HIC SCHEVRLINA SIMVL TVCHERINAQ, SIGNA REFVLGENT
QVE DOCTOR GEMINI SCHEVRLE PARENTIS HABES

Book cover plate showing the crest of the Scheurl family (16th century).
(Courtesy, Lutherhalle, Wittenberg, Germany.)

Whose copyright is it, and who can use it—legally? Watermarks are beginning to help confirm copyrights, authentic ownership, fingerprints, and data integrity.

Protecting Digital Media Content

Nasir Memon and Ping Wah Wong

Thanks to the proliferation of the World-Wide Web, a huge amount of multimedia content—text, graphics, images, video, and audio—are available for browsing and downloading by millions of users worldwide over the network. As a result, security and copyright issues have become important problems in research and applications. Digital watermarking is applicable in copyright protection, ownership assertion, and integrity checks in multimedia content.

Current copyright laws are perhaps inadequate for dealing with digital data, leading to interest in developing other copyright protection mechanisms, including digital watermarking techniques. A watermark is a signal added to digital data (audio, video, or still images) that can be detected or extracted later to make an assertion about the data. The watermark signal can serve various purposes, including:

- Ownership assertion. In order to establish ownership over some content (such as an image), Alice can use a private key to generate a watermark and embed it into the original image. She then makes the watermarked image publicly available. Later, when Bob claims he owns an image derived from this public image, Alice can produce the unmarked original and demonstrate the presence of her watermark in Bob's image. Since Alice's original image is unavailable to Bob, he cannot do the same. For such a scheme to work, the watermark has to survive common image-processing operations. It also needs to be a function of the original image to avoid counterfeiting attacks.
- Fingerprinting. To avoid unauthorized duplication and distribution of publicly available multimedia content, an author can embed a distinct watermark (or fingerprint) into each copy of the data. If unauthorized copies are found later, the origin of the copy can be determined by retrieving the fingerprint. In this application, the watermark needs to be invisible and invulnerable to deliberate attempts at forgery, removal, or invalidation.
- Authentication and integrity verification. When multimedia content is used for legal purposes, medical applications, news reporting, or commercial transactions, the originator of the content has

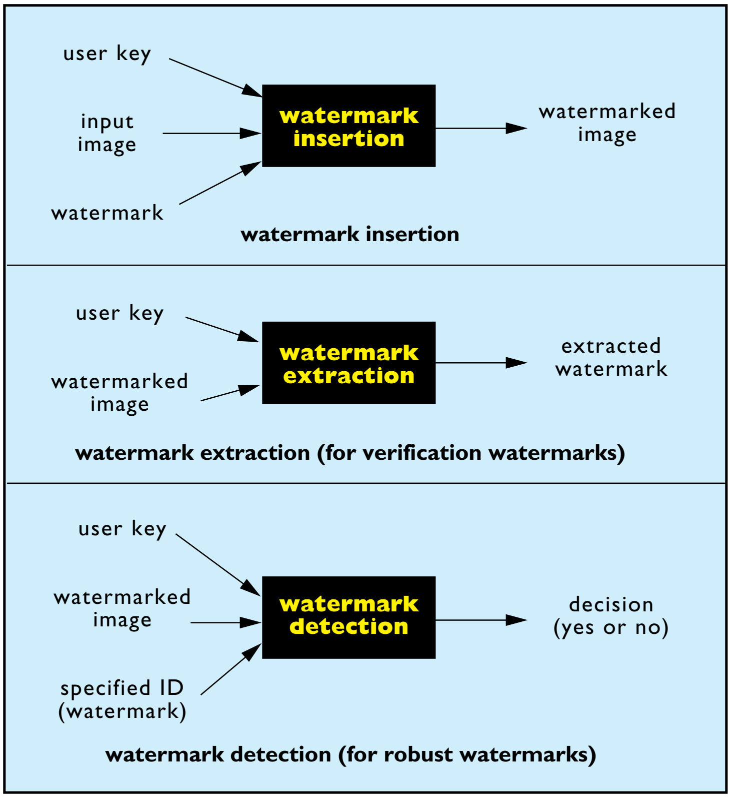

Figure 1. General watermarking procedures. Watermark insertion integrates the input image and a watermark to form the output watermarked image. Watermark extraction uncovers the watermark in watermarked images, a technique usually applicable in verification watermarks. For robust watermarks, the presence of a specified ID (watermark) can be detected using a predefined threshold; a yes or no answer indicates the presence of the ID, depending on whether the output from a signal detection block exceeds the given threshold. All of these procedures may also require a user key (either a public key or a secret key), depending on the specific algorithm.

to be verified while ensuring the content has not been changed, manipulated, or falsified. Although authentication of multimedia content can be done through conventional cryptographic techniques, the advantage of using a verification watermark is that the authenticator is inseparably bound to the content, hence simplifying the logistical problem of data handling. When the watermarked data is checked, the watermark is extracted using a unique key associated with the source, and the integrity of the data is verified through the integrity of the extracted watermark.

- Content labeling. The bits embedded into the data comprise an annotation, giving further information about the contents. For example, a photographic image could be annotated to describe the time and place the photograph was taken, a procedure that could be done automatically by the processor in the camera.
- Usage control. In a closed system in which the multimedia content needs special hardware for copying and viewing, a digital watermark can be inserted to indicate the number of copies permitted. Every time a copy is made, the watermark can be modified by the hardware, and at some point the hardware would not create any more copies of the data. An example is the digital video disc (DVD).
- Content protection. In certain applications, a content owner may want to publicly and freely provide a preview of the multimedia content being sold. To make the preview commercially worthless, the content could be stamped with a visible watermark very difficult to remove in an automated way.

The specific requirements of each watermarking technique vary with the application. There is no universal watermarking technique that satisfies all requirements of all applications. Consequently, each watermarking technique has to be designed within the context of the entire system in which it is to be used. Here we describe some watermarking techniques, focusing on image watermarking. However, many of these techniques are generally applicable to other forms of content, including video and audio.

General Framework

We use several parameters for categorizing watermarking techniques. First, a watermark can be visible or invisible.[1] A visible watermark typically contains a visual message or a company logo indicating ownership of the image. An invisible watermarked image is visually very similar but not necessarily identical to the original unmarked image. The existence of such a watermark can be determined only through a watermark extraction or detection algorithm.

Watermarking techniques can also be classified as fragile or robust. Fragile watermarks are easily corrupted by any form of image-processing procedure.

[1] In a strict sense, the terms "visible" and "invisible" are inappropriate in watermarking of multimedia data, such as sound clips. In a wider sense, we mean "perceptible" and "imperceptible." However, the ideas supporting watermarking of image data also apply to other forms of multimedia data.

Watermarks for image integrity checks, in which a change has to be detected or spatially localized, are necessarily fragile [10, 11]. Robust watermarks resist common image-manipulation procedures and are useful for ownership assertion purposes.

Watermark insertion. The top portion of Figure 1 shows the general procedure for inserting a watermark into an image. The watermark is generally a function of a number of factors, including user identity (an ID), a user key, the original image and its parameters (such as image size), and possibly others. The specific steps involved in watermark insertion depend on the watermarking technique being used; there are many applications of watermarking, and different watermarking schemes are generally necessary for different applications. Hence we can expect considerable variation in how watermark signals are inserted into images.

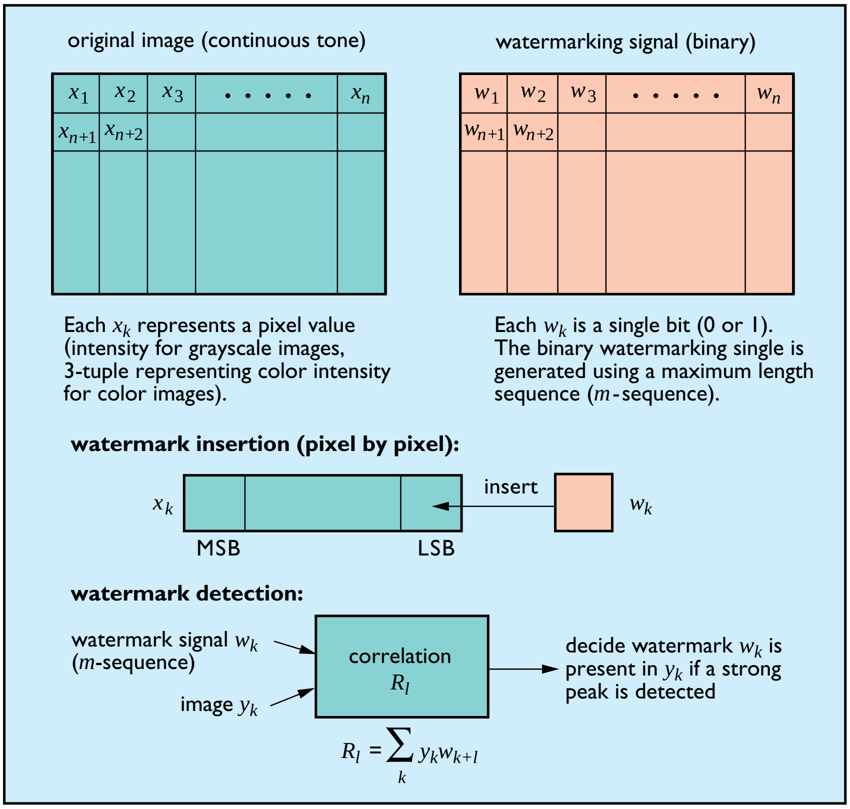
original image (continuous tone)
x_1 x_2 x_3 . . . . . x_n
x_{n+1} x_{n+2}
watermarking signal (binary)
w_1 w_2 w_3 . . . . . w_n
w_{n+1} w_{n+2}
Each x_k represents a pixel value (intensity for grayscale images, 3-tuple representing color intensity for color images).
Each w_k is a single bit (0 or 1). The binary watermarking single is generated using a maximum length sequence (m-sequence).
watermark insertion (pixel by pixel):
x_k
MSB
LSB
insert
w_k
watermark detection:
watermark signal w_k (m-sequence)
image y_k
correlation R_l
decide watermark w_k is present in y_k if a strong peak is detected
$R_l = \sum_k y_k w_{k+l}$

Figure 2. A simple watermarking technique proposed in [9]. Though not robust, it illustrates several important concepts in watermarking, such as using an m-sequence as the watermarking signal and a correlation function as the detection mechanism.

To illustrate how a watermark can be inserted in an image, we describe a simple watermarking method proposed in [9]. While it is useful for carrying ownership information, this watermark is not robust. In fact, it can be destroyed easily, because it is embedded in the least significant bit (LSB) of the image. We describe it in some detail here because [9] is historic as the first watermarking paper presented at a major conference and illustrates several important concepts in watermarking.

Let x_1, x_2, . . . be the pixels of an image (see Figure 2). To insert a watermark, we first generate a watermarking signal using a key to seed a generator for an *m*-sequence (a maximum-length random sequence). An *m*-sequence generated by an n^{th}-order linear shift register has a maximal length of $2^n - 1$, and its autocorrelation properties resemble those of random noise [6]. The elements (binary valued) of this *m*-sequence are then arranged into a 2D watermarking signal, as shown in Figure 2. This signal is then inserted pixel by pixel into the LSB position of the original image x_k. Since the watermarking signal is located at the LSB, it is invisible. For the same reason, it is also not very robust, in the sense that it can be removed easily.

The two major classes of robust watermarking techniques are private and oblivious, or public. A private scheme requires an original or reference image in the watermark detection procedure; an oblivious scheme does not.

Oblivious schemes are attractive for many applications. When we want to detect copyright violations in an image archive or in images posted on the Internet, we can use software agents, such as Webcrawlers, to perform identity checks for as many images as we can locate. Schemes requiring the original image to detect a watermark are not suitable for such applications.

Watermarking schemes that require a user key can be classified as either secret-key or public-key schemes. Secret-key schemes use the same user key for watermark insertion and extraction or detection procedures. As a result, secret-key schemes require secure communication between the image owner and the image receiver, or user, to pass the key information. Public-key schemes use separate keys for watermark insertion and extraction, or detection. A

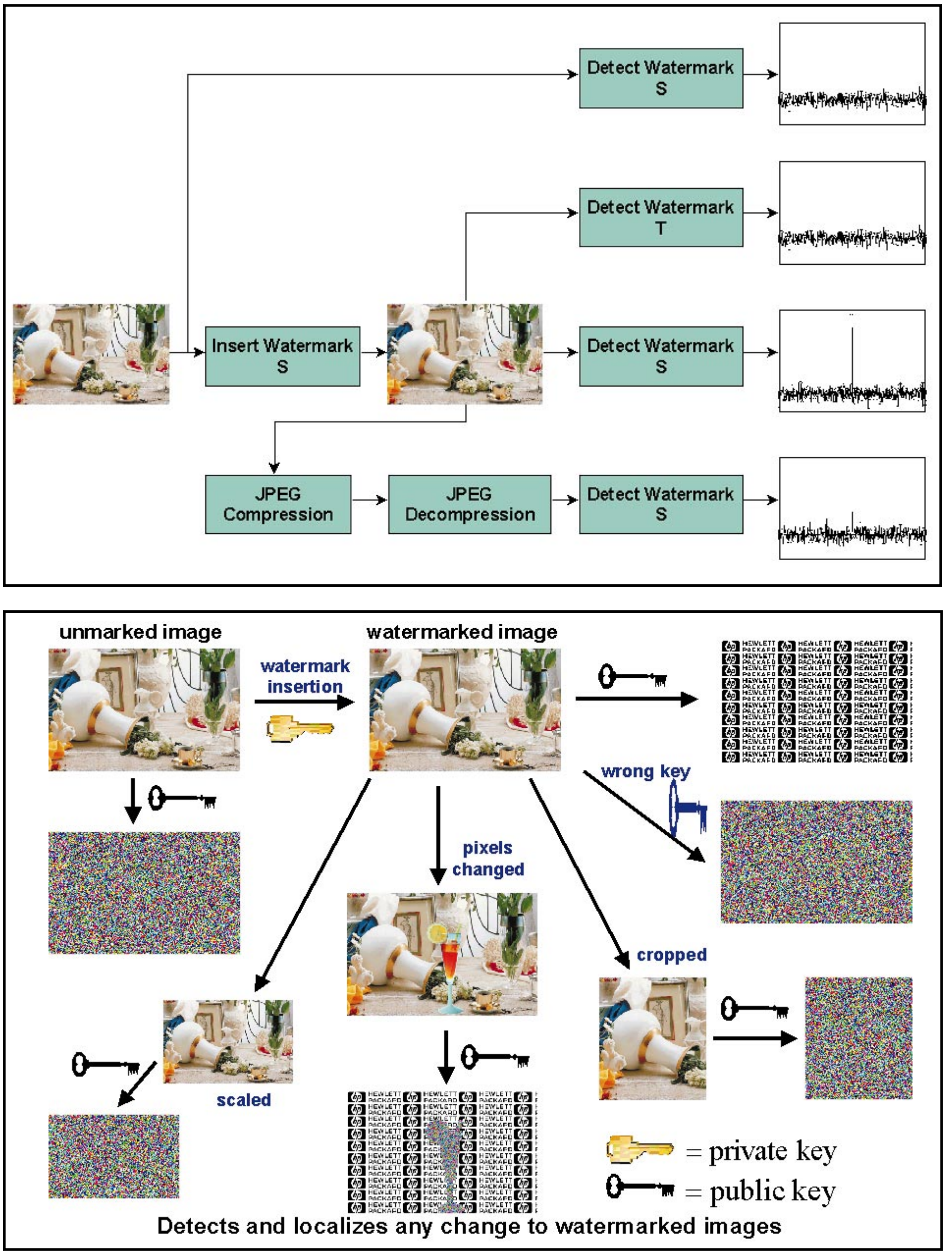

Figure 3. Properties of different kinds of watermarks. (a) A robust watermark, showing that an image with a watermark S exhibits strong correlation with S (even if the watermarked image is changed by lossy JPEG compression and decompression, there is still significant correlation). (b) A public-key verification watermark, showing that a verification watermark can detect and localize any change made to the watermarked image, whether the change is in pixel values or in image size. We can extract a completely correct watermark only if we have the appropriate key and if the watermarked image has not been changed in any way.

private key (known only to the owner) is typically used for watermark insertion. The watermarked image can be checked using a public key known to everybody. As a result, anyone can perform the watermark extraction or detection procedure, but no one but the owner can insert the watermark.

Watermark detection and extraction. A watermark must be detectable or extractable to be useful. Depending on the way the watermark is inserted, and depending on the nature of the watermarking algorithm, the method used can involve very distinct approaches. In some watermarking schemes, a watermark can be extracted in its exact form, a procedure we call watermark *extraction*. In other cases, we can detect only whether a specific given watermarking signal is present in an image, a procedure we call watermark *detection*.

In applications in which the purpose of a watermark is to make detectable any change made to the image, the watermark must be extracted in its exact form. The watermark extraction procedure is represented in the middle portion of Figure 1. As in the watermark insertion procedure, we specify a user key to extract the proper watermark. It is through the integrity of the extracted watermark that we can tell whether the watermarked image has been changed or manipulated. An implication is that we can embed a message, such as a serial number, within the watermark, and we can extract the message from the watermarked image. In other words, we can ask a question like, What is the serial number of this

IT IS THROUGH THE INTEGRITY OF THE EXTRACTED WATERMARK THAT WE CAN TELL WHETHER THE WATERMARKED IMAGE HAS BEEN CHANGED OR MANIPULATED.

image? The watermark extraction algorithm gives us the answer.

With robust watermarks, it is often not feasible to extract the exact watermark signal, so we use a watermark detection procedure represented by the bottom portion of Figure 1. The correlation test, a popular approach for detecting robust watermarks suggested in [9], is the basic detection mechanism for many robust watermarking techniques. We specify an ID (a watermark) and check whether there is sufficient correlation between the ID and the image, as in Figure 2. We conclude that a valid watermark is detected if the correlation is strong enough. However, several comments are in order. First, note the subtle difference between the extraction process for fragile watermarks and the detection process for robust watermarks. To start, an ID for the detection process has to be specified. Specifying an ID implies that in robust watermark detection we ask, Does the image contain the specified ID? instead of the stronger question asked in fragile watermarking detection, What is the ID of the image? Second, a threshold has to be specified in the detection process. Specifying a threshold implies that in the watermark detection procedure, we are asking, Is this image *likely* to contain the specified ID? The answer—the likelihood—depends on the threshold.

Watermarking Techniques

Each type of application imposes special requirements on the watermarking technique. Here we describe different types of proposed watermarks, classifying them from the point of view of the typical application for which they could be used. (For a survey of watermarking techniques, see [8].)

Robust watermarks. Robustness is a key requirement often imposed by an application. That is, the watermark embedded in the data must be recoverable despite intentional or unintentional modification of the image—a requirement generally very difficult to meet. For example, the watermark should be robust against such image-processing operations as filtering, requantization, dithering, scaling, cropping, and various common image compression techniques (see Figure 3).

Early robust watermarking techniques typically manipulated an image in the spatial domain. One simple spatial-domain watermarking technique [5], shown to be quite robust against lossy image compression, filtering, and scanning, consists of dividing the image pixels into two roughly equal sets—A and B—randomly selected by means of a secret key. A small integer k is then added to the intensity value of each pixel in set A and subtracted from the intensity of each pixel in set B. In practice, k is small enough that its addition or subtraction does not cause any visible degradation in the image. Detection consists of partitioning the given image pixels using the known division strategy and computing the mean intensities of the pixels in A and B separately. If the image is watermarked, the difference between these intensities is approximately $2k$; otherwise the difference is close to 0.

As opposed to spatial-domain techniques, which can embed a small number of bits, transform-domain techniques can embed a larger number of bits without incurring noticeable visual artifacts. Such techniques can be employed with common image transforms, such as discrete cosine transforms (DCTs), the wavelet transform, and Fourier transforms. A transform-domain-based technique, reported in [12], is tailored to JPEG lossy image compression, facilitating insertion of a watermark while an image is being compressed. The watermark is embedded in the DCT coefficients obtained by transforming nonoverlapping 8×8 image blocks. The specific blocks are pseudorandomly selected, and specific coefficients from a limited set are then made to conform to one of two relationships in order to embed one bit of information. If storing that bit requires a significant change in the coefficients of a block, the coefficients are manipulated to form an invalid pattern, indicating that no information is contained in the block. The invalid pattern generally requires a smaller change in the coefficients than that needed to encode a 1 bit or a 0 bit. During the extraction process, the same coefficients are pseudorandomly selected, and the relationship among the three coefficients is analyzed to extract the one bit of information.

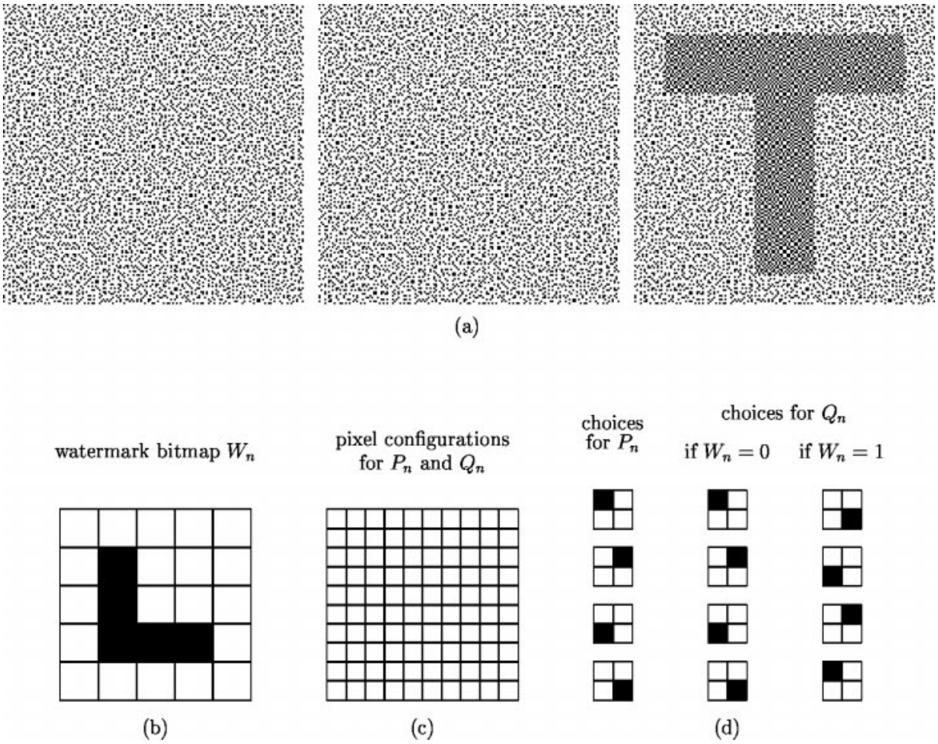

Figure 4. Visual cryptography. The first two bitmaps in (a) contain a message (the watermark) that cannot be read independently. When they are superimposed on each other (by way of a pixel-by-pixel logical AND operation), the secret message appears, as shown in the farthest-right bitmap of (a). Parts (b), (c), and (d) show a method for embedding a binary watermark in a two-out-of-two scheme in visual cryptography. Corresponding to each pixel location of the bitmap (b), we randomly pick one subpixel configuration from the four possible arrangements for Pn; we then pick the subpixel configurations for Qn according to the choice for Pn and the value of the bitmap (b). Note: When Wn = 0, then Pn and Qn are identical, giving an average intensity equal to 1/4 of the intensity of white. When Wn = 1, the black subpixels in Pn and Qn are at opposite corners, or out of phase, giving an average intensity equal to 1/2 of the intensity of white.

A significant advantage of DCT-based watermarking techniques over spatial-domain approaches is that watermarking insertion and detection can be done with the compressed representation of the image. Using the compressed representation makes it possible to directly insert watermarks into JPEG images, as well as into MPEG and MJPEG video streams. Watermark detection or extraction can also be done directly from the compressed data with minimal decompression. These watermarking techniques are suitable for such applications as digital TV broadcasting and video on demand. However, the problem with a watermarking technique that modifies quantized DCT coefficients is that changes in the coefficients are amplified when the image is decompressed. As a result, artifacts caused by the insertion of a watermark are common, and various adaptive techniques that try to minimize the artifacts have been reported in the literature. (For a survey of such techniques for image and video data, see [8].)

Although these techniques are reasonably robust in some simple image-processing operations, Cox et al. [1] observed that in order for a watermark to be robust it has to be embedded in the perceptually most significant parts of an image. Although this embedding scheme seems to contradict the requirement of invisibility, a reasonable balance between the two requirements can be obtained through principles from spread-spectrum communications. The scheme [1] embeds a set of independent and identically distributed samples *S* drawn from a Gaussian distribution into the data's most perceptually significant frequency components. Results reported by Cox et al. using the largest 1,000 DCT coefficients for watermark insertion show the technique to be remarkably robust against various image processing operations, as well as printing and rescanning. However, it has also been shown that Cox et al.'s technique is susceptible to collusion attacks. That is, users with copies of the same image with different watermarks can collude to produce a copy that would not show the presence of any of these watermarks. It has been conjectured in the literature that 8–10 copies may be sufficient for a collusion attack.

In addition to these techniques, many others have been developed in the past few years, and new ones continue to appear in the image-processing literature. However, these techniques indicate that robust watermarking is indeed a difficult problem. A single technique satisfying all requirements imposed on robust watermarking is quite difficult

NO KNOWN TECHNIQUE CAN SURVIVE A RESOURCEFUL AND CLEVER ATTACKER.

to achieve and is the subject of active research.

Authentication and verification watermarks. The objective of this type of watermark is to enable detection of any change made to an image. Furthermore, it is desirable to be able to spatially localize the changes in an image. For example, Friedman [2] proposed the idea of a "trustworthy digital camera," in which a digital signature scheme is incorporated to protect the images from unauthorized manipulation. His scheme uses a standard digital signature for general digital data. A signature for each image is computed in the same way a typical digital signature is computed; the digital signature is then stored together with the image data. Hence, a user can check the integrity of the image data in the usual way with the digital signature.

Yeung and Mintzer [11] recently suggested a way to incorporate a binary watermark into an image so any change to the image can be detected and localized. Their watermarking method is based on an extraction function in which a binary output is generated by using each input pixel value as an index to a fixed binary random sequence *R*. If the original image is colored, then the binary output is formed by combining the results from the three color planes with an exclusive-OR function. If the binary output is not the same as the corresponding value in the binary watermark, a small adjustment is made to the original pixel value so the two binary values are identical. Hence, the adjusted image is the watermarked image from which a binary watermark can be extracted using the random sequence *R*.

Wong [10] proposed a method using a cryptographic hash function, such as MD5, which is a standard way of computing a digest from a digital data stream. The image is partitioned into blocks (such as 8 × 8). The LSB of each pixel in a block is stripped off, and the remaining high-order bits, or most significant bits (MSBs), of the pixels are concatenated with the image-size parameter and a secret user key. The resulting bitstream is passed into a cryptographic hash function to generate a digital signature that is logically combined with a binary watermark image; the result is then inserted back into the image as the LSB. This scheme yields a watermarked image and allows detection of changes to any pixel up to the block level. Furthermore, the scheme allows for detection of any change to the image size, such as scaling or cropping.

More recently, Wong's watermarking algorithm was extended to a public-key scheme in which a user's private key is needed to insert the watermark, as in Figure 3. However, watermark extraction requires only the public key. The LSB of each pixel in a block is stripped off, and the remaining MSBs of the pixels and the image size parameters are hashed and the result encrypted using a public-key algorithm. The resulting cipher text and the binary watermark image are combined using an exclusive-OR function; the result is then embedded into the LSB of the block. In the extraction step, the same MSB data and the image size parameters are hashed. The LSB of the data block (cipher text) is decrypted using a corresponding public key decryption algorithm. The decrypted result and the hash output are combined using an exclusive-OR function to produce the visual watermark. The public-key extension certainly expands the practical applicability of this watermarking scheme.

Watermarks for binary images. The previous sections focused on digital watermarking for continuous-tone images. Here we consider embedding binary watermarks into binary or halftone images. This idea, originated under the name *visual cryptography*, was later applied to information hiding in binary images [7] (see Figure 4). A secret message (or watermark) is hidden in the two binary images on the left side of Figure 4(a). The secret message is not apparent to a human observer looking at either one of the bitmaps. However, when the two images are superimposed on each other (a process essentially equivalent to a pixel-by-pixel logical AND operation), the secret message is uncovered, as shown in the farthest-right bitmap of Figure 4(a). This method can be generalized into a *k*-out-of-*n* scheme (where, naturally, $k \leq n$), as in cryptography. We have a message embedded into *n* different copies of different binary images. If we want to uncover the secret message, we have to obtain *k*

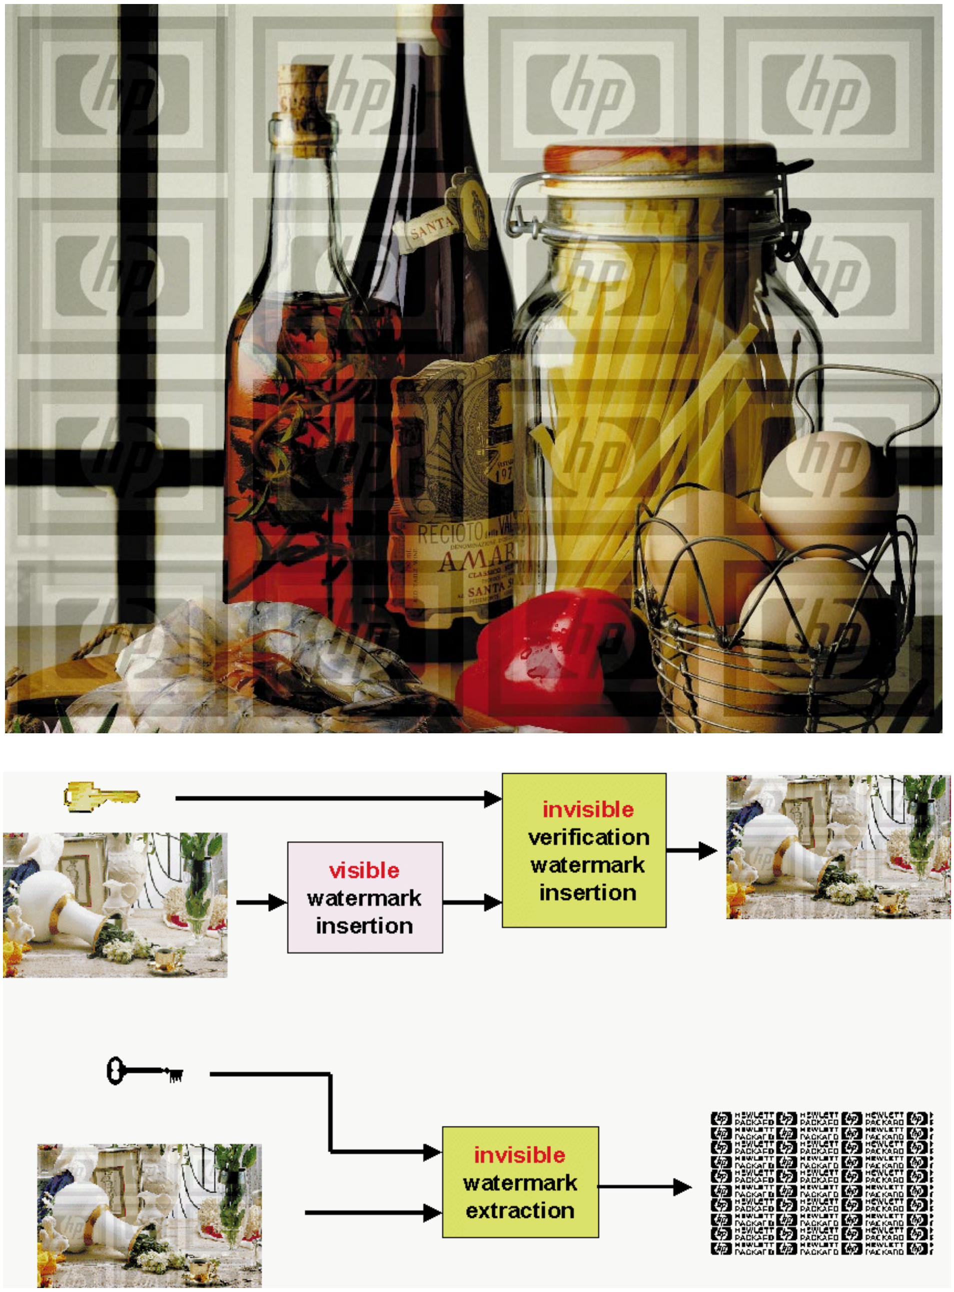

Figure 5. Visible watermarking. (a) An image watermarked using a secure visible watermarking algorithm [10]. (b) A block diagram of the algorithm, using an invisible authentication watermark to protect a visible watermark. The authentication watermark prevents a user Bob from inserting Alice's logo into an image and then claiming the visible watermark was inserted by Alice.

different copies of these binary images and superimpose them onto each other.

Consider a method for the special case of a two-out-of-two scheme. Here we first obtain a bitmap W_n of a message, such as the letter "L" in Figure 4(b). We want to create two binary images—P_n and Q_n—so each appears as random noise but form the letter L when they are superimposed on each other. First, we let P_n and Q_n be binary images, each containing four times the number of pixels as are in W_n; but each pixel is 25% the size of the pixels in W_n, as shown in Figure 4(c). To generate image P_n, we randomly select each group of four subpixels from the four possible configurations in the column for P_n in Figure 4(d). Hence the binary image P_n appears like random noise with an average intensity of 1/4 of the intensity of white because one of four subpixels is black and the others are white. (See the farthest-left bitmap in Figure 4(a) for an example.) Corresponding to each group of four subpixels in P_n, we select from Figure 4(d) a group of four subpixels for Q_n corresponding to the choice made earlier for P_n and the value (1 or 0) of W_n. Hence Q_n also appears as random noise with an average intensity of 1/4 of the intensity of white. However, when P_n and Q_n are superimposed on each other, a well-defined visual pattern (the watermark) appears.

The reason for the appearance of the underlying watermark is that when $W_n = 0$, the subpixels in P_n and Q_n are identical and still give an average intensity of 1/4 of the intensity of white when superimposed. When $W_n = 1$, however, the black pixels in P_n and Q_n are "out of phase" with each other. In this case, the human observer sees the union of the black pixels, giving an average intensity of 1/2 of the intensity of white. This method produces a visual watermark visually unrelated to the dots on each individual sheet P_n and Q_n but is related to the union of the dots in P_n and Q_n. (See the farthest-right bitmap in Figure 4(a) for an example.)

Knox and Wang [3] proposed a similar method for embedding a binary watermark into printed halftone images, based on a philosophy similar to that of Shamir's method [7]. The difference is that Knox and Wang used a stochastic halftoning screen and embedded the structure of the watermark into two dither matrices (arrays of threshold values used to generate halftones). Using each of these dither matrices to generate a halftone image and then superimposing the two printed images reveals the embedded watermark.

Visible watermarks. Visible watermarking is the insertion of a visible pattern or image into a source image (see Figure 5). There are two important criteria for a good visible watermark. First, it must be difficult for an unauthorized person to remove it. To this end, Magerlein, Braudaway, and Mintzer [4] suggested a visible watermarking technique whereby random noise is inserted to increase the difficulty of unauthorized removal. A good visible watermark also has to resist falsification. Since it is relatively easy to embed a pattern or logo into an image, we have to ensure the visible watermark in an image was indeed inserted by the claimed user. For example, Bob can insert Alice's logo in an image and claim the logo was inserted by Alice and that the image is an authorized copy. Alice would like to be able to deny such an act if Bob was the one who actually inserted the watermark. To this end, Wong [10] proposed protecting a visible watermark using an authentication-type invisible watermark (see Figure 5). Using this method, Alice can deny any association with the image containing her logo, provided the image does not contain her invisible authentication watermark. Since only Alice would have her secret key for the insertion of the authentication watermark, Bob cannot falsify an image simply by inserting Alice's logo.

Conclusion

Although robust watermarking techniques for copyright protection are an important area for future research, robust digital watermarking is a very difficult problem due to the numerous kinds of image manipulations a robust watermark has to be able to survive. Even specifying the set of image manipulations a robust image watermarking technique can *provably* survive is already nontrivial. Although many techniques have been proposed, there is still no known technique that can survive a resourceful and clever attacker. Software publicly available on the Internet has been shown to be effective in breaking several commercial watermarking techniques.

Solutions for both public- and secret-key verification watermarks have been proposed, allowing detection of any change in an image. Since images captured from a digital camera are often encoded using JPEG lossy compression, a verification watermark able to survive JPEG compression is highly desirable. However, this survivability requirement makes the problem of verification watermarking somewhat ill-defined, in the sense that it is very difficult to characterize precisely the way JPEG changes an image as opposed to the way other image-manipulation algorithms change an image.

Watermarking is undoubtedly important for protecting various forms of content in the digital age. Although the technology is still a relatively young research area, it has already attracted many first-rate researchers. We will see many new watermarking approaches and applications in the near future. C

References

1. Cox, I., Kilian, J., Leighton, T., and Shamoon, T. Secure spread spectrum watermarking for multimedia. Tech. Rep. 95-10, NEC Research Institute, Princeton, N.J., 1995.
2. Friedman, G. The trustworthy digital camera: Restoring credibility to the photographic image. *IEEE Trans. Consum. Electron. 39,* 4 (Nov. 1993), 905–910.
3. Knox, K., and Wang, S. Digital watermarks using stochastic screens. In *Proceedings of SPIE/IS&T Symposium on Electronic Imaging* (San Jose, Calif., Jan.), SPIE, Bellingham, Wa., 1997.
4. Magerlein, K., Braudaway, G., and Mintzer, F. Protecting publicly available images with a visible image watermark. In *Proceedings of Optical Security and Counterfeit Deterrence Techniques* (San Jose, Calif., Feb.). SPIE, Bellingham, Wa., 1996, pp. 126–132.
5. Nikolaidis, N., and Pitas, I. Copyright protection of images using robust digital signatures. In *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing* (Atlanta, Ga., May). IEEE Press, 1996, pp. 2,168–2,171.
6. Sarwate, D., and Pursley, M. Cross-correlation properties of pseudorandom and related sequences. *Proc. IEEE 68* (May 1980), 593–619.
7. Shamir, A. Method and apparatus for protecting visual information with printed cryptographic watermarks. U.S. patent 5488664, Jan. 1996.
8. Swanson, M., Kobayashi, M., and Tewfik, A. Multimedia data embedding and watermarking technologies. *IEEE Proc. 86,* 6 (June 1998) 1,064-1,087.
9. Van Schyndel, R., Tirkel, A., and Osborne, C. A digital watermark. In *Proceedings of ICIP* (Austin, Tex., Nov.). IEEE Press, 1994, pp. 86–90.
10. Wong, P. A watermark for image integrity and ownership verification. In *Proceedings of IS&T PIC Conference* (Portland, Oreg., May). 1998.
11. Yeung, M., and Mintzer, F. An invisible watermarking technique for image verification. In *Proceedings of ICIP* (Santa Barbara, Calif., Oct.). IEEE Press, 1997.
12. Zhao, J., and Koch, E. Embedding robust labels into images for copyright protection. In *Proceedings of the KnowRight'95 Conference on Intellectual Property Rights and New Technologies* (Vienna, Austria, Aug.), 1995, pp. 241–251.

Nasir Memon (memon@max.cs.niu.edu) is an associate professor of computer science at Northern Illinois University, currently on leave as visiting faculty at Hewlett-Packard Laboratories in Palo Alto, Calif.

Ping Wah Wong (pwong@cup.hp.com) is a manager in the Internet Imaging Operation at Hewlett-Packard in Cupertino, Calif.