DIGITAL WATERMARKS: SHEDDING LIGHT ON THE INVISIBLE

DIGITAL MEDIA PROMISE REWARDS—AND RISK, CHIEFLY IN THE POTENTIAL FOR UNAUTHORIZED REPLICATION AND DISTRIBUTION. IN THE CONTENT PROTECTION DEBATE, DIGITAL WATERMARKS OFFER ONE SOLUTION.

Minerva M. Yeung
Boon-Lock Yeo
Matthew Holliman
Intel Corporation

We are witnessing the increasing popularity and utilization of digital media (digital images, digital audio, and digital video clips). Not surprisingly, this is a mixed blessing. The computer and consumer electronics industries—and, by extension, the semiconductor industry—anticipate increased revenues from digital media. Higher revenues have been predicted for certain segments, including digital cameras, camcorders, digital versatile disk (DVD) players and writers, video CD devices, and media processors. On the other hand, digital media renew the media industry's worries that existing laws and protections for digital content cannot sufficiently protect the industry's interests.

In particular, there are inexpensive and readily available tools and equipment that can be used to replicate, manipulate, and distribute digital multimedia content with ease. Complicating the issue is the fact that the replication often produces perfect copies of the original. Moreover, manipulation and editing of the digital copies can deceive even the most professional eyes, and mass distribution in electronic forms can take place in a matter of seconds. Misappropriation of digital assets greatly concerns content creators and owners. This is particularly true when the content is available through the Internet, or on any programmable device like the PC.

For content creators and owners, protecting digital assets is vital. There is a consensus that more effective tools are needed to deter—better yet, prevent—misappropriation of digital media, from photographs and computer-generated graphics to CD audio and DVD video.

Interestingly, this need carries profound implications for the computer and consumer electronics industries in the pursuit of a digital and multimedia world. For example, content is the driving force behind DVD sales, which will in turn drive the sales of DVD players and PC upgrades. Unfortunately, many DVD titles available today are not first-tier hit movies. The movie "Titanic" is still unavailable in DVD format, even though the videotape version hit the stores some months ago. Without assurance of proper protection against lost revenues, content owners are reluctant to make available and promote digital media.

In addition to content protection, there are other questions to be addressed. What new features can be made available in digital devices? What features can be included to offset potential drawbacks by going digital? For example, how do we prove the authenticity of digital photos? Can we admit digital images as evidence in court, or trust news report contents, given that such images can easily be manipulated with inexpensive tools?

0272-1732/98/$10.00 © 1998 IEEE

Digital watermarking can answer some of the concerns regarding copyright protection of high-value digital materials such as CD-quality audio, publication-quality images, and digital video. Watermarks can offer value-added features for encoding user data and ownership information in the content capture and creation process. In this article we examine digital watermarking applications and limitations. We also describe the tests we conducted, and consider the implications of watermarking for digital imaging and media devices.

The watermarks are bound to and hidden in the source data (or object), inseparable from the source. The watermarks can thus survive operations that do not degrade the data.

Digital watermarks

Digital watermarking is the embedding of unobtrusive marks or labels (the watermarks), which can be represented as bits, into digital content. Usually, the embedded marks are invisible (or imperceptible), and can later be detected or extracted. The watermarks are bound to and hidden in the source data (or object), inseparable from the source. The watermarks can thus survive operations that do not degrade the data beyond the utility value for the intended applications. The object can be an image, an audio clip, a video clip, or a 3D model. (The terms "visible" and "invisible" are commonly used to describe the effects of watermarking multimedia objects, although they are more appropriate for visual media. For audio clips, we use the terms "perceptible" and "imperceptible".)

Some watermarks, designed to be vaguely visible, can serve as copyright notification. Invisible watermarks, which are our focus in this article, that are embedded into an object can serve the same purpose if they can be detected readily and displayed. Based on their properties and application domains, watermarking techniques can be further classified as fragile or robust.

In fragile watermarking, the embedded watermark changes or disappears if a watermarked object is altered. The watermark can thus be used to verify content to ensure its integrity. For example, trustworthy images, captured with a digital camera, can be provided for use with news articles or presented as evidence in a court of law. In this case, a content creator embeds an invisible watermark at capture or creation time. The watermark's presence at the time of publication or upon receipt by an end user is intended to indicate that the image or object has not been altered.

In another application of fragile watermarking, digital photos, 3D models, or human fingerprint images can be scanned, watermarked, and then stored in a digital archive. The watermarks let the content owner detect unauthorized alterations without having to compare the objects to the original scans or send separate signature files for authentication via digital signature schemes.

In robust watermarking, the embedded watermark persists even after attempted removal. Such attempts might be accomplished by common transformations such as filtering, cropping, translation, rotation, resizing, or lossy compression. Alternatively, removal of watermarks might be attempted by malicious attacks such as the use of sophisticated algorithms to process the objects. Intentional attacks can include any combinations of such transformations.

Prominent applications of robust watermarks are summarized here:

- *Evidence of ownership*—Robust watermarks can indicate ownership when the owner's label is detected from a suspected copy, provided the watermarking scheme is properly designed.
- *Fingerprinting*—Imprinting fingerprints into data as watermarks allows the intended recipient to be traced should the content be misappropriated. An object's seller, for instance, might insert a unique and invisible label into the object to indicate to whom the object is sold. Later, the seller may find that a copy of

the object has been published without royalty payment. The extracted watermark can serve to identify the misappropriator and act as a deterrent.

- *Tracing and infringement detection*—Embedded watermarks can serve as a form of indicators—the detection of these indicators can trigger protection or royalty collection mechanisms. Some commercial systems (for example, the one by Digimarc[1]) advertise a Web-crawling capability to find copyrighted images by searching for images on the Web to detect watermarks and inform registered owners.
- *Copy control*—Robust watermarks, if coupled with a detection enforcement mechanism, can be used in copy or usage control. The detection of the control data can trigger some particular content protection features in a device. For example, dedicated hardware can allow a device to play video clips, but forbid a user to copy the video if it detects a "No Copy" in the embedded watermark. The Copy Protection Technical Working Group—an ad hoc voluntary group comprised of Hollywood studio, computer, and consumer electronics industry representatives—is working on a potential watermark standard for DVD.
- *Labeling and metadata insertion*—Watermarks, if sufficiently robust, can function as fairly universal and format-independent descriptors. The embedded data can be descriptors of the content and auxiliary information such as time stamps, Global Positioning System data, and object descriptors with links to the creator's Web site. The embedded data, of course, can also be descriptors of copyright and usage control information. The watermark communicates the information as hidden data in the source signal, instead of in file headers, even after format conversions that may include compression.

Digital watermarking does not offer the same capability and level of security as data encryption. Watermarking does not prevent the viewing of (or listening to) content, nor does it prevent content access. Digital watermarking schemes are not immune to hackers' attacks. In addition, different watermarking techniques may produce watermarks with different properties that have to be used carefully. Researchers have found some techniques to be *invertible*, which could lead to counterclaims on the rightful ownership of a piece of watermarked content.[2] An invertible watermarking technique is one that is susceptible to an attack that creates multiple claims of ownership for the same watermarked content.

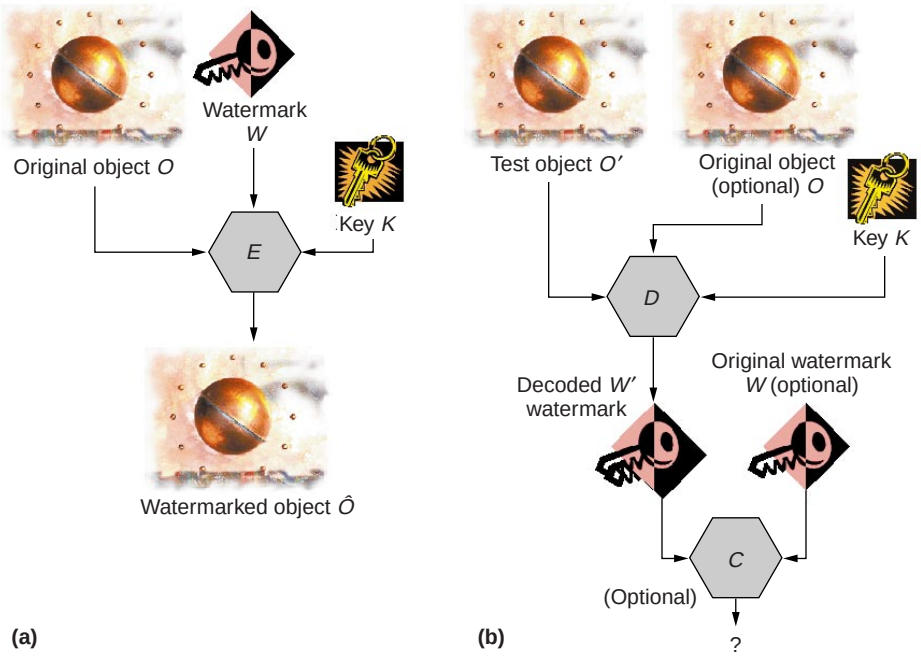

Figure 1. Encoding, decoding, and comparing embedded watermarks in an object. In the encoding process (a), a content creator/owner inserts a watermark into an original object. In (b), a content owner checks a test object to recover a watermark, then compares the recovered watermark to the original inserted watermark.

Watermarking techniques

A watermarking technique operating on a piece of digital media consists of a way to encode the watermark into the piece of content and, subsequently, a way to decode that

watermark. We describe a generalized formulation for encoding and decoding watermarks. Figure 1a illustrates the encoding process by which a watermark is inserted into an object. Figure 1b shows the decoding process by which a watermark is recovered and compared to the inserted watermark. The object can be an image, an audio or video clip, or a 3D model.

We denote an object by O, a watermark (comprised of a sequence of owner-supplied data bits $W = \{w_1, w_2, \cdots\}$), and a watermarked object by $\hat{O}$. We also define the key K as a sequence of bits that helps to define the specific mapping function for additional security in watermark insertion and extraction. Additionally, ε is an encoder function. This function takes the object O and a watermark W, incorporates the mapping supplied by K and generates a new object called the watermarked object $\hat{O}$:

$$\varepsilon_K(O, W) = \hat{O} \tag{1}$$

W can also be made to be dependent on O.

The decoder function D takes an object O' (O' is any object, watermarked or otherwise) and recovers a watermark W', or evidence of the original watermark's presence, from the object. If available, the decoding key K can help define the decoding function's specific mapping. In this process, the watermark decoder can use the original object O as a reference object for the extraction process.

Using O in watermark decoding typically provides extra robustness against intentional and unintentional corruption of an image's pixel values. More formally, if the decoding scheme involves a reference O, we write

$$D_K(O', O) = P(W) \tag{2}$$

Here P is a function indicating the presence of the watermark W in O'. We call this a private watermarking scheme because the decoding requires the use of the original O. When $P(W)$ equals W', the decoding process simply returns the extracted watermark W'. P may also take the form $P()$ equals $Evid()$ that returns a scalar value indicating the evidence of the presence of W in O'.

If the decoding does not need O in the watermark extraction, we can write a general decoding function as

> Robust watermarking is the insertion of a watermark such that the watermark of its presence can be reliably detected even after image modifications, which is crucial for content protection.

$$D_K(O') = P(W) \tag{3}$$

This is a public watermarking scheme because the decoding does not require the use of the original O. Note that in some watermarking schemes, D may be dependent on the specific watermark embedded in the encoding process.

When $P(W)$ equals W', a comparator function can further compare W' to the reference W, which generates a binary output decision indicating a match or otherwise.

The encoding and decoding functions are designed a little differently, depending on whether the watermarking is fragile or robust. In robust watermarking, given a test object O' derived from a watermarked object $\hat{O}$ ($O' = T(\hat{O})$ for some transformation T), the decoder should report strong indications that W is present. In fragile watermarking, on the other hand, given a test object O' differing from a watermarked object $\hat{O}$, the decoder should report mismatches even for small changes detected.

Fragile watermarks are intended for verification applications. The advantage here is the capability of verifying an object without resorting to comparing it with a reference original, as the watermark is directly detected from the watermarked object. In robust watermarking, some schemes require a reference original to robustly detect or extract the watermark, while others do not need the original.

Robust watermarking techniques

Let's say we denote a feature set F (a set of

values derived from the pixel values) of an image as $\{p_i\}$. Robust watermarking, then, refers to the insertion of a watermark $\{w_i\}$ into the feature set F such that the watermark or its presence can be reliably detected even after image modifications. We can denote the resulting watermarked image to have a new feature set $\{\hat{p}_i\}$, in which $\hat{p}_i = p_i + w_i$. Both p_i and w_i may be vector-valued, and w_i may be dependent on p_i.

Most robust watermarking techniques reported in recent years can be broadly classified as being either spatial domain or transform domain based. Spatial-domain-based techniques use the set of pixels as the feature set F for watermark encoding. Transform-domain-based techniques use, as the feature set, coefficients from transforms such as the Discrete Cosine Transform, wavelets, and Fourier transforms. Very recently, there has also been work on geometric watermarking schemes, in which F captures some geometric properties of the images.

A number of schemes of varying robustness have reportedly withstood common image processing tasks without incurring significant quality degradation. Many of the techniques share common properties or characteristics; some more extensively tested than others. Here we describe the basic techniques used for spatial and transform domain watermarking schemes, highlighting the key differences. We do not discuss geometric watermarking schemes, since they are still in the early stages of research.

Spatial domain watermarking. One class of techniques known as Patchwork-like are representative of this approach. Patchwork, first reported by Bender and colleagues, illustrates the basic principles of numerous schemes.[3] The algorithm pseudorandomly selects n patches of points, (a_i, b_i), where a private key generates the pseudorandom sequence.

Denoting the sets of a_i's as A and b_i's as B, the pixel values in A are incremented by k, while values in B are decremented by k. In this case, the feature set is the subset of pixels $A \cup B$. The net effect of this watermarking scheme is to assign w_i to be k or $-k$ depending on whether the pixel p_i in $A \cup B$ belongs to set A or B. Detection is accomplished by computing

$$S = \frac{\sum_i (a_i - b_i)}{n}$$

and comparing S to a threshold value. A value of S exceeding the threshold is deemed to indicate the presence of a watermark, whereas an unwatermarked image should result in a difference $S \approx 0$. Note that the decoding does not extract a watermark; rather, it computes the evidence of the original watermark's presence. Similar approaches have been reported, though the detection methods can vary.[4,5]

The most basic implementation of the algorithm chooses patches to be single pixels, which essentially adds high-frequency noise to the image. Such watermarks are vulnerable to filtering and lossy compression. In practice, larger clusters of points could be used. For example, we can treat p_i as a group of pixels (say, 4×4 blocks) and w_i as a vector. This effectively shifts the watermark's frequency content toward lower spatial frequencies, which in our experiments generally increases watermark robustness albeit at the expense of some loss in detection reliability. In addition, more features can potentially be built into the detection methods to detect the presence of the watermark under distortions such as cropping, scaling, and stretching.

Such a class of detection techniques are public schemes and do not require the use of the original, making them more suitable for applications such as tracing and copy control. But if the original is available, the decoder can use it to assist in detection. Use of the original in decoding results in a private scheme and potentially allows the watermark to be recovered, rather than its presence detected. Private schemes tend to be suitable for establishing evidence of ownership. To test an image J with features $\{p'_i\}$, we could compute

$$d_i = p'_i - p_i \tag{4}$$

If $p'_i = \hat{p}_i$, then $d_i = w_i$. However, $\hat{p}_i$'s would be very sensitive if p_i's are chosen to be single pixels, and the resulting d_i may not be reflective of w_i.

On the other hand, if the p_i's are transform domain coefficients, then $\hat{p}_i$'s would be much less sensitive, especially if p_i's represent low-frequency components. This is the basis for

many transform domain watermarking approaches.

Transform domain watermarking. One of the earliest examples of transform domain watermarking techniques that use Equation 4 for decoding was reported by Cox and colleagues.[6] This technique first computes a full-image frequency transform, followed by the watermark's insertion into several perceptually significant transform coefficients. The watermark itself is a set of independent, identically distributed samples drawn from a Gaussian distribution, generated as a pseudorandom sequence from a user's private key. The decoder extracts a watermark by first comparing a suspected watermarked image with an unwatermarked original using Equation 4. Second, the decoder correlates the extracted, possibly degraded, watermark sequence with the original uncorrupted watermark.

The transform domain approach is also robust during a variety of tests but is extremely computationally demanding. Researchers have proposed block-based variations, which demand significantly less computation. In addition, watermark encoders can exploit local masking effects to shape the watermark signal and to better ensure invisibility.[7]

We can also employ public spatial-domain techniques, such as Patchwork, on frequency coefficients. The difference here is that p_i is no longer treated as pixels, but rather as frequency coefficients.

Basic principle. Both spatial and transform domain watermarking techniques essentially add a watermark signal w_i to appropriately chosen feature sets such that w_i is preserved after image modification. Whether the original is required or not during decoding then depends on a combination of factors, such as the application, degree of robustness, and complexity. For a description of test results we achieved in testing watermarking techniques, see the box "Robustness test of watermarking techniques" on the next page.

Fragile watermarking techniques

The technique proposed by Friedman for trustworthy digital cameras is closely related to fragile watermarking.[8] A key difference is that content creators don't insert a watermark; rather, they compute and store a digital signature for the image. A user can check image integrity on the basis of the digital signature but cannot locate any altered regions.

> We can offer stronger protection if watermarking is incorporated inito the image capture pipeline, which would serve to verify content integrity and image authenticity.

In earlier works on fragile watermarking, the watermark was embedded by changing the least significant bits (LSB) of an image.[9] Subsequent techniques to improve LSB-based fragile watermarking have been proposed, for example, such as embedding a binary watermark into a source image that enables detection of subsequent unauthorized image alterations.[10] During this insertion procedure, a watermark extraction function is applied to each input pixel in turn. The unwatermarked pixel value is used as an index into a pseudorandom binary sequence. If the extracted bit value matches the corresponding value in the watermark plane, processing continues with the next pixel. Otherwise, the scheme adjusts the current pixel value until the extracted watermark value equals the desired bit value.

This process is repeated for each pixel in the image. Watermark extraction applies the same function to each pixel in turn, generating a binary image that can be used to identify and localize changes made to individual pixels.

A different fragile watermarking technique partitions an image into 8×8 blocks and strips the LSB of each pixel in a block.[11] The technique concatenates the remaining high-order bits of the pixels with the image size parameters, and hashes them using a cryptographic hash function such as MD5. The resulting ciphertext is then encrypted with a public key, XOR'ed with a binary watermark image, and inserted back into the LSB of the block. Watermark detection is accomplished

Robustness test of watermarking techniques

We conducted several experiments to test robust digital watermarking on two test images. Specifically, we conducted robustness testing of the scheme developed by the NEC Research Institute[1] and compared it to our version of Enhanced Patchwork.[2]

Tables A and B show sample test results we obtained for the two images, "Peppers" and "Town." In both cases, we achieved these results on the basis of the NEC scheme and our version of Enhanced Patchwork, implemented with 5×5 overlapping patches. The unwatermarked Peppers image is shown in Figure A, while the same image after watermarking with Patchwork is shown in Figure B for comparison.

Both watermarking schemes show robustness after a variety of attacks. The watermark can be detected and watermarked image differentiated after JPEG compression (even after 68:1 compression), downscaling, cropping, printing and scanning, and even after photocopying, the effects of which are seen in Figure C. The NEC technique is sensitive to alignment and scaling problems, and so requires careful registration with the original image for accurate watermark detection.

In contrast, Enhanced Patchwork requires only reasonably close (with-

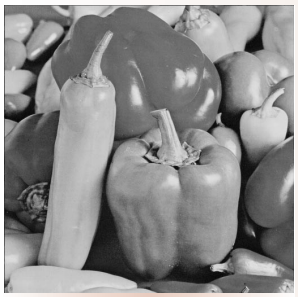

Figure A. Original image of peppers (512×512).

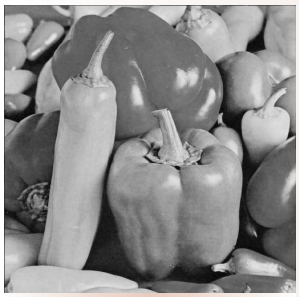

Figure B. Pepper image watermarked by the Enhanced Patchwork algorithm.

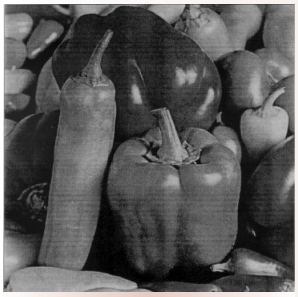

Figure C. Scanned version of Figure B after printing and photocopying.

Table A. "Peppers" image (512 x 512).

	NEC scheme		Enhanced Patchwork	
Image state	**Correlation**	**PSNR (dB)**	**Correlation**	**PSNR (dB)**
Unwatermarked image	0.0107	N/A	0.0156	N/A
Uncorrupted watermarked image	1.0000	33.25	1.0000	42.55
JPEG compression (15:1)	0.8398	29.72	0.8052	32.52
JPEG compression (26:1)	0.7986	28.78	0.6633	30.90
JPEG compression (68:1)	0.5316	25.42	0.4864	26.11
Downscaling (256×256)	0.2413	21.29	0.4255	21.57
Cropping (256×256)	0.4023	33.25	0.3647	42.55
Printing (HP Laserjet IIIsi) and scanning	0.2147	19.04	0.5558	19.00
Printing (QMS 3825) and scanning	0.1362	19.00	0.5089	21.13
Printing, photocopying, and scanning	0.1518	12.54	0.6907	11.81
Stirmark (default parameters)	0.1158	18.31	0.2187	18.42

Table B. "Town" image (640 x 480). This was the second image we tested, not shown here except for the results.

	NEC scheme		Enhanced Patchwork	
Image state	**Correlation**	**PSNR (dB)**	**Correlation**	**PSNR (dB)**
Unwatermarked image	–0.0204	N/A	–0.0131	N/A
Uncorrupted watermarked image	1.0000	35.83	1.0000	46.01
JPEG compression (15:1)	0.9264	25.52	0.7208	25.90
JPEG compression (26:1)	0.7344	23.25	0.5924	23.48
JPEG compression (68:1)	0.5628	22.27	0.5423	22.44
Downscaling (320×240)	0.1776	15.95	0.7481	15.98
Cropping (320×240)	0.5134	35.83	0.2489	46.01
Printing (HP Laserjet IIIsi) and scanning	0.2216	17.94	0.7739	19.14
Printing (QMS 3825) and scanning	0.2066	16.94	0.7587	16.28
Printing, photocopying, and scanning	0.1352	12.69	0.5303	13.53
Stirmark (default parameters)	0.1520	14.64	0.5070	14.67

The scanner we used for our experiments of Tables A and B was an HP ScanJet IIcx, operating at 72 dpi. Downscaling was accomplished by bilinear resampling. We upsampled the downscaled image by means of bicubic resampling, prior to watermark detection. We accomplished detection of the watermarks following cropping, by replacing missing portions of the cropped images with the corresponding portions of the original unwatermarked images.

in two or three pixels) alignment with the expected watermark signal and is less sensitive to imperfect registration. This is reflected in the two techniques' relative performance after printing and scanning, although the NEC detection following substantial watermark degradation can be improved with post-processing techniques.

Likewise, the results following an attack by Stirmark[3] indicate the NEC detector's potential vulnerability following subtle image distortions. Stirmark is a tool designed to test watermark robustness by introducing geometric distortions similar to those encountered when printing and scanning.

References

1. I.J. Cox et al., *Secure Spread Spectrum Watermarking for Multimedia*, Tech. Report 95-10, NEC Research Institute, Princeton, N.J., 1995.
2. W.R. Bender, D. Gruhl, and N. Morimoto, "Techniques for Data Hiding," *Proc. SPIE: Storage and Retrieval of Image and Video Databases*, Vol. 2,420, Soc. of Photo-Optical Instrumentation Engineers, Bellingham, Wash., Feb. 1995, pp. 164–173.
3. M. Kuhn, "Stirmark—Image Watermarking Robustness Test, Version 1.0, 1997-11-10," electronic copy available at *http://www.cl.cam.ac.uk/~mgk25/stirmark.html/*.

by decrypting the image LSBs, computing the same hash for each block as above, and XOR'ing the results to generate the binary watermark. This technique can be used to determine whether each 8×8 block has been modified.

Digital imaging devices

Thus far we have discussed watermarking as it is commonly performed with software techniques, on already created digital content. It is possible, however, to design watermarking schemes for digital imaging devices. In fact, we can potentially offer stronger protection if watermarking is incorporated into image capture devices.

Digital cameras and digital scanners are two major types of image capture devices. In both, system designers can incorporate watermarking into the image capture pipeline. The result is that digitized images are watermarked before being placed on or transferred to storage (flash memory on a digital camera or a computer's hard disk).

In a digital camera, we must perform watermarking entirely in hardware on the camera. In a digital scanner, which is typically connected to a computer, we can place the watermarking step either on the scanner hardware or in the software drivers on the computer. The goal is to automatically watermark any digital images captured.

Incorporating fragile watermarking into imaging devices serves to verify content integrity and image authenticity. Subsequent changes to the images could be detected and exposed, thus making forgery of digital images more difficult. Authentication features have already been proposed in new camera designs.[8] Fragile watermarking provides an alternative to these existing proposals.

Fragile watermarks offer a new set of features for digital imaging devices. Similarly, the capability of inserting robust watermarks in a digital imaging device offers the flexibility of incorporating authorship information and annotations of the captured image. For instance, the location at which an image is captured can be embedded into the image using an invisible watermarking technique to embed global positioning system (GPS) information. Such data could be invaluable for subsequent management of image archives—queries such as "find me all images I took at the Grand Canyon" can be handled via the embedded metadata.

Integrating digital watermarking and watermark detection technologies into imaging and media devices is not without controversy. To achieve copy control in DVD players or recorders with embedded watermarks, the definition of "public" watermarking schemes must be further extended beyond detection without an original. In these cases the detector must include the detection of copy control data in all video data whose underlying characteristics can be drastically different.

Narrowly defined, detection of a watermark's presence or absence may not be dictated by a universal threshold: A sufficient

difference in correlation values can indicate a watermark. Broadly defined, detection would require either a preset threshold across a wide spectrum of data characteristics or smart detectors that would adjust thresholds according to data characteristics. Consumers or end users will not tolerate false positives (a "No Copy" watermark is detected when in fact there is none present). Therefore, detection thresholds, if preset in devices, must be set significantly higher than the normal confidence level. This further implies that the robustness of watermark detection decreases significantly in the device because of detection failure in the face of even slight modification. In fact the watermark may still be present in the data, but the detector is just not smart enough to detect it.

Watermark robustness must be studied together with image quality measurements and detector reliability—a major challenge. Robustness concerns and the resilience against attacks have fueled the debate on whether public (in the broad sense) watermarking schemes are suitable candidates for effective copy control in imaging or media devices.

Protecting content against misappropriation or abuse is a key step toward providing a comprehensive information commerce infrastructure. Unless content owners are assured that their rights are protected, and their works are properly compensated and acknowledged, few will be willing to make their content available for others to access and enjoy. Data encryption and scrambling technology can offer secured content delivery, as well as the means to control access and collect revenues. The key to decode or descramble the secured data would be made available only to the content's (paid) patrons.

Unfortunately, little, if any, protection exists for the decrypted or descrambled content, which can be further redistributed or abused. Making matters worse is the fact that perfect digital copies are readily produced as end products. This contrasts with analog devices, which normally introduce quality degradation as a result of duplication.

Many content owners look to innovative technology to combat intellectual property theft, in lieu of or in addition to legislative efforts. Comprehensive content protection goes beyond data encryption technology. Protection must encompass ownership identification; fingerprinting, to allow tracing of source and recipient; audit trails; content branding and labeling; and usage control. In this regard, digital watermarking offers value-added protection on top of data encryption and scrambling for content protection.

Nevertheless, effective watermarking of multimedia objects is a challenge. One reason is that watermarking technology, currently under active research and development in both academia and industry, is cross-disciplinary—spanning the fields of electrical engineering, mathematics, and computer science.

A generic algorithm for digital watermarking in a variety of applications is unlikely. Instead, different classes of applications will impose different, as yet undefined, requirements. Different watermarking techniques are probable, ranging from slight modifications of existing algorithms to completely orthogonal methods in solving the problems. Nevertheless, the quest for effective watermarks, and more generally, the quest for better technologies to protect intellectual property rights, carries profound implications for computer and consumer electronic system design. MICRO

References

1. Digimarc Corporation, http://www.digimarc.com/, Jan 1997.
2. S. Craver, B.L. Yeo, and M.M. Yeung, "Digital Watermarking—Technical Trials and Legal Tribulations," *Comm. ACM*, Vol. 41, No. 7, July 1998, pp. 45–54.
3. W.R. Bender, D. Gruhl, and N. Morimoto, "Techniques for Data Hiding," *Proc. SPIE: Storage and Retrieval of Image and Video Databases*, Vol. 2,420, Soc. of Photo-Optical Instrumentation Engineers, Bellingham, Wash., Feb. 1995, pp. 164–173.
4. I. Pitas, "A Method for Signature Casting on Digital Images," *Proc. IEEE Int'l Conf. Image Processing (ICIP)*, Vol. 3, IEEE Press, Piscataway, N.J., 1996, pp. 215–218.
5. G.W. Braudaway, "Protecting Publicly-Available Images with an Invisible Image Watermark," *Proc. IEEE Int'l Conf. Image Processing (ICIP)*, Vol. 1, 1997, pp. 126–132.
6. I.J. Cox et al., *Secure Spread Spectrum Watermarking for Multimedia*, Tech. Report 95-10, NEC Research Institute, Princeton,

N.J., 1995.

7. M. Swanson, B. Zhu, and A. Tewfik, "Transparent Robust Image Watermarking," *Proc. IEEE Int'l Conf. on Image Processing*, Vol. 3, 1996, pp. 211–214.
8. G.L. Friedman, "The Trustworthy Digital Camera: Restoring Credibility to the Photographic Image," *IEEE Trans. Consumer Electronics*, Vol. 39, No. 4, Nov. 1993, pp. 905–910.
9. S. Walton, "Image Authentication for a Slippery New Age," *Dr. Dobb's J. Software Tools for Professional Programmers*, Vol. 20, No. 4, Apr. 1995.
10. M.M. Yeung and F.C. Mintzer, "An Invisible Watermarking Technique for Image Verification," *Proc. IEEE Int'l Conf. Image Processing 1997*, Vol. 2, Oct. 1997, pp. II.680–683.
11. N. Memon and P.W. Wong, "Protecting Digital Media Content," *Comm. ACM*, Vol. 41, No. 7, July 1998, pp. 35–43.

Minerva M. Yeung is a senior staff researcher and manager of a research group at Intel Corp. in Santa Clara, California, working on media protection and management. Her research interests are in image/video processing, content protection, computer-human interaction, and multimedia information systems. She is an associate editor of *IEEE Transactions on MultiMedia*. She has co-authored more than 30 papers, holds three patents, and has eight pending applications. Yeung received a BSEE degree from Purdue University, and MA and PhD degrees from Princeton University. She is a member of the IEEE.

Boon-Lock Yeo manages the Video Technology Department at Intel Microcomputer Research Labs in Santa Clara, California. He is an associate editor of *IEEE Transactions on Image Processing*. He has published more than 30 technical papers, holds four US patents, and has 13 pending. He received the 1996 IEEE Circuits and Systems Society Video Technology Transactions Best Paper Award. Yeo received a BSEE from Purdue University, and an MA and a PhD from Princeton University. He is a member of the IEEE.

Matthew Holliman is an intern with Intel's Microcomputer Research Labs. His research interests include image and video processing and applications, including digital watermarking and content protection. He received a BS degree in computer science from the University of Illinois at Urbana-Champaign and an MS in computer science from Northern Illinois University.

Direct questions concerning this article to Minerva Yeung at Intel Corp., Microcomputer Research Labs, SC12-303, 2200 Mission College Blvd., Santa Clara, CA 95052; minerva.yeung@intel.com.

For More Information

For more information on digital watermarking, see the July 1998 issue of the Communications of the ACM and the Jan.-Feb. 1999 issue of IEEE Computer Graphics and Applications.