks (e.g., block 0, block 1, block 2; 1000, 1002, and 1004 in FIG. 15). In a process referred to as interblock coincidence matching 1006, it looks for similarities between rotation-scale parameters that yielded the highest correlation in different blocks. To quantify this similarity, it computes the geometric distance between each peak in one block with every other peak in the other blocks. It then computes the probability that peaks will fall within this calculated distance. There are a variety of ways to calculate the probabil-

ity. In one implementation, the detector computes the geometric distance between two peaks, computes the circular area encompassing the two peaks ( $\pi(\text{geometric distance})^2$ ), and computes the ratio of this area to the total area of the block. Finally, it quantizes this probability measure for each pair of peaks (1008) by computing the log (base 10) of the ratio of the total area over the area encompassing the two peaks. At this point, the detector has calculated two detection values: quantized peak value, and the quantized distance metric.

[0187] The detector now forms multi-block grouping of rotation-scale vectors and computes a combined detection value for each grouping (1010). The detector groups vectors based on their relative geometric proximity within their respective blocks. It then computes the combined detection value by combining the detection values of the vectors in the group (1012). One way to compute a combined detection value is to add the detection values or add a weighted combination of them.

[0188] Having calculated the combined detection values, the detector sorts each grouping by its combined detection value (1014). This process produces a set of the top groupings of unrefined rotation-scale candidates, ranked by detection value 1016. Next, the detector weeds out rotation-scale vectors that are not promising by excluding those groupings whose combined detection values are below a threshold (the "refine threshold" 1018). The detector then refines each individual rotation-scale vector candidate within the remaining groupings.

[0189] The detector refines a rotation-scale vector by adjusting the vector and checking to see whether the adjustment results in a better correlation. As noted above, the detector may simply pick the best rotation-scale vector based on the evidence collected thus far, and refine only that vector. An alternative approach is to refine each of the top rotation-scale vector candidates, and continue to gather evidence for each candidate. In this approach, the detector loops over each vector candidate (1020), refining each one.

[0190] One approach of refining the orientation vector is as follows:

[0191] fix the orientation signal impulse functions ("points") within a valid boundary (1022);

[0192] pre-refine the rotation-scale vector (1024);

[0193] find the major axis and re-fix the orientation points (1026); and

[0194] refine each vector with the addition of a differential scale component (1028).

[0195] In this approach, the detector pre-refines a rotation-scale vector by incrementally adjusting one of the parameters (scale, rotation angle), adjusting the orientation points, and then summing a point-wise multiplication of the orientation pattern and the image block in the Fourier magnitude domain. The refiner compares the resulting measure of correlation with previous measures and continues to adjust one of the parameters so long as the correlation increases. After refining the scale and rotation angle parameters, the refiner finds the major axis, and re-fixes the orientation points. It then repeats the refining process with the introduction of differential scale parameters. At the end of this process, the refiner has converted each scale-rotation can-

didate to a refined 4D vector, including rotation, scale, and two differential scale parameters.

[0196] At this stage, the detector can pick a 4D vector or set of 4D vector and proceed to calculate the final remaining parameter, translation. Alternatively, the detector can collect additional evidence about the merits of each 4D vector.

[0197] One way to collect additional evidence about each 4D vector is to re-compute the detection value of each orientation vector candidate (1030). For example, the detector may quantize the correlation value associated with each 4D vector as described above for the rotation-scale vector peaks (see item 988, FIG. 14 and accompanying text). Another way to collect additional evidence is to repeat the coincidence matching process for the 4D vectors. For this coincidence matching process, the detector computes spatial domain vectors for each candidate (1032), determines the distance metric between candidates from different blocks, and then groups candidates from different blocks based on the distance metrics (1034). The detector then re-sorts the groups according to their combined detection values (1036) to produce a set of the top P groupings 1038 for the frame.

[0198] FIG. 16 is a flow diagram illustrating a method for aggregating evidence of the orientation signal from multiple frames. In applications with multiple frames, the detector collects the same information for orientation vectors of the selected blocks in each frame (namely, the top P groupings of orientation vector candidates, e.g., 1050, 1052 and 1054). The detector then repeats coincidence matching between orientation vectors of different frames (1056). In particular, in this inter-frame mode, the detector quantizes the distance metrics computed between orientation vectors from blocks in different frames (1058). It then finds inter-frame groupings of orientation vectors (super-groups) using the same approach described above (1060), except that the orientation vectors are derived from blocks in different frames. After organizing orientation vectors into super-groups, the detector computes a combined detection value for each super-group (1062) and sorts the super-groups by this detection value (1064). The detector then evaluates whether to proceed to the next stage (1066), or repeat the above process of computing orientation vector candidates from another frame (1068).

[0199] If the detection values of one or more super-groups exceed a threshold, then the detector proceeds to the next stage. If not, the detector gathers evidence of the orientation signal from another frame and returns to the inter-frame coincidence matching process. Ultimately, when the detector finds sufficient evidence to proceed to the next stage, it selects the super-group with the highest combined detection value (1070), and sorts the blocks based on their corresponding detection values (1072) to produce a ranked set of blocks for the next stage (1074).

[0200] 4.3 Estimating Translation Parameters

[0201] FIG. 17 is a flow diagram illustrating a method for estimating translation parameters of the orientation signal, using information gathered from the previous stages.

[0202] In this stage, the detector estimates translation parameters. These parameters indicate the starting point of a watermarked block in the spatial domain. The translation parameters, along with rotation, scale and differential scale, form a complete 6D orientation vector. The 6D vector

enables the reader to extract luminance sample data in approximately the same orientation as in the original watermarked image.

[0203] One approach is to use generalized match filtering to find the translation parameters that provide the best correlation. Another approach is to continue to collect evidence about the orientation vector candidates, and provide a more comprehensive ranking of the orientation vectors based on all of the evidence gathered thus far. The following paragraphs describe an example of this type of an approach.

[0204] To extract translation parameters, the detector proceeds as follows. In the multi-frame case, the detector selects the frame that produced 4D orientation vectors with the highest detection values (1080). It then processes the blocks 1082 in that frame in the order of their detection value. For each block (1084), it applies the 4D vector to the luminance data to generate rectified block data (1086). The detector then performs dual axis filtering (1088) and the window function (1090) on the data. Next, it performs an FFT (1092) on the image data to generate an array of Fourier data. To make correlation operations more efficient, the detector buffers the fourier values at the orientation points (1094).

[0205] The detector applies a generalized match filter 1096 to correlate a phase specification of the orientation signal (1098) with the transformed block data. The result of this process is a 2D array of correlation values. The peaks in this array represent the translation parameters with the highest correlation. The detector selects the top peaks and then applies a median filter to determine the center of each of these peaks. The center of the peak has a corresponding correlation value and sub-pixel translation value. This process is one example of getting translation parameters by correlating the Fourier phase specification of the orientation signal and the image data. Other methods of phase locking the image data with a synchronization signal like the orientation signal may also be employed.

[0206] Depending on the implementation, the detector may have to resolve additional ambiguities, such as rotation angle and flip ambiguity. The degree of ambiguity in the rotation angle depends on the nature of the orientation signal. If the orientation signal is octally symmetric (symmetric about horizontal, vertical and diagonal axes in the spatial frequency domain), then the detector has to check each quadrant (0-90, 90-180, 180-270, and 270-360 degrees) to find out which one the rotation angle resides in. Similarly, if the orientation signal is quad symmetric, then the detector has to check two cases, 0-180 and 180-270.

[0207] The flip ambiguity may exist in some applications where the watermarked image can be flipped. To check for rotation and flip ambiguities, the detector loops through each possible case, and performs the correlation operation for each one (1100).

[0208] At the conclusion of the correlation process, the detector has produced a set of the top translation parameters with associated correlation values for each block. To gather additional evidence, the detector groups similar translation parameters from different blocks (1102), calculates a group detection value for each set of translation parameters 1104, and then ranks the top translation groups based on their corresponding group detection values 1106.

#### [0209] 4.4 Refining Translation Parameters

[0210] Having gathered translation parameter estimates, the detector proceeds to refine these estimates. FIG. 18 is a flow diagram illustrating a process for refining orientation parameters. At this stage, the detector process has gathered a set of the top translation parameter candidates 1120 for a given frame 1122. The translation parameters provide an estimate of a reference point that locates the watermark, including both the orientation and message components, in the image frame. In the implementation depicted here, the translation parameters are represented as horizontal and vertical offsets from a reference point in the image block from which they were computed.

[0211] Recall that the detector has grouped translation parameters from different blocks based on their geometric proximity to each other. Each pair of translation parameters in a group is associated with a block and a 4D vector (rotation, scale, and 2 differential scale parameters). As shown in FIG. 18, the detector can now proceed to loop through each group (1124), and through the blocks within each group (1126), to refine the orientation parameters associated with each member of the groups. Alternatively, a simpler version of the detector may evaluate only the group with the highest detection value, or only selected blocks within that group.

[0212] Regardless of the number of candidates to be evaluated, the process of refining a given orientation vector candidate may be implemented in a similar fashion. In the refining process, the detector uses a candidate orientation vector to define a mesh of sample blocks for further analysis (1128). In one implementation, for example, the detector forms a mesh of 32 by 32 sample blocks centered around a seed block whose upper right corner is located at the vertical and horizontal offset specified by the candidate translation parameters. The detector reads samples from each block using the orientation vector to extract luminance samples that approximate the original orientation of the host image at encoding time.

[0213] The detector steps through each block of samples (1130). For each block, it sets the orientation vector (1132), and then uses the orientation vector to check the validity of the watermark signal in the sample block. It assesses the validity of the watermark signal by calculating a figure of merit for the block (1134). To further refine the orientation parameters associated with each sample block, the detector adjusts selected parameters (e.g., vertical and horizontal translation) and re-calculates the figure of merit. As depicted in the inner loop in FIG. 18 (block 1136 to 1132), the detector repeatedly adjusts the orientation vector and calculates the figure of merit in an attempt to find a refined orientation that yields a higher figure of merit.

[0214] The loop (1136) may be implemented by stepping through a predetermined sequence of adjustments to parameters of the orientation vectors (e.g., adding or subtracting small increments from the horizontal and vertical translation parameters). In this approach, the detector exits the loop after stepping through the sequence of adjustments. Upon exiting, the detector retains the orientation vector with the highest figure of merit.

[0215] There are a number of ways to calculate this figure of merit. One figure of merit is the degree of correlation

between a known watermark signal attribute and a corresponding attribute in the signal suspected of having a watermark. Another figure of merit is the strength of the watermark signal (or one of its components) in the suspect signal. For example, a figure of merit may be based on a measure of the watermark message signal strength and/or orientation pattern signal strength in the signal, or in a part of the signal from which the detector extracts the orientation parameters. The detector may compute a figure of merit based the strength of the watermark signal in a sample block. It may also compute a figure of merit based on the percentage agreement between the known bits of the message and the message bits extracted from the sample block.

[0216] When the figure of merit is computed based on a portion of the suspect signal, the detector and reader can use the figure of merit to assess the accuracy of the watermark signal detected and read from that portion of the signal. This approach enables the detector to assess the merits of orientation parameters and to rank them based on their figure of merit. In addition, the reader can weight estimates of watermark message values based on the figure of merit to recover a message more reliably.

[0217] The process of calculating a figure of merit depends on attributes the watermark signal and how the embedder inserted it into the host signal. Consider an example where the watermark signal is added to the host signal. To calculate a figure of merit based on the strength of the orientation signal, the detector checks the value of each sample relative to its neighbors, and compares the result with the corresponding sample in a spatial domain version of the orientation signal. When a sample's value is greater than its neighbors, then one would expect that the corresponding orientation signal sample to be positive. Conversely, when the sample's value is less than its neighbors, then one would expect that the corresponding orientation sample to be negative. By comparing a sample's polarity relative to its neighbors with the corresponding orientation sample's polarity, the detector can assess the strength of the orientation signal in the sample block. In one implementation, the detector makes this polarity comparison twice for each sample in an N by N block (e.g., N=32, 64, etc): once comparing each sample with its horizontally adjacent neighbors and then again comparing each sample with its vertically adjacent neighbors. The detector performs this analysis on samples in the mesh block after re-orienting the data to approximate the original orientation of the host image at encoding time. The result of this process is a number reflecting the portion of the total polarity comparisons that yield a match.

[0218] To calculate a figure of merit based on known signature bits in a message, the detector invokes the reader on the sample block, and provides the orientation vector to enable the reader to extract coded message bits from the sample block. The detector compares the extracted message bits with the known bits to determine the extent to which they match. The result of this process is a percentage agreement number reflecting the portion of the extracted message bits that match the known bits. Together the test for the orientation signal and the message signal provide a figure of merit for the block.

[0219] As depicted in the loop from blocks 1138 to 1130, the detector may repeat the process of refining the orienta-

tion vector for each sample block around the seed block. In this case, the detector exits the loop (1138) after analyzing each of the sample blocks in the mesh defined previously (1128). In addition, the detector may repeat the analysis in the loop through all blocks in a given group (1140), and in the loop through each group (1142).

[0220] After completing the analysis of the orientation vector candidates, the detector proceeds to compute a combined detection value for the various candidates by compiling the results of the figure of merit calculations. It then proceeds to invoke the reader on the orientation vector candidates in the order of their detection values.

[0221] 4.5 Reading the Watermark

[0222] FIG. 19 is a flow diagram illustrating a process for reading the watermark message. Given an orientation vector and the corresponding image data, the reader extracts the raw bits of a message from the image. The reader may accumulate evidence of the raw bit values from several different blocks. For example, in the process depicted in FIG. 19, the reader uses refined orientation vectors for each block, and accumulates evidence of the raw bit values extracted from the blocks associated with the refined orientation vectors.

[0223] The reading process begins with a set of promising orientation vector candidates 1150 gathered from the detector. In each group of orientation vector candidates, there is a set of orientation vectors, each corresponding to a block in a given frame. The detector invokes the reader for one or more orientation vector groups whose detection values exceed a predetermined threshold. For each such group, the detector loops over the blocks in the group (1152), and invokes the reader to extract evidence of the raw message bit values.

[0224] Recall that previous stages in the detector have refined orientation vectors to be used for the blocks of a group. When it invokes the reader, the detector provides the orientation vector as well as the image block data (1154). The reader scans samples starting from a location in a block specified by the translation parameters and using the other orientation parameters to approximate the original orientation of the image data (1156).

[0225] As described above, the embedder maps chips of the raw message bits to each of the luminance samples in the original host image. Each sample, therefore, may provide an estimate of a chip's value. The reader reconstructs the value of the chip by first predicting the watermark signal in the sample from the value of the sample relative to its neighbors as described above (1158). If the deduced value appears valid, then the reader extracts the chip's value using the known value of the pseudo-random carrier signal for that sample and performing the inverse of the modulation function originally used to compute the watermark information signal (1160). In particular, the reader performs an exclusive OR operation on the deduced value and the known carrier signal bit to get an estimate of the raw bit value. This estimate serves as an estimate for the raw bit value. The reader accumulates these estimates for each raw bit value (1162).

[0226] As noted above, the reader computes an estimate of the watermark signal by predicting the original, un-watermarked signal and deriving an estimate of the watermark

signal based on the predicted signal and the watermarked signal. It then computes an estimate of a raw bit value based on the value of the carrier signal, the assignment map that maps a raw bit to the host image, and the relationship among the carrier signal value, the raw bit value, and the watermark signal value. In short, the reader reverses the embedding functions that modulate the message with the carrier and apply the modulated carrier to the host signal. Using the predicted value of the original signal and an estimate of the watermark signal, the reader reverses the embedding functions to estimate a value of the raw bit.

[0227] The reader loops over the candidate orientation vectors and associated blocks, accumulating estimates for each raw bit value (1164). When the loop is complete, the reader calculates a final estimate value for each raw bit from the estimates compiled for it. It then performs the inverse of the error correction coding operation on the final raw bit values (1166). Next, it performs a CRC to determine whether the read is valid. If no errors are detected, the read operation is complete and the reader returns the message (1168).

[0228] However, if the read is invalid, then the detector may either attempt to refine the orientation vector data further, or start the detection process with a new frame. Preferably, the detector should proceed to refine the orientation vector data when the combined detection value of the top candidates indicates that the current data is likely to contain a strong watermark signal. In the process depicted in FIG. 19, for example, the detector selects a processing path based on the combined detection value (1170). The combined detection value may be calculated in a variety of ways. One approach is to compute a combined detection value based on the geometric coincidence of the top orientation vector candidates and a compilation of their figures of merit. The figure of merit may be computed as detailed earlier.

[0229] For cases where the read is invalid, the processing paths for the process depicted in FIG. 19 include: 1) refine the top orientation vectors in the spatial domain (1172); 2) invoke the translation estimator on the frame with the next best orientation vector candidates (1174); and 3) re-start the detection process on a new frame (assuming an implementation where more than one frame is available)(1176). These paths are ranked in order from the highest detection value to the lowest. In the first case, the orientation vectors are the most promising. Thus, the detector re-invokes the reader on the same candidates after refining them in the spatial domain (1178). In the second case, the orientation vectors are less promising, yet the detection value indicates that it is still worthwhile to return to the translation estimation stage and continue from that point. Finally, in the final case, the detection value indicates that the watermark signal is not strong enough to warrant further refinement. In this case, the detector starts over with the next new frame of image data.

[0230] In each of the above cases, the detector continues to process the image data until it either makes a valid read, or has failed to make a valid read after repeated passes through the available image data.

[0231] 5.0 Operating Environment for Computer Implementations

[0232] FIG. 20 illustrates an example of a computer system that serves as an operating environment for software

implementations of the watermarking systems described above. The embedder and detector implementations are implemented in C/C++ and are portable to many different computer systems. FIG. 20 generally depicts one such system.

[0233] The computer system shown in FIG. 20 includes a computer 1220, including a processing unit 1221, a system memory 1222, and a system bus 1223 that interconnects various system components including the system memory to the processing unit 1221.

[0234] The system bus may comprise any of several types of bus structures including a memory bus or memory controller, a peripheral bus, and a local bus using a bus architecture such as PCI, VESA, Microchannel (MCA), ISA and EISA, to name a few.

[0235] The system memory includes read only memory (ROM) 1224 and random access memory (RAM) 1225. A basic input/output system 1226 (BIOS), containing the basic routines that help to transfer information between elements within the computer 1220, such as during start-up, is stored in ROM 1224.

[0236] The computer 1220 further includes a hard disk drive 1227, a magnetic disk drive 1228, e.g., to read from or write to a removable disk 1229, and an optical disk drive 1230, e.g., for reading a CD-ROM or DVD disk 1231 or to read from or write to other optical media. The hard disk drive 1227, magnetic disk drive 1228, and optical disk drive 1230 are connected to the system bus 1223 by a hard disk drive interface 1232, a magnetic disk drive interface 1233, and an optical drive interface 1234, respectively. The drives and their associated computer-readable media provide non-volatile storage of data, data structures, computer-executable instructions (program code such as dynamic link libraries, and executable files), etc. for the computer 1220.

[0237] Although the description of computer-readable media above refers to a hard disk, a removable magnetic disk and an optical disk, it can also include other types of media that are readable by a computer, such as magnetic cassettes, flash memory cards, digital video disks, and the like.

[0238] A number of program modules may be stored in the drives and RAM 1225, including an operating system 1235, one or more application programs 1236, other program modules 1237, and program data 1238.

[0239] A user may enter commands and information into the computer 1220 through a keyboard 1240 and pointing device, such as a mouse 1242. Other input devices may include a microphone, joystick, game pad, satellite dish, digital camera, scanner, or the like. A digital camera or scanner 43 may be used to capture the target image for the detection process described above. The camera and scanner are each connected to the computer via a standard interface 44. Currently, there are digital cameras designed to interface with a Universal Serial Bus (USB), Peripheral Component Interconnect (PCI), and parallel port interface. Two emerging standard peripheral interfaces for cameras include USB2 and 1394 (also known as firewire and iLink).

[0240] Other input devices may be connected to the processing unit 1221 through a serial port interface 1246 or

other port interfaces (e.g., a parallel port, game port or a universal serial bus (USB)) that are coupled to the system bus.

[0241] A monitor 1247 or other type of display device is also connected to the system bus 1223 via an interface, such as a video adapter 1248. In addition to the monitor, computers typically include other peripheral output devices (not shown), such as speakers and printers.

[0242] The computer 1220 operates in a networked environment using logical connections to one or more remote computers, such as a remote computer 1249. The remote computer 1249 may be a server, a router, a peer device or other common network node, and typically includes many or all of the elements described relative to the computer 1220, although only a memory storage device 1250 has been illustrated in FIG. 20. The logical connections depicted in FIG. 20 include a local area network (LAN) 1251 and a wide area network (WAN) 1252. Such networking environments are commonplace in offices, enterprise-wide computer networks, intranets and the Internet.

[0243] When used in a LAN networking environment, the computer 1220 is connected to the local network 1251 through a network interface or adapter 1253. When used in a WAN networking environment, the computer 1220 typically includes a modem 1254 or other means for establishing communications over the wide area network 1252, such as the Internet. The modem 1254, which may be internal or external, is connected to the system bus 1223 via the serial port interface 1246.

[0244] In a networked environment, program modules depicted relative to the computer 1220, or portions of them, may be stored in the remote memory storage device. The processes detailed above can be implemented in a distributed fashion, and as parallel processes. It will be appreciated that the network connections shown are exemplary and that other means of establishing a communications link between the computers may be used.

[0245] While the computer architecture depicted in FIG. 20 is similar to typical personal computer architectures, aspects of the invention may be implemented in other computer architectures, such as hand-held computing devices like Personal Digital Assistants, audio and/video players, network appliances, telephones, etc.

[0246] 6.0 Concluding Remarks

[0247] Having described and illustrated the principles of the technology with reference to specific implementations, it will be recognized that the technology can be implemented in many other, different, forms. The techniques for embedding and detecting a watermark may be applied to various types of watermarks, including those encoded using linear or non-linear functions to apply a watermark message to a host signal. As one example, embedding methods, such as methods for error correction coding, methods for mapping watermark messages to the host signal, and methods for redundantly encoding watermark messages apply whether the watermarking functions are linear or non-linear. In addition, the techniques for determining and refining a watermark's orientation apply to linear and non-linear watermark methods. For example, the methods described above for detecting orientation of a watermark signal in a potentially transformed version of the watermarked signal apply to water-

mark systems that use different methods for embedding and reading messages, including, but not limited to, techniques that modulate spatial or temporal domain intensity values, that modulate transform coefficients, that employ dither modulation or quantization index modulation.

[0248] Some of the detector methods described above invoke a watermark message reader to assess the merits of a given orientation of a watermark signal in a potentially transformed version of the watermarked signal. In particular, some of these techniques assess the merits of an orientation by invoking a reader to determine the extent to which known message bits agree with those read from the watermarked signal using the orientation. These techniques are not specific to the type of message encoding or reading as noted in the previous paragraph. The merits of a given estimate of a watermark signal's orientation may be assessed by selecting an orientation parameter that increases correlation between the watermark signal (or known watermark signal attributes) and the watermarked signal, or that improves recovery of known watermark message bits from the watermark signal.

[0249] Some watermark readers extract a message from a watermarked signal by correlating known attributes of a message symbol with the watermarked signal. For example, one symbol might be associated with a first pseudorandom noise pattern, while another symbol is associated with another pseudorandom noise pattern. If the reader determines that a strong correlation between the known attribute and the watermark signal exists, then it is likely that the watermarked signal contains the message symbol.

[0250] Other watermark readers analyze the watermarked signal to identify attributes that are associated with a message symbol. Generally speaking, these watermark readers are using a form of correlation, but in a different form. If the reader identifies evidence of watermark signal attributes associated with a message symbol, it notes that the associated message symbol has likely been encoded. For example, readers that employ quantization index modulation analyze the watermarked signal by applying quantizers to the signal to determine which quantizer was most likely used in the embedder to encode a message. Since message symbols are associated with quantizers, the reader extracts a message by estimating the quantizer used to encode the message. In these schemes, the signal attribute associated with a message symbol is the type of quantization applied to the signal. Regardless of the signal attributes used to encode and read a watermark message, the methods described above for determining watermark orientation and refining orientation parameters still apply.

[0251] Symmetric Patterns and Noise Patterns: Toward a Robust Universal Coding System

[0252] The placement of identification patterns into images is certainly not new. Logos stamped into corners of images, subtle patterns such as true signatures or the wall-papering of the copyright circle-C symbol, and the watermark proper are all examples of placing patterns into images in order to signify ownership or to try to prevent illicit uses of the creative material.

[0253] What does appear to be novel is the approach of placing independent "carrier" patterns, which themselves are capable of being modulated with certain information, directly into images and audio for the purposes of transmis-



sion and discernment of said information, while effectively being imperceptible and/or unintelligible to a perceiving human. Steganographic solutions currently known to the inventor all place this information “directly” into empirical data (possibly first encrypted, then directly), whereas the methods of this disclosure posit the creation of these (most-often) coextensive carrier signals, the modulation of those carrier signals with the information proper, THEN the direct application to the empirical data.

[0254] In extending these concepts one step further into the application arena of universal code systems, where a sending site transmits empirical data with a certain universal coding scheme employed and a receiving site analyzes said empirical data using the universal coding scheme, it would be advantageous to take a closer look at the engineering considerations of such a system designed for the transmission of images or motion images, as opposed to audio. Said more clearly, the same type of analysis of a specific implementation such as is contained in FIG. 9 of U.S. Pat. No. 5,748,783 and its accompanying discussion on the universal codes in audio applications should as well be done on imagery (or two dimensional signals). This section is such an analysis and outline of a specific implementation of universal codes and it attempts to anticipate various hurdles that such a method should clear.

[0255] The unifying theme of one implementation of a universal coding system for imagery and motion imagery is “symmetry.” The idea driving this couldn’t be more simple: a prophylactic against the use of image rotation as a means for less sophisticated pirates to bypass any given universal coding system. The guiding principle is that the universal coding system should easily be read no matter what rotational orientation the subject imagery is in. These issues are quite common in the fields of optical character recognition and object recognition, and these fields should be consulted for further tools and tricks in furthering the engineering implementation of this invention. As usual, an immediate example is in order.

[0256] Digital Video And Internet Company XYZ has developed a delivery system of its product which relies on a non-symmetric universal coding which double checks incoming video to see if the individual frames of video itself, the visual data, contain XYZ’s own relatively high security internal signature codes using the methods of this invention. This works well and fine for many delivery situations, including their Internet tollgate which does not pass any material unless both the header information is verified AND the in-frame universal codes are found. However, another piece of their commercial network performs mundane routine monitoring on Internet channels to look for unauthorized transmission of their proprietary creative property. They control the encryption procedures used, thus it is no problem for them to unencrypt creative property, including headers, and perform straightforward checks. A pirate group that wants to traffic material on XYZ’s network has determined how to modify the security features in XYZ’s header information system, and they have furthermore discovered that by simply rotating imagery by 10 or 20 degrees, and transmitting it over XYZ’s network, the network doesn’t recognize the codes and therefore does not flag illicit uses of their material, and the receiver of the pirate’s rotated material simply unrotates it.

[0257] Summarizing this last example via logical categories, the non-symmetric universal codes are quite acceptable for the “enablement of authorized action based on the finding of the codes,” whereas it can be somewhat easily by-passed in the case of “random monitoring (policing) for the presence of codes.” [Bear in mind that the non-symmetric universal codes may very well catch 90% of illicit uses, i.e. 90% of the illicit users wouldn’t bother even going to the simple by-pass of rotation.] To address this latter category, the use of quasi-rotationally symmetric universal codes is called for. “Quasi” derives from the age old squaring the circle issue, in this instance translating into not quite being able to represent a full incrementally rotational symmetric 2-D object on a square grid of pixels. Furthermore, basic considerations must be made for scale/magnification changes of the universal codes. It is understood that the monitoring process must be performed when the monitored visual material is in the “perceptual” domain, i.e. when it has been unencrypted or unscrambled and in the form with which it is (or would be) presented to a human viewer. Would-be pirates could attempt to use other simple visual scrambling and unscrambling techniques, and tools could be developed to monitor for these telltale scrambled signals. Said another way, would-be pirates would then look to transform visual material out of the perceptual domain, pass by a monitoring point, and then transform the material back into the perceptual domain; tools other than the monitoring for universal codes would need to be used in such scenarios. The monitoring discussed here therefore applies to applications where monitoring can be performed in the perceptual domain, such as when it is actually sent to viewing equipment.

[0258] The “ring” is the only full rotationally symmetric two dimensional object. The “disk” can be seen as a simple finite series of concentric and perfectly abutted rings having width along their radial axis. Thus, the “ring” needs to be the starting point from which a more robust universal code standard for images is found. The ring also will fit nicely into the issue of scale/magnification changes, where the radius of a ring is a single parameter to keep track of and account for. Another property of the ring is that even the case where differential scale changes are made to different spatial axes in an image, and the ring turns into an oval, many of the smooth and quasi-symmetric properties that any automated monitoring system will be looking for are generally maintained. Likewise, appreciable geometric distortion of any image will clearly distort rings but they can still maintain gross symmetric properties. Hopefully, more pedestrian methods such as simply “viewing” imagery will be able to detect attempted illicit piracy in these regards, especially when such lengths are taken to by-pass the universal coding system.

[0259] Rings to Knots

[0260] Having discovered the ring as the only ideal symmetric pattern upon whose foundation a full rotationally robust universal coding system can be built, we must turn this basic pattern into something functional, something which can carry information, can be read by computers and other instrumentation, can survive simple transformations and corruptions, and can give rise to reasonably high levels of security (probably not unbreakable, as the section on universal codes explained) in order to keep the economics of subversion as a simple incremental cost item.

[0261] One current preferred embodiment of the “ring-based” universal codes is what the inventor refers to as “knot patterns” or simply “knots,” in deference to woven Celtic knot patterns which were later refined and exalted in the works of Leonardo Da Vinci (e.g. Mona Lisa, or his knot engravings). Some rumors have it that these drawings of knots were indeed steganographic in nature, i.e. conveying messages and signatures; all the more appropriate. **FIGS. 21 and 22** explore some of the fundamental properties of these knots.

[0262] Two simple examples of knot patterns are depicted by the supra-radial knots, **850** and the radial knots **852**. The names of these types are based on the central symmetry point of the splayed rings and whether the constituent rings intersect this point, are fully outside it, or in the case of sub-radial knots the central point would be inside a constituent circle. The examples of **850** and **852** clearly show a symmetrical arrangement of **8** rings or circles. “Rings” is the more appropriate term, as discussed above, in that this term explicitly acknowledges the width of the rings along the radial axis of the ring. It is each of the individual rings in the knot patterns **850** and **852** which will be the carrier signal for a single associated bit plane in our N-bit identification word. Thus, the knot patterns **850** and **852** each are an **8**-bit carrier of information. Specifically, assuming now that the knot patterns **850** and **852** are luminous rings on a black background, then the “addition” of a luminous ring to an independent source image could represent a “1” and the “subtraction” of a luminous ring from an independent source image could represent a “0.” The application of this simple encoding scheme could then be replicated over and over as in **FIG. 22** and its mosaic of knot patterns, with the ultimate step of adding a scaled down version of this encoded (modulated) knot mosaic directly and coextensively to the original image, with the resultant being the distributable image which has been encoded via this universal symmetric coding method. It remains to communicate to a decoding system which ring is the least significant bit in our N-bit identification word and which is the most significant. One such method is to make a slightly ascending scale of radii values (of the individual rings) from the LSB to the MSB. Another is to merely make the MSB, say, 10% larger radius than all the others and to pre-assign counterclockwise as the order with which the remaining bits fall out. Yet another is to put some simple hash mark inside one and only one circle. In other words, there are a variety of ways with which the bit order of the rings can be encoded in these knot patterns.

[0263] The preferred embodiment for the decoding of, first of all checking for the mere existence of these knot patterns, and second, for the reading of the N-bit identification word, is as follows. A suspect image is first fourier transformed via the extremely common 2D FFT computer procedure. Assuming that we don’t know the exact scale of the knot patterns, i.e., we don’t know the radius of an elemental ring of the knot pattern in the units of pixels, and that we don’t know the exact rotational state of a knot pattern, we merely inspect (via basic automated pattern recognition methods) the resulting magnitude of the Fourier transform of the original image for telltale ripple patterns (concentric low amplitude sinusoidal rings on top of the spatial frequency profile of a source image). The periodicity of these rings, along with the spacing of the rings, will inform us that the universal knot patterns are or are not likely present, and their scale in pixels. Classical small signal detection methods can

be applied to this problem just as they can to the other detection methodologies of this disclosure. Common spatial filtering can then be applied to the fourier transformed suspect image, where the spatial filter to be used would pass all spatial frequencies which are on the crests of the concentric circles and block all other spatial frequencies. The resulting filtered image would be fourier transformed out of the spatial frequency domain back into the image space domain, and almost by visual inspection the inversion or non-inversion of the luminous rings could be detected, along with identification of the MSB or LSB ring, and the (in this case **8**) N-bit identification code word could be found. Clearly, a pattern recognition procedure could perform this decoding step as well.

[0264] The preceding discussion and the method it describes has certain practical disadvantages and shortcomings which will now be discussed and improved upon. The basic method was presented in a simple-minded fashion in order to communicate the basic principles involved.

[0265] Let’s enumerate a few of the practical difficulties of the above described universal coding system using the knot patterns. For one (1), the ring patterns are somewhat inefficient in their “covering” of the full image space and in using all of the information carrying capacity of an image extent. Second (2), the ring patterns themselves will almost need to be more visible to the eye if they are applied, say, in a straightforward additive way to an 8-bit black and white image. Next (3), the “**8**” rings of **FIG. 21, 850** and **852**, is a rather low number, and moreover, there is a 22 and one half degree rotation which could be applied to the figures which the recognition methods would need to contend with (360 divided by 8 divided by 2). Next (4), strict overlapping of rings would produce highly condensed areas where the added and subtracted brightness could become quite appreciable. Next (5), the 2D FFT routine used in the decoding is notoriously computationally cumbersome as well as some of the pattern recognition methods alluded to. Finally (6), though this heretofore described form of universal coding does not pretend to have ultra-high security in the classical sense of top security communications systems, it would nevertheless be advantageous to add certain security features which would be inexpensive to implement in hardware and software systems which at the same time would increase the cost of would-be pirates attempting to thwart the system, and increase the necessary sophistication level of those pirates, to the point that a would-be pirate would have to go so far out of their way to thwart the system that willfulness would be easily proven and hopefully subject to stiff criminal liability and penalty (such as the creation and distribution of tools which strip creative property of these knot pattern codes).

[0266] All of these items can be addressed and should continue to be refined upon in any engineering implementation of the principles of the invention. This disclosure addresses these items with the following current preferred embodiments.

[0267] Beginning with item number 3, that there are only **8** rings represented in **FIG. 21** is simply remedied by increasing the number of rings. The number of rings that any given application will utilize is clearly a function of the application. The trade-offs include but are not limited to: on the side which argues to limit the number of rings utilized,

there will ultimately be more signal energy per ring (per visibility) if there are less rings; the rings will be less crowded so that their discernment via automated recognition methods will be facilitated; and in general since they are less crowded, the full knot pattern can be contained using a smaller overall pixel extent, e.g. a 30 pixel diameter region of image rather than a 100 pixel diameter region. The arguments to increase the number of rings include: the desire to transmit more information, such as ASCII information, serial numbers, access codes, allowed use codes and index numbers, history information, etc.; another key advantage of having more rings is that the rotation of the knot pattern back into itself is reduced, thereby allowing the recognition methods to deal with a smaller range of rotation angles (e.g., 64 rings will have a maximum rotational displacement of just under 3 degrees, i.e. maximally dissimilar to its original pattern, where a rotation of about 5 and one half degrees brings the knot pattern back into its initial alignment; the need to distinguish the MSB/LSB and the bit plane order is better seen in this example as well). It is anticipated that most practical applications will choose between 16 and 128 rings, corresponding to N=16 to N=128 for the choice of the number of bits in the N-bit identification code word. The range of this choice would somewhat correlate to the overall radius, in pixels, allotted to an elemental knot pattern such as **850** or **852**.

[**0268**] Addressing the practical difficulty item number 4, that of the condensation of rings patterns at some points in the image and lack of ring patterns in others (which is very similar, but still distinct from, item 1, the inefficient covering), the following improvement can be applied. **FIG. 21** shows an example of a key feature of a "knot" (as opposed to a pattern of rings) in that where patterns would supposedly intersect, a virtual third dimension is posited whereby one strand of the knot takes precedence over another strand in some predefined way; see item **854**. In the terms of imagery, the brightness or dimness of a given intersection point in the knot patterns would be "assigned" to one and only one strand, even in areas where more than two strands overlap. The idea here is then extended, **864**, to how rules about this assignment should be carried out in some rotationally symmetric manner. For example, a rule would be that, travelling clockwise, an incoming strand to a loop would be "behind" an outgoing strand. Clearly there are a multitude of variations which could be applied to these rules, many which would critically depend on the geometry of the knot patterns chosen. Other issues involved will probably be that the finite width, and moreover the brightness profile of the width along the normal axis to the direction of a strand, will all play a role in the rules of brightness assignment to any given pixel underlying the knot patterns.

[**0269**] A major improvement to the nominal knot pattern system previously described directly addresses practical difficulties (1), the inefficient covering, (2) the unwanted visibility of the rings, and (6) the need for higher levels of security. This improvement also indirectly addresses item (4) the overlapping issue, which has been discussed in the last paragraph. This major improvement is the following: just prior to the step where the mosaic of the encoded knot patterns is added to an original image to produce a distributable image, the mosaic of encoded knot patterns, **866**, is spatially filtered (using common 2D FFT techniques) by a standardized and (generally smoothly) random phase-only spatial filter. It is very important to note that this phase-only

filter is itself fully rotationally symmetric within the spatial frequency domain, i.e. its filtering effects are fully rotationally symmetric. The effect of this phase-only filter on an individual luminous ring is to transform it into a smoothly varying pattern of concentric rings, not totally dissimilar to the pattern on water several instances after a pebble is dropped in, only that the wave patterns are somewhat random in the case of this phase-only filter rather than the uniform periodicity of a pebble wave pattern. **FIG. 23** attempts to give a rough (i.e. non-greyscale) depiction of these phase-only filtered ring patterns. The top figure of **FIG. 23** is a cross section of a typical brightness contour/profile **874** of one of these phase-only filtered ring patterns. Referenced in the figure is the nominal location of the pre-filtered outer ring center, **870**. The center of an individual ring, **872**, is referenced as the point around which the brightness profile is rotated in order to fully describe the two dimensional brightness distribution of one of these filtered patterns. Yet another rough attempt to communicate the characteristics of the filtered ring is depicted as **876**, a crude greyscale image of the filtered ring. This phase-only filtered ring, **876** will can be referred to as a random ripple pattern.

[**0270**] Not depicted in **FIG. 23** is the composite effects of phase-only filtering on the knot patterns of **FIG. 21**, or on the mosaic of knot patterns **866** in **FIG. 22**. Each of the individual rings in the knot patterns **850** or **852** will give rise to a 2D brightness pattern of the type **876**, and together they form a rather complicated brightness pattern. Realizing that the encoding of the rings is done by making it luminous (1) or "anti-luminous" (0), the resulting phase-only filtered knot patterns begin to take on subtle characteristics which no longer make direct sense to the human eye, but which are still readily discernable to a computer especially after the phase-only filtering is inverse filtered reproducing the original rings patterns.

[**0271**] Returning now to **FIG. 22**, we can imagine that an 8-bit identification word has been encoded on the knot patterns and the knot patterns phase-only filtered. The resulting brightness distribution would be a rich tapestry of overlapping wave patterns which would have a certain beauty, but would not be readily intelligible to the eye/brain. [An exception to this might draw from the lore of the South Pacific Island communities, where it is said that sea travelers have learned the subtle art of reading small and multiply complex ocean wave patterns, generated by diffracted and reflected ocean waves off of intervening islands, as a primary navigational tool.] For want of a better term, the resulting mosaic of filtered knot patterns (derived from **866**) can be called the encoded knot tapestry or just the knot tapestry. Some basic properties of this knot tapestry are that it retains the basic rotational symmetry of its generator mosaic, it is generally unintelligible to the eye/brain, thus raising it a notch on the sophistication level of reverse engineering, it is more efficient at using the available information content of a grid of pixels (more on this in the next section), and if the basic knot concepts **854** and **864** are utilized, it will not give rise to local "hot spots" where the signal level becomes unduly condensed and hence objectionably visible to a viewer.

[**0272**] The basic decoding process previously described would now need the additional step of inverse filtering the phase-only filter used in the encoding process. This inverse filtering is quite well known in the image processing indus-

try. Provided that the scale of the knot patterns are known a priori, the inverse filtering is straightforward. If on the other hand the scale of the knot patterns is not known, then an additional step of discovering this scale is in order. One such method of discovering the scale of the knot patterns is to iteratively apply the inverse phase-only filter to variously scaled version of an image being decoded, searching for which scale-version begins to exhibit noticeable knot patterning. A common search algorithm such as the simplex method could be used in order to accurately discover the scale of the patterns. The field of object recognition should also be consulted, under the general topic of unknown-scale object detection.

[0273] An additional point about the efficiency with which the knot tapestry covers the image pixel grid is in order. Most applications of the knot tapestry method of universal image coding will posit the application of the fully encoded tapestry (i.e. the tapestry which has the N-bit identification word embedded) at a relative low brightness level into the source image. In real terms, the brightness scale of the encoded tapestry will vary from, for example,  $-5$  grey scale values to  $5$  grey scale values in a typical  $256$  grey scale image, where the preponderance of values will be within  $-2$  and  $2$ . This brings up the purely practical matter that the knot tapestry will be subject to appreciable bit truncation error. Put as an example, imagine a constructed knot tapestry nicely utilizing a full  $256$  grey level image, then scaling this down by a factor of  $20$  in brightness including the bit truncation step, then rescaling this truncated version back up in brightness by the same factor of  $20$ , then inverse phase-only filtering the resultant. The resulting knot pattern mosaic will be a noticeably degraded version of the original knot pattern mosaic. The point of bringing all of this up is the following: it will be a simply defined, but indeed challenging, engineering task to select the various free parameters of design in the implementation of the knot tapestry method, the end goal being to pass a maximum amount of information about the N-bit identification word within some pre-defined visibility tolerance of the knot tapestry. The free parameters include but would not be fully limited to: the radius of the elemental ring in pixels,  $N$  or the number of rings, the distance in pixels from the center of a knot pattern to the center of an elemental ring, the packing criteria and distances of one knot pattern with the others, the rules for strand weaving, and the forms and types of phase-only filters to be used on the knot mosaics. It would be desirable to feed such parameters into a computer optimization routine which could assist in their selection. Even this would begin surely as more of an art than a science due to the many non-linear free parameters involved.

[0274] A side note on the use of phase-only filtering is that it can assist in the detection of the ring patterns. It does so in that the inverse filtering of the decoding process tends to "blur" the underlying source image upon which the knot tapestry is added, while at the same time "bringing into focus" the ring patterns. Without the blurring of the source image, the emerging ring patterns would have a harder time "competing" with the sharp features of typical images. The decoding procedure should also utilize the gradient thresholding method described in another section. Briefly, this is the method where if it is known that a source signal is much larger in brightness than our signature signals, then an image being decoded can have higher gradient areas thresholded in

the service of increasing the signal level of the signature signals relative to the source signal.

[0275] As for the other practical difficulty mentioned earlier, item (5) which deals with the relative computational overhead of the 2D FFT routine and of typical pattern recognition routines, the first remedy here posited but not filled is to find a simpler way of quickly recognizing and decoding the polarity of the ring brightnesses than that of using the 2D FFT. Barring this, it can be seen that if the pixel extent of an individual knot pattern (**850** or **852**) is, for example,  $50$  pixels in diameter, than a simple  $64$  by  $64$  pixel 2D FFT on some section of an image may be more than sufficient to discern the N-bit identification word as previously described. The idea would be to use the smallest image region necessary, as opposed to being required to utilize an entire image, to discern the N-bit identification word.

[0276] Another note is that those practitioners in the science of image processing will recognize that instead of beginning the discussion on the knot tapestry with the utilization of rings, we could have instead jumped right to the use of 2D brightness distribution patterns **876**, **QUA** bases functions. The use of the "ring" terminology as the baseline invention is partly didactic, as is appropriate for patent disclosures in any event. What is more important, perhaps, is that the use of true "rings" in the decoding process, post-inverse filtering, is probably the simplest form to input into typical pattern recognition routines.

[0277] Method for Embedding Subliminal Registration Patterns into Images and Other Signals

[0278] The very notion of reading embedded signatures involves the concept of registration. The underlying master noise signal must be known, and its relative position needs to be ascertained (registered) in order to initiate the reading process itself (e.g. the reading of the 1's and 0's of the N-bit identification word). When one has access to the original or a thumbnail of the unsigned signal, this registration process is quite straightforward. When one doesn't have access to this signal, which is often the case in universal code applications of this technology, then different methods must be employed to accomplish this registration step. The example of pre-marked photographic film and paper, where by definition there will never be an "unsigned" original, is a perfect case in point of the latter.

[0279] Many earlier sections have variously discussed this issue and presented certain solutions. Notably, the section on "simple" universal codes discusses one embodiment of a solution where a given master code signal is known a priori, but its precise location (and indeed, its existence or non-existence) is not known. That particular section went on to give a specific example of how a very low level designed signal can be embedded within a much larger signal, wherein this designed signal is standardized so that detection equipment or reading processes can search for this standardized signal even though its exact location is unknown. The brief section on 2D universal codes went on to point out that this basic concept could be extended into 2 dimensions, or, effectively, into imagery and motion pictures. Also, the section on symmetric patterns and noise patterns outlined yet another approach to the two dimensional case, wherein the nuances associated with two dimensional scale and rotation were more explicitly addressed. Therein, the idea was not merely to determine the proper orientation and scale

of underlying noise patterns, but to have information transmitted as well, e.g., the N rings for the N-bit identification word.

[0280] This section now attempts to isolate the sub-problem of registering embedded patterns for registration's sake. Once embedded patterns are registered, we can then look again at how this registration can serve broader needs. This section presents yet another technique for embedding patterns, a technique which can be referred to as "subliminal digital graticules". "Graticules"—other words such as fiducials or reticles or hash marks could just as well be used—conveys the idea of calibration marks used for the purposes of locating and/or measuring something. In this case, they are employed as low level patterns which serve as a kind of gridding function. That gridding function itself can be a carrier of 1 bit of information, as in the universal 1 second of noise (its absence or presence, copy me, don't copy me), or it can simply find the orientation and scale of other information, such as embedded signatures, or it can simply orient an image or audio object itself.

[0281] FIGS. 24 and 25 visually summarize two related methods which illustrate applicant's subliminal digital graticules. As will be discussed, the method of FIG. 24 may have slight practical advantages over the method outlined in FIG. 25, but both methods effectively decompose the problem of finding the orientation of an image into a series of steps which converge on a solution. The problem as a whole can be simply stated as the following: given an arbitrary image wherein a subliminal digital graticule may have been stamped, then find the scale, rotation, and origin (offset) of the subliminal digital graticule.

[0282] The beginning point for subliminal graticules is in defining what they are. Simply put, they are visual patterns which are directly added into other images, or as the case may be, exposed onto photographic film or paper. The classic double exposure is not a bad analogy, though in digital imaging this specific concept becomes rather stretched. These patterns will generally be at a very low brightness level or exposure level, such that when they are combined with "normal" images and exposures, they will effectively be invisible (subliminal) and just as the case with embedded signatures, they will by definition not interfere with the broad value of the images to which they are added.

[0283] FIGS. 24 and 25 define two classes of subliminal graticules, each as represented in the spatial frequency domain, also known as the UV plane, 1000. Common two dimensional fourier transform algorithms can transform any given image into its UV plane conjugate. To be precise, the depictions in FIGS. 24 and 25 are the magnitudes of the spatial frequencies, whereas it is difficult to depict the phase and magnitude which exists at every point.

[0284] FIG. 24 shows the example of six spots in each quadrant along the 45 degree lines, 1002. These are exaggerated in this figure, in that these spots would be difficult to discern by visual inspection of the UV plane image. A rough depiction of a "typical" power spectrum of an arbitrary image as also shown, 1004. This power spectrum is generally as unique as images are unique. The subliminal graticules are essentially these spots. In this example, there are six spatial frequencies combined along each of the two 45 degree axes. The magnitudes of the six frequencies can be the same or different (we'll touch upon this refinement

later). Generally speaking, the phases of each are different from the others, including the phases of one 45 degree axis relative to the other. FIG. 26 depicts this graphically. The phases in this example are simply randomly placed between  $\text{PI}$  and  $-\text{PI}$ , 1008 and 1010. Only two axes are represented in FIG. 26—as opposed to the four separate quadrants, since the phase of the mirrored quadrants are simply  $\text{PI}/2$  out of phase with their mirrored counterparts. If we turned up the intensity on this subliminal graticule, and we transformed the result into the image domain, then we would see a weave-like cross-hatching pattern as related in the caption of FIG. 24. As stated, this weave-like pattern would be at a very low intensity when added to a given image. The exact frequencies and phases of the spectral components utilized would be stored and standardized. These will become the "spectral signatures" that registration equipment and reading processes will seek to measure.

[0285] Briefly, FIG. 25 has a variation on this same general theme. FIG. 25 lays out a different class of graticules in that the spectral signature is a simple series of concentric rings rather than spots along the 45 degree axes. FIG. 27 then depicts a quasi-random phase profile as a function along a half-circle (the other half of the circle then being  $\text{PI}/2$  out of phase with the first half). These are simple examples and there are a wide variety of variations possible in designing the phase profiles and the radii of the concentric rings. The transform of this type of subliminal graticule is less of a "pattern" as with the weave-like graticule of FIG. 24, where it has more of a random appearance like a snowy image.

[0286] The idea behind both types of graticules is the following: embed a unique pattern into an image which virtually always will be quite distinct from the imagery into which it will be added, but which has certain properties which facilitate fast location of the pattern, as well as accuracy properties such that when the pattern is generally located, its precise location and orientation can be found to some high level of precision. A corollary to the above is to design the pattern such that the pattern on average minimally interferes with the typical imagery into which it will be added, and has maximum energy relative to the visibility of the pattern.

[0287] Moving on to the gross summary of how the whole process works, the graticule type of FIG. 24 facilitates an image processing search which begins by first locating the rotation axes of the subliminal graticule, then locating the scale of the graticule, then determining the origin or offset. The last step here identifies which axes is which of the two 45 degree axes by determining phase. Thus even if the image is largely upside down, an accurate determination can be made. The first step and the second step can both be accomplished using only the power spectrum data, as opposed to the phase and magnitude. The phase and magnitude signals can then be used to "fine tune" the search for the correct rotation angle and scale. The graticule of FIG. 25 switches the first two steps above, where the scale is found first, then the rotation, followed by precise determination of the origin. Those skilled in the art will recognize that determining these outstanding parameters, along two axes, are sufficient to fully register an image. The "engineering optimization challenge" is to maximize the uniqueness and brightness of the patterns relative to their visibility, while minimizing the computational overhead in reaching some

specified level of accuracy and precision in registration. In the case of exposing photographic film and paper, clearly an additional engineering challenge is the outlining of economic steps to get the patterns exposed onto the film and paper in the first place, a challenge which has been addressed in previous sections.

[0288] The problem and solution as above defined is what was meant by registration for registration's sake. It should be noted that there was no mention made of making some kind of value judgement on whether or not a graticule is indeed being found or not. Clearly, the above steps could be applied to images which do not in fact have graticules inside them, the measurements then simply chasing noise. Sympathy needs to be extended to the engineer who is assigned the task of setting "detection thresholds" for these types of patterns, or any others, amidst the incredibly broad range of imagery and environmental conditions in which the patterns must be sought and verified. [Ironically, this is where the pure universal one second of noise stood in a previous section, and that was why it was appropriate to go beyond merely detecting or not detecting this singular signal, i.e. adding additional information planes]. Herein is where some real beauty shows up: in the combination of the subliminal graticules with the now-registered embedded signatures described in other parts of this disclosure. Specifically, once a "candidate registration" is found—paying due homage to the idea that one may be chasing noise—then the next logical step is to perform a reading process for, e.g., a 64 bit universal code signature. As further example, we can imagine that 44 bits of the 64 bit identification word are assigned as an index of registered users—serial numbers if you will. The remaining 20 bits are reserved as a hash code—as is well known in encryption arts—on the 44 bit identification code thus found. Thus, in one swoop, the 20 bits serve as the "yes, I have a registered image" or "no, I don't" answer. More importantly, perhaps, this allows for a system which can allow for maximum flexibility in precisely defining the levels of "false positives" in any given automated identification system. Threshold based detection will always be at the mercy of a plethora of conditions and situations, ultimately resting on arbitrary decisions. Give me N coin flips any day.

[0289] Back on point, these graticule patterns must first be added to an image, or exposed onto a piece of film. An exemplary program reads in an arbitrarily sized digital image and adds a specified graticule to the digital image to produce an output image. In the case of film, the graticule pattern would be physically exposed onto the film either before, during, or after exposure of the primary image. All of these methods have wide variations in how they might be accomplished.

[0290] The searching and registering of subliminal graticules is the more interesting and involved process. This section will first describe the elements of this process, culminating in the generalized flow chart of FIG. 32.

[0291] FIG. 28 depicts the first major "search" step in the registration process for graticules of the type in FIG. 24. A suspect image (or a scan of a suspect photograph) is first transformed in its fourier representation using well known 2D FFT routines. The input image may look like the one in FIG. 31, upper left image. FIG. 28 conceptually represents the case where the image and hence the graticules have not

been rotated, though the following process fully copes with rotation issues. After the suspect image has been transformed, the power spectrum of the transform is then calculated, being simply the square root of the addition of the two squared moduli. It is also a good idea to perform a mild low pass filter operation, such as a 3x3 blur filter, on the resulting power spectrum data, so that later search steps don't need incredibly fine spaced steps. Then the candidate rotation angles from 0 through 90 degrees (or 0 to  $\pi/2$  in radian) are stepped through. Along any given angle, two resultant vectors are calculated, the first is the simple addition of power spectrum values at a given radius along the four lines emanating from the origin in each quadrant. The second vector is the moving average of the first vector. Then, a normalized power profile is calculated as depicted in both 1022 and 1024, the difference being that one plot is along an angle which does not align with the subliminal graticules, and the other plot does align. The normalization stipulates that the first vector is the numerator and the second vector is the denominator in the resultant vector. As can be seen in FIG. 28, 1022 and 1024, a series of peaks (which should be "six" instead of "five" as is drawn) develops when the angle aligns along its proper direction. Detection of these peaks can be effected by setting some threshold on the normalized values, and integrating their total along the whole radial line. A plot, 1026, from 0 to 90 degrees is depicted in the bottom of FIG. 28, showing that the angle 45 degrees contains the most energy. In practice, this signal is often much lower than that depicted in this bottom figure, and instead of picking the highest value as the "found rotation angle," one can simply find the top few candidate angles and submit these candidates to the next stages in the process of determining the registration. It can be appreciated by those practiced in the art that the foregoing was simply a known signal detection scheme, and that there are dozens of such schemes that can ultimately be created or borrowed. The simple requirement of the first stage process is to whittle down the candidate rotation angles to just a few, wherein more refined searches can then take over.

[0292] FIG. 29 essentially outlines the same type of gross searching in the power spectral domain. Here instead we first search for the gross scale of the concentric rings, stepping from a small scale through a large scale, rather than the rotation angle. The graph depicted in 1032 is the same normalized vector as in 1022 and 1024, but now the vector values are plotted as a function of angle around a semi-circle. The moving average denominator still needs to be calculated in the radial direction, rather than the tangential direction. As can be seen, a similar "peaking" in the normalized signal occurs when the scanned circle coincides with a graticule circle, giving rise to the plot 1040. The scale can then be found on the bottom plot by matching the known characteristics of the concentric rings (i.e. their radii) with the profile in 1040.

[0293] FIG. 30 depicts the second primary step in registering subliminal graticules of the type in FIG. 24. Once we have found a few rotation candidates from the methods of FIG. 28, we then take the plots of the candidate angles of the type of 1022 and 1024 and perform what the inventor refers to as a "scaled kernel" matched filtering operation on those vectors. The scaled kernel refers to the fact that the kernel in this case is a known non-harmonic relationship of frequencies, represented as the lines with x's at the top in 1042 and 1044, and that the scale of these frequencies is swept

through some pre-determined range, such as 25% to 400% of some expected scale at 100%. The matched filter operation simply adds the resultant multiplied values of the scaled frequencies and their plot counterparts. Those practiced in the art will recognize the similarity of this operation with the very well known matched filter operation. The resulting plot of the matched filter operation will look something like **1046** at the bottom of **FIG. 30**. Each candidate angle from the first step will generate its own such plot, and at this point the highest value of all of the plots will become our candidate scale, and the angle corresponding to the highest value will become our primary candidate rotation angle. Likewise for graticules of the type in **FIG. 25**, a similar "scaled-kernel" matched filtering operation is performed on the plot **1040** of **FIG. 29**. This generally provides for a single candidate scale factor. Then, using the stored phase plots **1012**, **1014** and **1016** of **FIG. 27**, a more traditional matched filtering operation is applied between these stored plots (as kernels), and the measured phase profiles along the half-circles at the previously found scale.

[0294] The last step in registering graticules of the type in **FIG. 24** is to perform a garden variety matched filter between the known (either spectral or spatial) profile of the graticule with the suspect image. Since both the rotation, scale and orientation are now known from previous steps, this matched filtering operation is straightforward. If the accuracies and precision of preceding steps have not exceeded design specifications in the process, then a simple micro-search can be performed in the small neighborhood about the two parameters scale and rotation, a matched filter operation performed, and the highest value found will determine a "fine tuned" scale and rotation. In this way, the scale and rotation can be found to within the degree set by the noise and cross-talk of the suspect image itself. Likewise, once the scale and rotation of the graticules of the type in **FIG. 25** are found, then a straightforward matched filter operation can complete the registration process, and similar "fine tuning" can be applied.

[0295] Moving on to a variant of the use of the graticules of the type in **FIG. 24**, **FIG. 31** presents the possibility for finding the subliminal graticules without the need for performing a computationally expensive 2 dimensional FFT (fast fourier transform). In situations where computational overhead is a major issue, then the search problem can be reduced to a series of one-dimensional steps. **FIG. 31** broadly depicts how to do this. The figure at the top left is an arbitrary image in which the graticules of the type of **FIG. 24** have been embedded. Starting at angle 0, and finishing with an angle just below 180 degrees, and stepping by, for example 5 degrees, the grey values along the depicted columns can be simply added to create a resulting column-integrated scan, **1058**. The figure in the top right, **1052**, depicts one of the many angles at which this will be performed. This column-integrated scan then is transformed into its fourier representation using the less computationally expensive 1 dimensional FFT. This is then turned into a magnitude or "power" plot (the two are different), and a similar normalized vector version created just like **1022** and **1024** in **FIG. 28**. The difference now is that as the angle approaches the proper angles of the graticules, slowly the tell-tale peaks begin to appear in the **1024**-like plots, but they generally show up at higher frequencies than expected for a given scale, since we are generally slightly off on our rotation. It remains to find the angle which maximizes the

peak signals, thus zooming in on the proper rotation. Once the proper rotation is found, then the scaled kernel matched filter process can be applied, followed by traditional matched filtering, all as previously described. Again, the sole idea of the "short-cut" of **FIG. 31** is to greatly reduce the computational overhead in using the graticules of the type in **FIG. 24**. The inventor has not reduced this method of **FIG. 31** to practice and currently has no data on precisely how much computational savings, if any, will be realized. These efforts will be part of application specific development of the method.

[0296] **FIG. 32** simply summarizes, in order of major process steps, the methods revolving around the graticules of the type in **FIG. 24**.

[0297] In another variant embodiment, the graticule energy is not concentrated around the 45 degree angles in the spatial frequency domain. (Some compression algorithms, such as JPEG, tend to particularly attenuate image energy at this orientation.) Instead, the energy is more widely spatially spread. **FIG. 24A** shows one such distribution. The frequencies near the axes, and near the origin are generally avoided, since this is where the image energy is most likely concentrated.

[0298] Detection of this energy in a suspect image again relies on techniques like that reviewed above. However, instead of first identifying the axes, then the rotation, and then the scale, a more global pattern matching procedure is performed in which all are determined in a brute force effort. Those skilled in the art will recognize that the Fourier-Mellin transform is well suited for use in such pattern matching problems.

[0299] The foregoing principles find application, for example, in photo-duplication kiosks. Such devices typically include a lens for imaging a customer-provided original (e.g. a photographic print or film) onto an opto-electronic detector, and a print-writing device for exposing and developing an emulsion substrate (again photographic paper or film) in accordance with the image data gathered by the detector. The details of such devices are well known to those skilled in the art, and are not belabored here.

[0300] In such systems, a memory stores data from the detector, and a processor (e.g. a Pentium microprocessor with associated support components) can be used to process the memory data to detect the presence of copyright data steganographically encoded therein. If such data is detected, the print-writing is interrupted.

[0301] To avoid defeat of the system by manual rotation of the original image off-axis, the processor desirably implements the above-described technique to effect automatic registration of the original, notwithstanding scale, rotation, and origin offset factors. If desired, a digital signal processing board can be employed to offload certain of the FFT processing from the main (e.g. Pentium) processor. After a rotated/scaled image is registered, detection of any steganographically encoded copyright notice is straightforward and assures the machine will not be used in violation of a photographer's copyright.

[0302] While the techniques disclosed above have make use of applicant's preferred steganographic encoding methods, the principles thereof are more widely applicable and

can be used in many instances in which automatic registration of an image is to be effected.

1. An image processor for embedding in image data a digital watermark including position and rotation information, using an irrotationally symmetric pattern arrangement.

2. The image processor according to claim 1, wherein said irrotationally symmetric pattern arrangement is a two-dimensional matrix constituted by  $m \times n$  elements.

3. The image processor according to claim 1, wherein said irrotationally symmetric pattern arrangement is a pattern arrangement for which the positive or negative symbols of each corresponding elements are not wholly the same if the pattern arrangement is rotated at an arbitrary angle (except for angles of 360 degrees multiplied by an integer number).

4. An image processor capable of extracting digital watermark information from image data in which said digital watermark information including position and rotation information is embedded, comprising: rotation information searching means for performing processing for extracting said rotation information from the image data with said digital watermark information embedded therein, for a plurality of rotation angles different from one another; position information searching means for performing processing for extracting said position information from said image data, for a plurality of start-of-extraction positions different from one another; calculating means for calculating confidence coefficients indicating accuracy as to whether said position and rotation information is extracted, for each information searched by said rotation information searching means and position information searching means and extracted as position and rotation information; and determining means for determining the position and rotation angle at which said digital watermark information is embedded in said image data, based on the confidence coefficient calculated by said calculating means.

5. The image processor according to claim 4, wherein said digital watermark information includes said position information and rotation information and usage information, and said usage information includes the ID of a device or the user ID.

6. The image processor according to claim 4, wherein said digital watermark information includes said position information and rotation information and usage information, and said usage information includes information for controlling a device.

7. The image processor according to claim 4, wherein said calculating means calculates confidence coefficients by performing computation of said image data with a matrix constituted by  $m \times n$  coefficients.

8. The image processor according to claim 7, wherein said matrix computation processing is convolution computation.

9. The image processor according to claim 4, further comprising: extracting means for extracting the digital watermark information embedded in said image data, based on the position in said image data on the basis of the result of determination by said determining means.

10. An image processing method of embedding in an image digital watermark information including position and rotation information, using an irrotationally symmetric pattern arrangement.

11. An image processing method of extracting digital watermark information from image data in which said digital watermark information including position and rotation information is embedded, comprising: a rotation infor-

mation searching step of performing processing for extracting said rotation information from the image data with said digital watermark information embedded therein, for a plurality of rotation angles different from one another; a position information searching step of performing processing for extracting said position information from said image data, for a plurality of start-of-extraction positions different from one another; a calculating step of calculating confidence coefficients indicating accuracy as to whether said position and rotation information is extracted, for each information searched in said rotation information searching step and position information searching step and extracted as position and rotation information; and a determining step of determining the position and rotation angle at which said digital watermark information is embedded in said image data, based on the confidence coefficient calculated in said calculating step.

12. A computer program product embodying a program for implementing an image processing method of embedding in an image digital watermark information including position and rotation information, using an irrotationally symmetric pattern arrangement.

13. A computer program product embodying a program for implementing an image processing method of extracting digital watermark information from image data in which said digital watermark information including position and rotation information is embedded, the program comprising: program codes for a rotation information searching step of performing processing for extracting said rotation information from the image data with said digital watermark information embedded therein, for a plurality of rotation angles different from one another; program codes for a position information searching step of performing processing for extracting said position information from said image data, for a plurality of start-of-extraction positions different from one another; program codes for a calculating step of calculating confidence coefficients indicating accuracy as to whether said position and rotation information is extracted, for each information searched in said rotation information searching step and position information searching step and extracted as position and rotation information; and program codes for a determining step of determining the position and rotation angle at which said digital watermark information is embedded in said image data, based on the confidence coefficient calculated in said calculating step.

14. A computer data signal used for implementing an image processing method of embedding in an image digital watermark information including position and rotation information, using an irrotationally symmetric pattern arrangement.

15. The computer data signal of claim 14 embodied in a propagating wave.

16. A computer data signal used for implementing an image processing method of extracting digital watermark information from image data in which said digital watermark information including position and rotation information is embedded, comprising: code signals for use in a rotation information searching step of performing processing for extracting said rotation information from the image data with said digital watermark information embedded therein, for a plurality of rotation angles different from one another; code signals for use in a position information searching step of performing processing for extracting said position information from said image data, for a plurality of



start-of-extraction positions different from one another; code signals for use in a calculating step of calculating confidence coefficients indicating accuracy as to whether said position and rotation information is extracted, for each information searched in said rotation information searching step and position information searching step and extracted as position and rotation information; and code signals for use in a determining step of determining the position and rotation

angle at which said digital watermark information is embedded in said image data, based on the confidence coefficient calculated in said calculating step.

17. The computer data signal of claim 16, embodied in a propagating wave.

\* \* \* \* \*