

Digital Watermarking: Algorithms and Applications


*Christine I. Podilchuk
and Edward J. Delp*

Digital watermarking of multimedia content has become a very active research area over the last several years. A general framework for watermark embedding and detection/decoding is presented here along with a review of some of the algorithms for different media types described in the literature. We highlight some of the differences based on application such as copyright protection, authentication, tamper detection, and data hiding as well as differences in technology and system requirements for different media types such as digital images, video, audio and text.

Introduction

The success of the Internet, cost-effective and popular digital recording and storage devices, and the promise of higher bandwidth and quality of service (QoS) for both wired and wireless networks has made it possible to create, replicate, transmit, and distribute digital content in an effortless way. The protection and enforcement of intellectual property rights for digital media has become an important issue. In 1998, Congress passed the Digital Millennium Copyright Act (DMCA) which makes it illegal to circumvent any technological measure that protects an owner's intellectual property rights of digital content. The headline news regarding Napster made the general public aware of the issues regarding intellectual property rights and the impact of current technology.

In recent years, the research community has seen much activity in the area of digital watermarking as an additional



© 1998 CORBIS CORP.

There should be no perceptible difference between the watermarked and original signal, and the watermark should be difficult to remove or alter without damaging the host signal.

tool in protecting digital content and many excellent papers have appeared in special issues [1], [2], as well as dedicated conferences and workshops [3]-[5]. New companies dedicated to watermarking technology are emerging and products like Digimarc's MediaBridge are appearing [6]. Unlike encryption, which is useful for transmission but does not provide a way to examine the original data in its protected form, the watermark remains in the content in its original form and does not prevent a user from listening to, viewing, examining, or manipulating the content. Also, unlike the idea of steganography, where the method of hiding the message may be secret and the message itself is secret, in watermarking, typically the watermark embedding process is known and the message (except for the use of a secret key) does not have to be secret. In steganography, usually the message itself is of value and must be protected through clever hiding techniques and the "vessel" for hiding the message is not of value. In watermarking, the effective coupling of message to the "vessel," which is the digital content, is of value and the protection of the content is crucial. Watermarking is the direct embedding of additional information into the original content or host signal. Ideally, there should be no perceptible difference between the watermarked and original signal [7], [8] and the watermark should be difficult to remove or alter without damaging the host signal. In some instances, the amount of information that can be hidden and detected reliably is important. It is easy to see that the requirements of imperceptibility, robustness, and capacity conflict with each other. For instance, a straightforward way to provide an imperceptible watermark is to embed the watermark signal into the *perceptually insignificant* portion of the host data. However, this makes the watermark vulnerable to attack because it is fairly easy to remove or alter the watermark without affecting the host signal. To provide a robust watermark, a good strategy is to embed the watermark signal into the significant portion of the host signal. This portion of the host data is highly sensitive to alterations, however, and may produce very audible or visible distortions in the host data. Applications for digital watermarking include copyright protection, fingerprinting, authentication, copy control, tamper detection, and data hiding applications such as broadcast monitoring. Watermarking algorithms have

been proposed for audio, still images, video, graphics, and text, and excellent review articles on multimedia watermarking can be found in [9]-[13].

Visible watermarks which do not interfere with the intelligibility of the host signal have also been proposed [14]. In this article, we limit the scope of our review to transparent marking techniques. Transparent watermarking techniques can be *fragile*, *robust*, or *semifragile*. Fragile watermarks do not survive lossy transformations to the original host signal and their purpose is tamper detection of the original signal. There are many effective ways to insert a fragile watermark into digital content while preserving the imperceptibility requirement. Placing the watermark information into the *perceptually insignificant* portions of the data guarantees imperceptibility and provides fragile marking capabilities. For instance, early watermark techniques for still image data propose inserting watermark information into the least significant bits of the pixel values. This results in an imperceptible mark which can detect lossy transformations performed on the watermarked content. For security applications and copyright protection, robust watermarking techniques have been proposed. Here the technical challenge is to provide transparency and robustness which are conflicting requirements. Ideally, an effective, robust watermarking scheme provides a mark that can only be removed when the original content is destroyed as well. The degree of robustness and distortion necessary to alter the value of the original content can vary for different applications. Typically, many of the applications for copyright protection involve relatively high quality original content and the imperceptibility criterion is critical for such applications. The authors in [15] and [16] were the first to describe that in order for a watermarking technique to be robust, the watermark should be embedded in the *perceptually significant* portion of the data. Some typical distortions or attacks that digital watermarking schemes are expected to survive include resampling, rescaling, compression, linear and nonlinear filtering, additive noise, A/D and D/A conversion, and transcoding. Applications for robust watermarking include copyright protection where each copy gets a unique watermark (commonly referred to as a fingerprint) to identify the end-user so that tracing is possible for cases of illegal use; authentication, where the watermark can represent a signature and copy control for digital recording devices. Within the class of robust watermarking techniques there are several different constraints on encoder and decoder design which depends on the particular application. The differences are discussed in detail later in this paper. Semifragile watermarking techniques differentiate between lossy transformations that are "information preserving" and lossy transformations which are "information altering." Lossy transformations include any signal processing step that alters the original signal values and is not invertible. For example, in authentication applications it may be desirable to have a watermark

that can distinguish between a lossy transformation such as compression which does not alter the integrity of the content and an alteration which does alter the integrity, such as manipulating or replacing objects within the scene.

Requirements and design of watermarking techniques are impacted by the different types of content in two major ways: imperceptibility and robustness requirements. The first challenge is designing a watermark embedding algorithm which provides an imperceptible mark, that is, one which does not noticeably degrade the original host signal. By taking advantage of psychovisual and psycho-auditory properties, we can design effective watermarking schemes which remain transparent under particular conditions [7], [8], [17]-[22]. Ideally, the marking algorithm should be adapted by using perceptual models appropriate for the different media types. The perceptual models used for representations of continuous tone images are not appropriate for text or graphics. The other factor for designing watermarking schemes for multimedia is the type of degradations that the watermark is expected to survive and system requirements for media specific applications. For instance, it may be desirable for a still image watermarking technique to be able to survive JPEG compression and photocopying while for some video watermarking applications, it may be important to do watermark embedding and detection in real time on a compressed bit stream.

In the next section we describe watermarking for different media types including an overview of some sample algorithms proposed in the literature. This is followed by a description of a general framework for watermark embedding and watermark detection and decoding, outlining some of the differences for different applications. We then review some work on modeling the general watermarking problem and drawing parallels to communication and information theory to help understand the fundamental properties and limitations of a watermarking system. This work is very useful for future algorithm design and helping to define open areas of research. Lastly, we review and summarize future directions in this new and exciting area.

Media Requirements

Here we explore the requirements for watermarking systems designed for different media types and review some of the algorithms covered in the literature.

Image Watermarking

Many techniques have been developed for the watermarking of still image data. For grey-level or color-image watermarking, watermark embedding techniques are designed to insert the watermark directly into the original image data, such as the luminance or color components or into some transformed version of the original data to take advantage of perceptual properties or

robustness to particular signal manipulations. Requirements for image watermarking include imperceptibility, robustness to common signal processing operations, and capacity. Common signal processing operations which the watermark should survive include compression (such as JPEG), filtering, rescaling, cropping, A/D and D/A conversion, geometric distortions, and additive noise. Capacity refers to the amount of information (or payload) that can be hidden in the host image and detected reliably under normal operating conditions. Many of the watermarking techniques are additive, where the watermark signal is added directly to the host signal or transformed host signal. The watermark may be scaled appropriately to minimize noticeable distortions to the host. Perceptual models may be used to determine and adapt the watermark scale factor appropriately to the host data. The watermark itself is a function of the watermark information, a secret or public key and perhaps the original host data. Some examples of watermark information includes a binary sequence representing a serial number or credit card number, a logo, a picture, or a signature. Many of the current watermarking techniques insert one bit of information over many pixels or transform coefficients and use classical detection schemes to recover the watermark information. These types of watermarking techniques are usually referred to as spread-spectrum approaches, due to their similarity to spread-spectrum communication systems. For still image watermarking, watermark embedding is applied directly to the pixel values in the spatial domain or to transform coefficients in a transform domain such as the discrete cosine transform (DCT) or discrete wavelet transform (DWT). Watermark detection usually consists of some preprocessing step (which may include removal of the original host signal if it is available for detection) followed by a correlation operator. More details on watermark embedding and detection appear later. Spatial-domain watermarking techniques for image data include [23]-[28]. Some of the earliest techniques [23], [29], [28] embed m -sequences into the least significant bit (LSB) of the data to provide an effective transparent embedding technique. M -sequences are chosen due to their good correlation properties so that a correlation operation can be used for watermark detection. Furthermore, these techniques are computationally inexpensive to implement. Such a scheme was first proposed in [23] and extended to two dimensions in [29]. In [28] the authors reshape the m -sequence into two-dimensional watermark blocks which are added and detected on a block-by-block basis. The block-based method, referred to as *variable- w two-dimensional watermark* (VW2D) is shown to be robust to JPEG compression. This technique has also been shown to be an effective fragile watermarking scheme which can detect image alterations on a block basis [30]. Other early work [31] suggests using check sums for LSB watermark embedding.

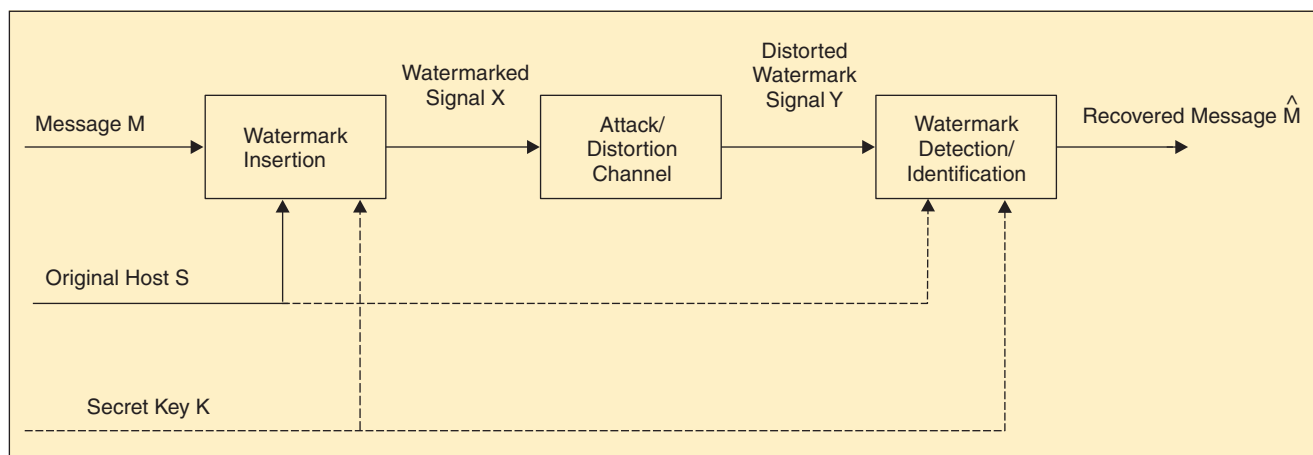
Several spatial-domain watermarking techniques for images are proposed in [25]. One technique consists of embedding a texture-based watermark into a portion of the image with similar texture. The idea here is that due to the similarity in texture, it will be difficult to perceive the watermark. The watermark is detected using a correlation detector. Another technique described as the *patchwork* method divides the image into two subsets A and B where the brightness of one subset is incremented by a small amount and the brightness of the other set is decremented by the same amount. The incremental brightness level is chosen so that the change in intensity remains imperceptible. The location of the subsets is secret and assuming certain properties for image data, the watermark is easily located by averaging the difference between the values in the two subsets. It is assumed that, on average, without the watermark, this value will go to zero for image data. In the example where the pixels in Set A are incremented by one and the pixels in set B are decremented by one, with N locations in the set, the expected value of the sum of differences between the sets is given by $2N$. For nonwatermarked data, this value should go to zero. A variation of this approach is described in [27], where more information can be inserted in the host signal. Another spatial-domain technique is proposed in [32], where the blue component of an image in RGB format is watermarked to ensure robustness while remaining fairly insensitive to human visual system (HVS) factors.

Transform domain watermarking is useful for taking advantage of perceptual criteria in the embedding process, for designing watermarking techniques which are robust to common compression techniques, and for direct watermark embedding of compressed bit streams. A common transform framework for images is the block-based DCT which is a fundamental building block of current image coding standards such as JPEG and video coding standards such as the MPEG video coders [33] and the ITU H.26x family of codecs. One of the first block-based DCT watermarking technique is proposed in [34]. The DCT is performed on 8×8 blocks of data, a pseudorandom subset of the blocks are chosen and a tri-

plet of midrange frequencies are slightly altered to encode a binary sequence. This is a reasonable heuristic watermarking approach since watermarks inserted in the high frequencies are vulnerable to attack whereas the low frequency components are perceptually significant and sensitive to alterations. One of the most influential watermarking works [15], [16] was first to describe how spread spectrum principles borrowed from communication theory can be used in the context of watermarking. The published results show that the technique is very effective both in terms of image quality and robustness to signal processing and attempts to remove the watermark. The technique is motivated by both perceptual transparency and watermark robustness. One of the significant contributions in this work is the realization that the watermark should be inserted in the *perceptually significant* portion of the image in order for it to be robust to attack. A DCT is performed on the whole image and the watermark is inserted in a predetermined range of low frequency components minus the DC component. The watermark consists of a sequence of real numbers generated from a Gaussian distribution which is added to the DCT-coefficients. The watermark signal is scaled according to the signal strength of the particular frequency component. This is a reasonable and simple way to introduce some type of perceptual weighting into the watermarking scheme. The watermark embedding algorithm could be described as

$$X = S(1 + \alpha W) \quad (1)$$

where S is the original host signal, X is the watermarked signal, and W is the watermark consisting of a random, Gaussian distributed sequence. α is a scaling factor which the authors suggest to set to 0.1 to provide a good trade-off between imperceptibility and robustness. Referring to Fig. 1 for a block diagram of a general watermarking system, the secret key is used to generate the random sequence W in this case. Also note that this particular algorithm addresses the case of *watermark detection* where you would like to detect whether a particular watermark is or is not present in the host signal at the receiver. The



▲ 1. Block diagram of a watermarking system.

watermark detector for this scheme [15] is described by the similarity measure

$$\text{sim}(W, \hat{W}) = \frac{\hat{W} \cdot W}{\sqrt{\hat{W} \cdot \hat{W}}} \quad (2)$$


where \hat{W} is the extracted watermark from the received, possibly distorted signal Y . The authors show that the similarity measure is also normally distributed so that a high similarity value is extremely unlikely for $\hat{W} \neq W$. Other postfiltering operations could be performed to undo possible distortions, improve performance, and get a better similarity measure. More details on improving detection results can be found later.

A variation on this idea is variable length DCT-based watermarking [35], where the DCT coefficients are sorted by magnitude and only the n largest coefficients are marked that correspond to a user specified percent of the total energy. This allows the user to trade off imperceptibility and robustness to attack. Other DCT-based watermarking schemes use more elaborate models of the human visual system to incorporate an image adaptive watermark of maximum strength subject to the imperceptibility criterion [17], [7], [8]. Two image-adaptive watermarking schemes are described in [19] and [7], which are based on a block-based DCT framework and wavelet framework. The perceptual models used here can be described in terms of three different properties of the human visual system that have been studied in the context of image coding: frequency sensitivity, luminance sensitivity, and contrast masking [36]. Frequency sensitivity describes the human eye's sensitivity to sine wave gratings at various frequencies. This component only depends on the modulation transfer function (MTF) of the eye and is independent of the image data. Luminance sensitivity measures the effect of the detectability threshold of noise on a constant background. For the human visual system, this is a nonlinear function and depends on local image characteristics. Contrast masking refers to the detectability of one signal in the presence of another signal and the effect is strongest when both signals are of the same spatial frequency, orientation, and location. A combination of the three components results in *just noticeable distortion* (JND) thresholds for the entire image. These models were first developed to design more efficient image compression schemes than waveform techniques alone could provide. This model was derived for the baseline mode of JPEG and showed a significant improvement in compression performance when used to derive an



▲ 2. Watermarked images (first, third row) and corresponding image-adaptive watermarks using perceptual models (second, fourth row).

image-adaptive quantization table [36]. A similar model was developed for wavelet-based compression using only frequency sensitivity to derive perceptual weights for each of the subbands [37]. This model was used for a wavelet-based watermarking scheme [19], [7]. Unlike compression, where the amount of perceptual information that can be incorporated into the encoder is limited to the amount of side information that is necessary to transmit this information to the decoder, all of the perceptual information can be utilized in a watermarking scheme. For instance, in JPEG, we are limited to one quantization matrix for the entire image which cannot take full advantage of local visual threshold characteristics. The image dependent masking thresholds are used to determine the location and maximum strength of the watermark signal that can be tolerated in every location of the host image under the constraint of imperceptibility at some specified viewing condition. Examples of the image-adaptive watermarks described in [7] are illustrated in Fig. 2.



▲ 3. Watermark example for different viewing distances.

Note how the watermark structure is similar to the local image properties. Fig. 3 illustrates several watermarked images and the corresponding watermark to the right of the watermarked image. The three examples show the watermarked image for different viewing conditions where the top image corresponds to a viewing distance of four times the image height, the middle image corresponds to a viewing distance of two times the image height and the bottom image corresponds to a viewing distance of one times the image height. Modifying the viewing conditions of the watermark embedding algorithm allows for a tradeoff between imperceptibility and robustness to certain types of attacks. These examples are for illustrative purposes and the viewing conditions here are not based on viewing printed images.

Two DCT-based approaches were described in [38] and [39] where watermark detection does not require the original image. The method in [40] is an extension of the method proposed in [19] and [7] to the case where the original host signal is not available for watermark detection. This is an important feature for some applications such as authentication and will be covered in more detail later. Another block-based frequency domain technique described in [41] is based on inserting a watermark into

the phase components of the image data using the same motivation as in [16], that for the watermark to be robust to attack, it must be embedded in the *perceptually significant* portion of the data. It has been established that for image data, the phase information is perceptually more important than the magnitude data. Other novel approaches for watermarking image data include fractal-based approaches [42], [43] and geometric feature based watermarking [44]. In [44], salient points in an image are found and warped according to a dense line pattern representing the watermark and generated randomly. Detection consists of determining whether a significantly large number of points are within the vicinity of the line patterns.

The type of distortions or attacks that image watermarking techniques are designed to survive fall into two broad categories—noise type distortions like compression and geometrical distortions which cause loss of synchronization for detection, such as resampling and rotations. Watermarking schemes for tamper detection and tamper estimation to be able to differentiate between lossy attacks which alter the information and lossy attacks which do not alter the information [45]-[47] have also been proposed.

Document Watermarking

Much of the early work on recognizing the potential problems with intellectual property rights of digital content and addressing these issues with early watermarking techniques was in the area of document watermarking [48]-[50]. These techniques were devised for watermarking electronic versions of text documents which are in some formatted version such as postscript or PDF. Most of this work is based on hiding the watermark information into the layout and formatting of the document directly. In [48]-[50], the authors develop document watermarking schemes based on line shifts, word shifts as well as slight modifications to the characters. These techniques are focused on watermarking the binary-valued text regions of a document. Watermark detection consists of postprocessing steps to try to remove noise and correct for skew. These techniques are quite effective against some common attacks such as multigenerational photocopying. The authors point out that optical character recognition can remove the layout information and, for such schemes, remove the watermark information. An open area of research remains in how to formulate format deviations in a perceptual framework. Fig. 4 illustrates an example from [49] of word shift coding; (a) shows where the space has been added before the word “for,” and (b) contains the unwatermarked and watermarked versions in their natural state to illustrate that the word shift is not noticeable. Fig. 5 shows an example from [49] on character alteration for watermark embedding.

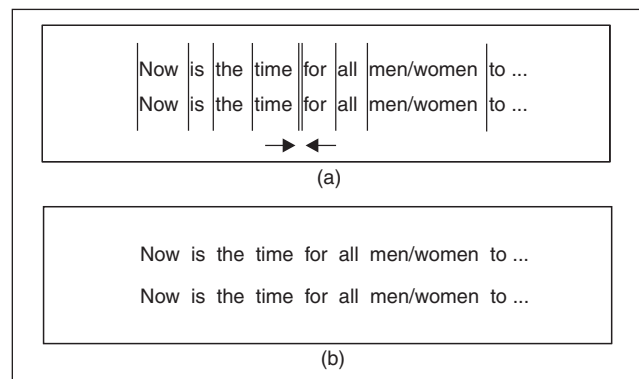
Graphics Watermarking

There has been some work on effective watermarking of graphics, motivated in part by such standards as MPEG-4. In [51], the authors address watermarking three-dimensional polygonal models. The work in [52] addresses the watermarking of facial animation parameters as defined by the MPEG-4 standard. The watermark is embedded directly into the parameters and can be extracted from the watermarked parameters directly or from video sequences rendered using the parameter bit stream where the parameters are estimated using a model-based approach. One bit of watermark information is embedded in a block of facial animation parameter (FAP) data using a pseudonoise sequence that is generated from the secret key. The authors limit the amount of deviation the watermark signal has on the FAPs empirically to minimize visible distortion. For instance, global FAPs like head rotation are limited to deviate by 1% of their dynamic range while local FAPs such as lip motion is limited to 3%. Watermark detection can be done directly on the watermarked FAPs through a traditional correlation detector. The authors demonstrate that they are able to recover the watermark information without error using both the FAPs directly or by estimating them from a rendered sequence. They also show that their method is robust to moderate compression using MPEG-2. In general, watermarking of graphics data remains an interesting research topic since our understanding of perceptual models in this domain is not yet fully recognized.


Video Watermarking

The Copy Protection Technical Working Group (CPTWG), an ad hoc group consisting of the Motion Picture Association of America, the Consumer Electronics Manufacturers Association, and members of the computer industry, is examining digital video protection as it applies to digital versatile disk (DVD) technology [53], [54]. The current plan is to adopt a de facto standard for a DVD copy protection system which includes watermarking. The watermark component of the system, besides the usual requirements of robustness and transparency, must satisfy other constraints and system requirements unique to this application. In this case, the watermark is designed to support copy generation management, and the minimum information that the watermark must convey is: copy never, copy once, copy no more, and copy freely. A cost-effective solution for watermark detection is a critical requirement for DVD watermarking so that real-time decoding with no frame buffer (no reference to previous frames) is required. Low false positive rates is a critical component for the consumer driven DVD market and is much more important than the security risks associated with false negatives. Other issues that have arisen in the design of an effective copy control system for DVD includes the placement of the detector. The two remaining proposals have two different approaches for detector placement—watermark

detection in the drive and watermark detection within the application (within the MPEG decoder). The drive-based solution has the advantage that as long as the watermark exists, pirated content cannot leave the drive in playback mode or recording mode. There is some added complexity with detection in the drive versus detection in the application, for example, a partial decode of the MPEG bit stream is necessary. Watermark detection in the MPEG decoder is not as secure as the drive-based solution and some added features that have been suggested for detection in the application include a protocol to recognize a compliant device, a bidirectional link with authentication, encryption and data integrity, and a protocol between source and sink which informs the drive whether to stop transmitting data. Advantages of application-based detection are the ability to provide a more complex detector and the flexibility of extending the scheme to other data types. The other unique requirement for DVD applications is copy generation management, that is, the ability to detect the *copy once* state and change it to a *copy no more* state after the recording. The two proposals have different approaches for this feature as well—secondary watermarks and tickets. The secondary watermark approach adds a second watermark after the recording. The second-



▲ 4. Example of word-shift coding. In (a), the top text line has added spacing before the “for,” the bottom text line has the same spacing after the “for.” In (b), these same text lines are shown again without the vertical lines to demonstrate that either spacing appears natural. From [49].



▲ 5. Example shows feature coding performed on a portion of text from a journal table of contents. In (a), no coding has been applied. In (b), feature coding has been applied to select characters. In (c), the feature coding has been exaggerated to show feature alterations. From [49].

ary watermark embedder must be computationally inexpensive, must be applicable in the baseband and compressed video domains, and should not alter the bit rate in MPEG embedding. The second approach uses a ticket which is a cryptographic counter implemented as a multibit random number. The recorder modifies the ticket by passing it through a cryptographic one-way function (hash function) where each time it goes through a player, it gets decremented by one. An excellent review article on this topic can be found in [53].

Other general video watermarking techniques have also been described in the literature. A scene-adaptive video watermarking technique is proposed in [18] where the watermarking scheme is based on a temporal wavelet decomposition. The wavelet decomposition separates static areas from dynamic areas so that separate watermarking strategies can be applied to the different areas. The authors propose a constant watermark for the static area and a varying watermark for the dynamic areas to defeat watermark deletion through frame averaging.

Many times, digital video will already be in a compressed format at the point where watermarking is applied, and it is desirable to be able to embed the watermark directly into the compressed bit stream without going through a full decoding, watermarking, and reencoding step which adds considerable complexity and additional delay. Interesting work [55], [56] on watermarking of uncompressed and compressed video has been studied. One of the issues addressed in this work is the direct embedding of watermark information in a compressed video bit stream, subject to the imperceptibility constraint as well as an additional constraint that the total bit rate of the watermarked compressed bit stream cannot exceed the total bit rate of the unwatermarked bit stream. This is an important requirement because for many applications, bandwidth limitations dictate the total bit rate possible for the video stream. Current video compression standards such as MPEG or ITU H.26x standards consist of the same general framework which includes block based motion compensation which takes advantage of temporal correlation and block-based DCT coding which takes advantage of local spatial correlations. The watermarking technique does not alter the motion vector information which is used for the motion compensation and is encoded in a lossless manner or any of the critical side information. The watermark signal is only embedded into the DCT coefficients so that only partial decoding of the block DCT is necessary for watermark embedding. Only nonzero DCT coefficients are marked and if constant bit rate is required, DCT coefficients are marked only if the bit rate for the quantized representation is equal or less than the bit rate needed for the unmarked quantized coefficients. This is possible due to variable length coding. The watermark embedding process consists of inverse entropy coding and inverse quantization, embedding the watermark in the DCT coefficients and checking for bit rate compliance. Al-

though much of the video may not be marked due to this additional constraint, it is still possible to embed a few bytes of information per second, which is useful for many applications.

In other work [57], [58], two techniques are introduced for real-time watermark embedding of compressed video. One technique adds the watermark by modifying the fixed length and variable length codes in the compressed video bit stream. This allows for a computationally efficient way of real-time watermark insertion and allows for a relatively high payload. The drawback of this technique is that decoding the bit stream removes the watermark completely. A more robust technique is also proposed which adds a watermark by enforcing energy differences between various video regions. This is done by discarding high frequency components so that only partial decoding of a compressed video bit stream is necessary to apply this watermark. This technique results in a watermark that is still present after decompressing the video bit stream.

In [59] a video watermarking method is proposed for broadcast monitoring where encoder and decoder complexity are critical requirements. The low complexity scheme consists of spatial domain encoding and decoding with a perceptually based scaling factor that depends on a simple measure of local activity. Other techniques proposed for video watermarking of compressed bit streams includes embedding the watermark information in the motion vectors [60]. A DCT-based watermarking scheme for video which is motivated by previous still image watermarking techniques is introduced in [61].

Other requirements for video watermarking may include real-time watermark detection/identification and perhaps real-time watermark embedding, robustness to NTSC/PAL conversion, MPEG compression, frame averaging attack, A/D and D/A conversion, and rate control. Other broadcast applications for hiding additional information are described in [62].

Audio Watermarking

Most of the research on audio watermarking has been focused on either direct watermarking of the audio signal or bit stream embedding where the audio is represented in a compressed format. Just as in image and video watermarking, the use of perceptual models is an important component in generating an effective and acceptable watermarking scheme for audio [21], [63], [25]. Many of the requirements for audio watermarking are similar to image watermarking, such as imperceptibility (inaudibility), robustness to signal alterations such as compression, filtering, and A/D and D/A conversion. In [25], the authors propose three techniques for audio watermarking—a spread spectrum technique, echo coding, and phase coding. The approach described in [21] and [63] consists of generating a PN-sequence for the watermark and processing it with a filter that approximates the fre-

quency masking properties of the human auditory system (HAS), followed by a time-domain weighting for temporal masking. Correlation properties of PN-sequences are desirable for detection and applying an auditory model guarantees imperceptibility—a critical feature for high quality audio clips where copyright protection may be most critical. Masking is the phenomena where the detectibility of a signal component depends on the presence or absence of other signal components in its immediate vicinity either in the frequency domain or temporal or spatial domain. Here, detectibility refers to audibility for audio or visibility for image and video signals. An overview paper on how perceptual models have been exploited for signal compression can be found in [64]. The audio watermarking technique in [21] and [63] uses the frequency masking model proposed in MPEG. More details on generating the thresholds can be found in [63]. Watermark embedding consists of adding a perceptually weighted PN-sequence to the audio file while watermark detection consists of a correlation detector to determine whether the watermark is or is not present in the received signal.

The Secure Digital Music Initiative (SDMI) that consists of companies and organizations in information technology, consumer electronics, security technology, the recording industry, and ISPs has been formed to examine technology which provides some security features for digital music and copyright protection for next-generation portable digital music devices. Phase I screening looks for a watermark in the content but allows all music that is compatible with the device to be playable. Phase II will incorporate watermark detection which will allow new releases to play while filtering out pirated copies of music. After extensive testing of imperceptibility and robustness, SMDI has chosen ARIS audio watermarking technology for Phase I screening technology which will be used to indicate when the software used by Phase I devices should be upgraded to incorporate Phase II technology. Some of the requirements particular to music as seen by the SDMI group includes inaudibility, robustness, tamper resistance, reliability (no false positives), ease of implementation, cost, and ability to compress the content. Details of other watermarking technology for audio can be found in [10].

Watermark Embedding

The watermark embedding scheme can either embed the watermark directly into the host data or to a transformed version of the host data. Some common transform domain watermarking for image data can be DCT based [7], [34], [65], [16], [17], [38] or wavelet based [18], [7]. Transform-domain techniques are popular due to the natural framework for incorporating perceptual knowledge into the embedding algorithm and because many of the state-of-the-art compression techniques such as JPEG

work in the same framework (block-based DCT) and this allows for watermarking of the compressed bit stream with only partial decoding. A simple way of applying some perceptual knowledge is to watermark the midfrequency components, since the low frequency components are very sensitive to distortion and the high frequency components can be removed without significantly affecting the original image quality. Use of more formal perceptual models for watermark embedding have also been developed [7], [17], [18], [8]. A review article on using perceptual models for watermarking can be found in [8]. The earliest watermarking techniques involved embedding a low energy pseudorandom noise pattern directly to the digital host signal (for example, image luminance values) [28], [23]-[25].

A basic block diagram of a watermark system is illustrated in Fig. 1 where S denotes the original host signal and can represent image luminance values or some transform domain signal such as the DCT coefficients. M denotes the watermark message which, for example, can be a sequence of bits representing a serial number or credit card number, one bit in the case of a signature for authentication applications, a logo or picture. When the message M is used to identify the destination or end-user to help track illegal usage later, M is sometimes referred to as a fingerprint and recovering M is known as *identification*. The watermark signal can either represent a signature where the goal is to determine whether or not the signature is present in the content (*detection* or *verification*) or a sequence of information bits or other data where the goal is to extract the bit pattern with low probability of bit error or to identify one out of N possible watermark messages. The watermark can be binary or real valued. The watermark is usually parameterized by a key K which is secret and could be used to generate a random sequence to embed in the host signal as described in [15] and expressed in (1). This key could also be used to determine a random sequence which identifies locations in the host signal for watermark embedding. Without knowledge of the key, it should be difficult to remove or alter the embedded message without destroying the original content. For many applications, just as in cryptography, watermarking algorithms follow Kerckhoff's principle, that is, the watermark embedding process is public and security is based only on choosing a secret key. The watermark information or key could also be dependent on the host signal. For instance, the secret key may depend on a hash of the host signal. This is a particularly useful feature for the invertible watermark attack outlined in [66].

There are also applications where a “no-key” or “public-key” system may be desirable [67], [68]. The dotted lines in Fig. 1 represent optional components that may or may not be present depending on the application. In summary, from Fig. 1, a secret key may or may not be present at the encoder and decoder and the original host signal may or may not be present at the decoder.

Transform domain watermarking is useful for taking advantage of perceptual criteria in the embedding process, for designing watermarking techniques which are robust to common compression techniques, and for direct watermark embedding of compressed bit streams.

Some of the watermarking techniques described in the literature are simple additive watermarking schemes expressed as

$$X = S + W \quad (3)$$

where W is the watermark signal and could depend on the secret key K and the message to be embedded M . Most spread spectrum techniques, however, use some sort of perceptual weighting and modulate the watermark signal according to some properties of the host signal itself so that the simple expression in (3) does not hold for watermark embedding. An example of a spread-spectrum technique which uses the magnitude of the DCT coefficients to modulate the watermark signal is described in [15] and [16] and the embedding algorithm is expressed in (1). In the image-adaptive schemes described in [7] and [8], the watermark signal strength is modulated for every DCT or wavelet coefficient based on the local properties of the host data. Examples of image-adaptive watermarked images and corresponding watermarks based on the algorithms in [7] are illustrated in Fig. 2. Note how the structure of the watermark strength is highly correlated with the structure of the underlying host signal. The watermark signal strength is strongest in the high frequency details, edges and textures.

Another type of watermark embedding technique outside the family of spread spectrum watermarking and LSB watermarking is quantization index modulation (QIM) which was first introduced in [67] and [68]. The authors describe a dither modulation approach as a particular example of QIM where the watermark embedding step can be described by

$$X = Q(S + d(M)) - d(M) \quad (4)$$

where $d(M)$ is a dither vector and Q represents a quantization operator. The watermark information is conveyed in the choice of quantizer. QIM systems are especially useful for applications where it is desirable to have a “no-key” system, that is, a system where the de-

coder is public and no secret key is required so that anyone can embed and detect a watermark.

Besides embedding the original message, many techniques also embed some form of redundancy such as simple repetition codes or more complex channel codes such as Reed-Solomon codes designed for traditional communications systems to provide better detection capabilities and lower probability of bit error.

Watermark Detection

In keeping consistent with the taxonomy of earlier detection and estimation problems, we differentiate between detection and identification at the watermark receiver. Detection or verification refers to the process of making a binary decision at the decoder—whether a specific watermark is or is not present in the received data. This may be appropriate for authentication applications where you would like to verify that a signature is present in the received content. This problem lends itself to a hypothesis testing formulation, and the effectiveness of the watermark scheme can be measured in terms of *Type I* and *Type II* errors. Type I errors or false positives refer to the case where a watermark is detected when it does not exist, and Type II errors or false negatives refer to the case when an existing watermark is not detected. For many applications, especially in the consumer markets, it is more important to have low or zero false positives at the risk of higher false negatives rates. This is also referred in the literature as the probability of false alarm and probability of detection. Plots of probability of detection versus the probability of false alarm are referred to as receiver operating characteristics (ROC) curves. *Identification* refers to the process of being able to decode one of N possible choices (messages) at the receiver. An application for this includes copyright protection where multiple copies of the same content get a unique label so that misuse of one of the copies can be traced back to its owner. Identification problems can be categorized as “open set” or “closed set.” Open set identification refers to the possibility that one of N or no watermark exists in the data. Closed set refers to problems where one of N possible watermarks is known to be in the received data and the detector has to pick the most likely one. For identification problems where the goal is to extract a watermark sequence, for instance a binary sequence of length B where one of $N = 2^B$ watermark patterns is present, the bit error rate (BER) is a very useful measure of performance. The effectiveness of a watermarking scheme can be illustrated by plotting the BER versus SNR in the case of an additive noise watermarking attack as shown in Fig. 6. In this example, 32 bits of watermark information were inserted into the “Lena” image with a simple repetition code for protection. The effectiveness of the watermarking scheme was tested by detecting the bits in the presence of an additive noise attack (this is a good model for some common transformations such as compression). The BER is ap-

proximately zero until $\sigma = 3.5$ for the noise term which results in severely distorted image quality as shown in Fig. 7. Refer to [69] for the details of the experiment. Another important factor for multiple bit watermarking is payload or capacity—how many bits can be reliably detected (low BER) for a given application. This is especially useful for data hiding applications where security (robustness) may not be critical. Watermark detection may include a secret key and the original content, it may include the secret key and no original (blind detection), or it may involve no key or a public key and no original where it is desirable to allow anyone to mark or detect a watermark.

Many of the watermarking schemes are based on the general concepts of spread spectrum communications and a classical correlation detector is used for watermark detection or decoding [70]. In Fig. 1, \hat{M} is the recovered message, Y is the received, possible distorted watermarked signal, S is the original content, and K is the secret key. For many applications, S is not available for watermark decoding and this is referred to as blind detection. In this case, the original signal S acts as an additive noise component in the watermark detection process for the simple additive watermarking scheme. Also, when the original is available at the decoder, it could be used to estimate the channel distortions and invert them to provide better detection performance. For the case when the original is available for detection, and the watermark embedding algorithm is a simple additive process, a typical watermark detector can be described by the normalized correlation operation


$$\hat{W} = Y - S$$

$$\rho_{w\hat{w}} = \frac{\hat{W} \cdot W}{\sqrt{E_{\hat{w}} \cdot E_w}} \quad (5)$$

If W is identical to \hat{W} and normally distributed, the correlation coefficient goes to one. Watermark detection is performed by comparing the correlation coefficient to a threshold value which can be modified according to the tradeoff between probability of detection and the probability of false alarm that is appropriate for a particular application. The final step for watermark detection for the binary case is

$$\rho_{w\hat{w}} > T_\rho \text{ watermark } W \text{ detected}$$

$$\rho_{w\hat{w}} \leq T_\rho \text{ watermark } W \text{ is not detected.} \quad (6)$$




▲ 6. On the robustness of our watermarking technique: (a) estimated probability of decoding error; (b) PSNR numbers for different amounts of jamming noise. From [69].

It is advantageous to process the received signal Y , to try to estimate and inverse any distortions introduced by the channel.

Work exploring attacks and counterattacks for watermarking schemes [66], [71], [72] has been very useful in understanding and helping evolve the state-of-the-art in watermarking algorithms. For instance, the authors in [66] describe a way to defeat additive watermarks through invertible watermarking which leads to an ambiguity attack. Many techniques have been proposed to overcome such an attack. Some methods propose making the watermark depend on the original data (for example through a hash function) [66], [73], [74] while others propose using secure time stamps provided by third parties [74].

Several watermark attacks are based on losing synchronization so that effective detection through a correlation operator fails. Using portions of the watermark to embed a known synchronization marker has been proposed to overcome this problem due to either cropping or translation [75]. Dealing with more general geometric distortions is also addressed in [76] by considering affine transformations which can be used to model scaling, rotation, or shearing. The basic idea here is to insert a reference pattern along with the watermark into the original image to be able to identify the geometric transformation from the distorted reference signal and invert it. The reference pattern proposed is multiple embedding of the same watermark at different locations. Other methods to deal with synchronization attacks for a correlation-based detector have been suggested. In [77], the authors pro-



▲ 7. Visual quality assessment: (a), a section of the original Lena; (b) a section of the jammed Lena. For this much degradation, over 1000 tests, all watermarks were correctly retrieved. From [69].

pose using transformation invariant domains for watermark embedding. For instance, watermark embedding in the Fourier-Mellin transformation domain is invariant to translation, scale, and rotation. Using a calibration signal as part of the watermark has also been suggested for geometrical distortions by Digimarc Corporation [6].

Fundamental Properties and Limitations of Watermarking

Much of the work on trying to model and understand some of the fundamental properties and limitations of watermarking algorithms is based on drawing parallels to communications systems. We have already mentioned that many of the popular watermark embedding algorithms are variations on the idea of spread-spectrum techniques for secure communication systems where an information bearing narrowband signal is converted into a wideband signal prior to transmission, by modulating the information waveform with a wideband noiselike waveform. As a result of the bandwidth expansion, within any narrow spectral band, the total amount of energy from the information signal is small. By appropriately combining all the weak narrowband signals at the demodulator, the original information signal is recovered.

There has been some interesting work in trying to model and understand some of the fundamental properties and limitations of watermarking algorithms [78]-[82], [69]. An information theoretic analysis of watermarking is presented [78] where an elegant framework is proposed for the hiding capacity problem (watermark payload). The framework shows the tradeoff between achievable information hiding rates and allowed distortions for the information hider (watermark embedder) and the attacker (possible distortions to remove or alter the watermark). Under particular conditions, the optimum marking strategy and optimum attacking strategy are shown. A similar approach was outlined in [69].

The work described in [80] derives a simple model for watermark embedding and attacks. The attack is a Wiener

estimate of the actual watermark signal which leads to an effective watermark design which attempts to match the power spectrum of the watermark as a scaled version of the power spectrum of the original host signal. Intuitively, this says that the watermark should look like the original signal. This also supports the use of visual models for watermark embedding where the watermark signal very closely matches the general characteristics of the host signal.

In [82], the authors address the problem of how to effectively model quantization (typically the lossy step to any data compression scheme) as a form of attack on watermarked data. Unlike previous work where quantization is usually modeled as additive noise which is adequate for fine quantization or high data rates, the authors look at modeling the watermarking and quantization effect as dithered quantization where the dither is represented by the watermark. They show that for the quantization attack, a Gaussian distributed watermark is more robust than a uniform or bipolar one and even more importantly, a Gaussian distributed host signal provides better detection results than a Laplacian source.

Other theoretical work addressing watermark detection can be found in [83]-[85].

Conclusion and Future Directions

We have reviewed the basic watermarking algorithms as they apply to different applications and media types. Although many technical problems have been addressed, there are many more yet to be solved. Many of the techniques developed for watermarking are based on a solid understanding of communications and signal processing principles, but there are still many technical challenges to be solved. It is difficult to model the distortions introduced by common signal processing transformations, which either intentionally or unintentionally affect the watermark detection or identification capabilities. Although very nice work exists in trying to understand the fundamental limitations of watermark embedding and detection, attack channels such as geometrical distortions cannot be described by these models. Other areas have not been resolved as well. Besides the obvious caveat of whether watermarking technology will be effective in a court of law, other questions remain. What are reasonable distortions for particular applications that the watermark is expected to survive? What is a meaningful measure of distortion that can be used to determine the effectiveness of a watermarking scheme? How is monitoring and policing for copyright infringement done? How much are potential customers for watermarking technology willing to pay for it?

These questions and many interesting technical challenges remain in this new and exciting field. The overview we have presented is meant to summarize the salient features and directions of watermarking research and tech-

nology and the interested reader is encouraged to explore the references for more details.

Acknowledgement

The authors would like to thank J. Brassil, L. O'Gorman, S. Low, and N. Maxemchuk for providing sample images of their work published in [49].

Christine I. Podilchuk received the B.S., M.S. and Ph.D. degrees in electrical engineering from Rutgers University, New Brunswick, NJ, in 1984, 1986, and 1988, respectively. Since then, she has been with Bell Laboratories, Murray Hill, NJ. Her research interests are in signal processing, particularly for image and video applications. Her most recent work includes digital watermarking for multimedia, video compression, and video transmission over wireless networks.

Edward J. Delp received the B.S.E.E. (cum laude) and M.S. degrees from the University of Cincinnati, and the Ph.D. degree from Purdue University. From 1980-1984, he was with The University of Michigan, Ann Arbor. Since 1984, he has been a Professor of Electrical and Computer Engineering and Professor of Biomedical Engineering at Purdue University, West Lafayette, IN. His research interests include image and video compression, multimedia security, medical imaging, communication, and information theory. He is a member of Tau Beta Pi, Eta Kappa Nu, Phi Kappa Phi, Sigma Xi, ACM, and the Pattern Recognition Society. He is a Fellow of the IEEE, a Fellow of the SPIE, and a Fellow of the Society for Imaging Science and Technology (IS&T).

References

- [1] Special Issue on Copyright and Privacy Protection, *IEEE J. Select. Areas Commun.*, vol. 16, May 1998.
- [2] Special Issue on Identification and Protection of Multimedia Information, *Proc. IEEE*, vol. 87, no. 7, July 1999.
- [3] "IS&T and SPIE electronic imaging," presented at the Conference on Security and Watermarking of Multimedia Contents, San Jose, CA, 1999, 2000.
- [4] Erlangen Watermarking Workshop, Erlangen, Germany, 5-6 Oct. 1999.
- [5] International Workshop on Information Hiding, 1996.
- [6] <http://www.digimarc.com>.
- [7] C. Podilchuk and W. Zeng, "Image-adaptive watermarking using visual models," *IEEE J. Select. Areas Commun.*, vol. 16, pp. 525-539, May 1998.
- [8] R.B. Wolfgang, C.I. Podilchuk, and E.J. Delp, "Perceptual watermarks for digital images," *Proc. IEEE*, no. 7, pp. 1108-1126, July 1999.
- [9] R. Wolfgang and E. Delp, "Overview of image security techniques with applications in multimedia systems," in *Proc. SPIE Int. Conf. on Multimedia Networks: Security, Displays, Terminals and Gateways*, Nov. 1997, vol. 3228, pp. 297-308.
- [10] M. Swanson, M. Kobayashi, and A. Tewfik, "Multimedia data-embedding and watermarking strategies," *Proc. IEEE*, vol. 86, pp. 1064-1087, June 1998.
- [11] F. Hartung and M. Kutter, "Multimedia watermarking techniques," *Proc. IEEE*, vol. 87, pp.1-79-1107, July 1999.
- [12] J. Su, F. Hartung, and B. Girod, "Digital watermarking of text, image, and video documents," *Comput. Graph.*, vol. 22, no. 6, pp. 687-695, Feb. 1999.

- [13] F. Petitcolas, R. Anderson, and M. Kuhn, "Information hiding—A survey," *Proc. IEEE*, vol. 87, pp. 1062-1078, July 1999.
- [14] G. Braudway, K. Magerlein, and F. Mintzer, "Protecting publicly available images with a visible watermark," in *Proc. SPIE Conf. Optical Security and Counterfeit Deterrence Techniques*, vol. 2659, Feb. 1996, pp. 126-132
- [15] I.J. Cox, J. Kilian, T. Leighton, and T. Shamoon, "Secure spread spectrum watermarking for multimedia," NEC Research Institute, Princeton, NJ, Technical Report 95-10, 1995.
- [16] I. Cox, J. Kilian, T. Leighton, and T. Shamoon, "Secure spread spectrum watermarking for multimedia," *IEEE Trans. Image Processing*, vol. 6, pp. 1673-1687, Dec. 1997.
- [17] M.D. Swanson, B. Zhu, and A.H. Tewfik, "Transparent robust image watermarking," in *IEEE Proc. Int. Conf. Image Processing*, Lausanne, Switzerland, 1996, pp. 211-214.
- [18] M. Swanson, B. Zhu, and A. Tewfik, "Multiresolution scene-based video watermarking using perceptual models," *IEEE J. Select. Areas Commun.*, vol. 16, pp. 525-539, May 1998.
- [19] C.I. Podilchuk and W. Zeng, "Digital image watermarking using visual models," *IS&T, SPIE Human Vision and Electronic Imaging II*, pp. 100-111, Feb. 1997.
- [20] C. Podilchuk and W. Zeng, "Perceptual watermarking of still images," in *Proc. IEEE Workshop Multimedia Signal Processing*, pp. 363-368, June 1997.
- [21] M. Swanson, B. Zhu, A. Tewfik, and L. Boney, "Robust audio watermarking using perceptual masking," *Signal Process., Special Issue on Watermarking*, 1997, pp. 337-355.
- [22] J.F. Delaigle, C.D. Vleeschouwer, and B. Macq, "A psychovisual approach for digital picture watermarking," *J. Electron. Imaging*, vol. 7, no. 3, pp. 628-640, July 1998.
- [23] R.G. Schyndel, A. Tirkel, and C. Osborne, "A digital watermark," in *Proc. IEEE Int. Conf. Image Processing (ICIP)*, 1994, pp. 86-90.
- [24] G. Caronni, "Assuring ownership rights for digital images," in *Proc. Reliable IT Systems, VIS '95*, 1995, pp. 251-263.
- [25] W. Bender, D. Gruhl, and N. Morimoto, "Techniques for data hiding," *IBM Syst. J.*, vol. 35, nos. 3-4, pp. 313-336, 1996.
- [26] K. Tanaka, Y. Nakamura, and K. Matsui, "Embedding secret information into a dithered multi-level image," in *IEEE Proc. Military Communications Conf.'90*, 1990, pp. 216-220.
- [27] I. Pitas, "A method for signature casting on digital images," in *IEEE Proc. Int. Conf. Image Processing*, Lausanne, Switzerland, Sept. 1996, vol. 3, pp. 215-218.
- [28] R.B. Wolfgang and E.J. Delp, "A watermark for digital images," in *IEEE Proc. Int. Conf. Image Processing*, Lausanne, Switzerland, Sept. 1996, vol. 3, pp. 219-222.
- [29] R.G. van Schyndel and C. Osborne, "A two-dimensional watermark," in *Proc. DICTA*, 1993, pp. 378-383.
- [30] R. Wolfgang and E. Delp, "Fragile watermarking using the vw2d watermark," in *Proc. SPIE, Security and Watermarking of Multimedia Contents*, San Jose, CA, Jan. 1999, vol. 3657, pp. 204-213.
- [31] S. Walton, "Information authentication for a slippery new age," *Dr. Dobbs J.*, vol. 20, no. 4, pp. 18-26, Apr. 1995.
- [32] M. Kutter, F. Jordan, and F. Bossen, "Digital watermarking of color images using amplitude modulation," *J. Electron. Imaging*, vol. 7, no. 2, pp. 326-332, Apr. 1998.
- [33] B.G. Haskell, A. Puri, and A.N. Netravali, *Digital Video: An Introduction to MPEG-2*. New York: Chapman & Hall, 1997.
- [34] E. Koch and J. Zhao, "Towards robust and hidden image copyright labeling," presented at Nonlinear Signal Processing Workshop, Thessaloniki, Greece, 1995.
- [35] F.M. Boland, J.J.K. O'Ruanaidh, and C. Dautzenberg, "Watermarking digital images for copyright protection," in *IEEE Int. Conf. Image Proc. and its Applications*, 1995, pp. 321-326.
- [36] A.B. Watson, "DCT quantization matrices visually optimized for individual images," in *Proc. SPIE, Conf. on Human Vision, Visual Processing and Digital Display IV*, 1992, pp. 202-216.
- [37] A.B. Watson, G.Y. Yang, J.A. Solomon, and J. Villasenor, "Visibility of wavelet quantization noise," *IEEE Trans. Image Processing*, vol. 6, pp. 1164-1175, Aug. 1997.

- [38] A.G. Bors and I. Pitas, "Image watermarking using DCT domain constraints," in *IEEE Proc. Int. Conf. on Image Processing*, Lausanne, Switzerland, Sept. 1996, vol. 3, pp. 231-234.
- [39] A. Piva, M. Barni, F. Bartolini, and V. Cappellini, "DCT-based watermarking recovering without resorting to the uncorrupted original image," in *Proc. IEEE Int. Conf. Image Processing*, Santa Barbara, CA, 1997, vol. 1, pp. 520-523.
- [40] W. Zeng and B. Liu, "A statistical watermark detection technique without using original images for resolving rightful ownerships of digital images," *IEEE Trans. Image Processing*, vol. 8, no. 11, pp. 1534-1548, Nov. 1999.
- [41] J.O. Ruanaidh, W.J. Dowling, and F.M. Boland, "Phase watermarking of digital images," in *Proc. IEEE Int. Conf. Image Processing (ICIP)*, 1996, pp. 239-242.
- [42] P. Davern and M. Scott, "Fractal based image steganography," in *Lecture Notes in Computer Science: Information Hiding*, vol. 1174, pp. 279-294, 1996.
- [43] J. Puate and F. Jordan, "Using fractal compression scheme to embed a digital signature into an image," in *Proc. SPIE Photonics East*, Boston, MA, Nov. 1996, pp. 108-119.
- [44] M.J.J.B. Maes and C.W.A.M. van Overveld, "Digital watermarking by geometric warping," in *Proc. IEEE Int. Conf. Image Processing*, Chicago, IL, Oct. 1998, pp. 424-426.
- [45] E. Lin, C. Podilchuk, and E. Delp, "Detection of image alterations using semi-fragile watermarks," in *Proc. SPIE, Security and Watermarking of Multimedia Contents II*, San Jose, CA, Jan. 2000, vol. 3971, pp. 152-163.
- [46] J. Fridrich, "Image watermarking for tamper detection," in *Proc. IEEE Int. Conf. Image Processing*, Chicago, IL, Oct. 1998, pp. 404-408.
- [47] D. Kundur and D. Hatzinakos, "Digital watermarking for telltale tamper proofing and authentication," *Proc. IEEE*, vol. 7, pp. 1167-1180, July 1999.
- [48] J. Brassil, S. Low, N. Maxemchuk, and L. O'Gorman, "Electronic marking and identification techniques to discourage document copying," in *Proc. Infocom'94*, 1994, pp. 1278-1287.
- [49] J. Brassil, S. Low, N. Maxemchuk, and L. O'Gorman, "Electronic marking and identification techniques to discourage document copying," *IEEE J. Select. Areas Commun.*, vol. 13, pp. 1495-1504, 1995.
- [50] J. Brassil, S. Low, and N. Maxemchuk, "Copyright protection for the electronic distribution of text documents," *Proc. IEEE*, no. 7, pp. 1181-1196, July 1999.
- [51] R. Ohbuchi, H. Masuda, and M. Aono, "Watermarking three-dimensional polygonal models through geometric and topological modifications," *IEEE J. Select. Areas Commun.*, vol. 16, pp. 551-560, May 1998.
- [52] F. Hartung, P. Eistert, and B. Girod, "Digital watermarking of mpeg-4 facial animation parameters," *Comput. Graph.*, vol. 22, no. 4, pp. 425-435, Aug. 1998.
- [53] J. Bloom, I. Cox, T. Kalker, J.P. Linnartz, M. Miller, and C. Traw, "Copy protection for dvd video," *Proc. IEEE*, no. 87, pp. 12667-1276, July, 1999.
- [54] Data Hiding Subgroup; <http://www.dvcc.comdhsq>.
- [55] F. Hartung and B. Girod, "Watermarking of uncompressed and compressed video," *Signal Process.*, vol. 66, no. 3, pp. 283-301, May 1998.
- [56] F. Hartung, "Digital watermarking and fingerprinting of uncompressed and compressed video," Ph.D. dissertation, Shaker Verlag, Univ. of Erlangen-Nuremberg, Aachen, Germany, Oct. 1999.
- [57] C. Langelaar, R. Legendijk, and J. Biemond, "Realtime labeling of mpeg-2 compressed video," *J. Vis. Commun. Image Represent.*, vol. 9, no. 4, pp. 256-270, Dec. 1998.
- [58] G. C. Langelaar, "Real-time watermarking techniques for compressed video data," Ph.D. thesis, Delft Univ. of Technology, 2000.
- [59] T. Kalker, G. Depovere, J. Haitsma, and M. Maes, "A video watermarking system for broadcast monitoring," in *Proc. SPIE, Security and Watermarking of Multimedia Contents*, San Jose, CA, Jan. 1999, vol. 3657, pp. 103-112.
- [60] F. Jordan, M. Kutter, and T. Ebrahimi, "Proposal of a watermarking technique for hiding/retrieving data in compressed and decompressed video," in *ISO/IEC Doc. JTC1/SC29/WG11/MPEG97/M2281*, July 1997.
- [61] J. Dittmann, M. Stabenau, and R. Steinmetz, "Robust mpeg video watermarking technologies," in *Proc. ACM Multimedia*, Sept. 1998, pp. 71-80.
- [62] B. Macq and J.-J. Quisquater, "Cryptology for digital tv broadcasting," *Proc. IEEE*, vol. 83, pp. 944-957, 1995.
- [63] L. Boney, A. Tewfik, and K. Hamdy, "Digital watermarks for audio signals," in *IEEE Proc. Multimedia*, 1996, pp. 473-480.
- [64] N.D. Jayant, J.D. Johnston, and R.J. Safranek, "Signal compression based on models of human perception," *Proc. IEEE*, vol. 81, pp. 1385-1422, Oct. 1993.
- [65] I. J. Cox, J. Kilian, T. Leighton, and T. Shamoon, "Secure spread spectrum watermarking for images, audio and video," in *IEEE Proc. Int. Conf. Image Processing*, Lausanne, Switzerland, 1996, pp. 243-246.
- [66] S. Craver, N. Memon, B.-L. Yeo, and M. Yeung, "Can invisible watermarks resolve rightful ownerships?" in *IS&T, SPIE Electronic Imaging'97: Storage and Retrieval of Image and Video Databases*, 1997, pp. 310-323.
- [67] B. Chen and G. Wornell, "Dither modulation: A new approach to digital watermarking and information embedding," in *Proc. SPIE, Security and Watermarking of Multimedia Contents*, San Jose, CA, Jan. 1999, vol. 3657, pp. 342-353.
- [68] B. Chen and G. Wornell, "Preprocessed and postprocessed quantization index modulation methods for digital watermarking," in *Proc. SPIE, Security and Watermarking of Multimedia Contents II*, San Jose, CA, 2000, vol. 3971, pp. 48-59.
- [69] S. Seretto, C. Podilchuk, and K. Ramchandran, "Capacity issues in digital image watermarking," in *Proc. IEEE Int. Conf. Image Processing (ICIP)*, 1998, pp. 445-449.
- [70] T.M. Cover and J.A. Thomas, *Elements of Information Theory*. New York: Wiley, 1991.
- [71] I. Cox and J.-P. Linnartz, "Some general methods for tampering with watermarks," *IEEE J. Select. Areas Commun.*, vol. 16, pp. 587-593, May 1998.
- [72] Petitcolas, F.A.P., Stirmark, www.cl.cam.ac.uk.
- [73] L. Qiao and K. Nahrstedt, "Watermarking schemes and protocols for protecting rightful ownership and customer's rights," *J. Vis. Commun. Image Represent.*, vol. 9, pp. 194-210, Sept. 1998.
- [74] R. Wolfgang and E. Delp, "A watermarking technique for digital imagery: further studies," in *Proc. Int. Conf. Imaging Science, Systems, and Applications*, Las Vegas, NV, June 1997, pp. 279-287.
- [75] M. Kutter, "Digital image watermarking: hiding information in images," Ph.D. dissertation, Ecole Polytechnique Federale De Lausanne, 1999.
- [76] M. Kutter, "Watermarking resisting to translation, rotation and scaling," in *Proc. SPIE*, vol. 3528, pp. 423-431, Nov. 1998.
- [77] J. O'Ruanaidh and T. Pun, "Rotation, scale and translation invariant digital image watermarking," in *Proc. IEEE Int. Conf. Image Processing*, Santa Barbara, CA, Oct. 1997, vol. 1, pp. 536-539.
- [78] P. Moulin and J. O'Sullivan, "Information-theoretic analysis of information hiding," *IEEE Trans. Inform. Theory*, to be published.
- [79] P. Moulin and J. O'Sullivan, "Information-theoretic analysis of watermarking," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing*, Istanbul, Turkey, June 2000, pp. 161-164.
- [80] J.K. Su and B. Girod, "Power-spectrum condition for energy-efficient watermarking," in *Proc. IEEE Int. Conf. Image Processing (ICIP)*, Kobe, Japan, Oct. 1999, pp. 301-305.
- [81] J. Su and B. Girod, "Fundamental performance limits of power-spectrum condition-compliant watermarks," in *Proc. SPIE, Security and Watermarking of Multimedia Contents II*, San Jose CA, Jan. 2000, vol. 3971, pp. 314-325.
- [82] J. Eggers and B. Girod, "Quantization watermarking," in *Proc. SPIE, Security and Watermarking of Multimedia Contents II*, San Jose, CA., Jan. 2000, pp. 60-73.
- [83] T. Kalker, J.-P. Linnartz, G. Depovere, and M. Maes, "On the reliability of detecting electronic watermarks," in *Proc. European Signal Processing Conf. (EUSIPCO98)*, Sept. 1998, pp. 13-16.
- [84] J.-P. Linnartz, T. Kalker, and J. Haitsma, "Detecting electronic watermarks in digital video," in *Proc. IEEE Int. Conf. Acoustics, Speech and Signal Processing (ICASSP99)*, Apr. 1999, pp. 2071-2074.
- [85] J. Hernandez and F. Perez-Gonzalez, "Statistical analysis of watermarking schemes for copyright protection," *Proc. IEEE*, vol. 87, pp. 1142-1166, July 1999.