Optimal Differential Energy Watermarking of DCT Encoded Images and Video

Gerrit C. Langelaar and Reginald L. Lagendijk

Abstract—This paper proposes the differential energy watermarking (DEW) algorithm for JPEG/MPEG streams. The DEW algorithm embeds label bits by selectively discarding high frequency discrete cosine transform (DCT) coefficients in certain image regions. The performance of the proposed watermarking algorithm is evaluated by the robustness of the watermark, the size of the watermark, and the visual degradation the watermark introduces. These performance factors are controlled by three parameters, namely the maximal coarseness of the quantizer used in pre-encoding, the number of DCT blocks used to embed a single watermark bit, and the lowest DCT coefficient that we permit to be discarded. In this paper, we follow a rigorous approach to optimizing the performance and choosing the correct parameter settings by developing a statistical model for the watermarking algorithm. Using this model, we can derive the probability that a label bit cannot be embedded. The resulting model can be used, for instance, for maximizing the robustness against re-encoding and for selecting adequate error correcting codes for the label bit string.

Index Terms—Data hiding, discrete cosine transform, digital watermark, image coding, image communication, image content control.

I. INTRODUCTION

THE RAPID growth of digital media and communication networks has created an urgent need for self-contained data identification schemes to create adequate intellectual property right (IPR) protection technology in particular for image and video data. In addition to conventional identification solutions such as the insertion of visual logos into the image or video data, and protection of the data through scrambling or encryption of the imagery or bit streams, the recently introduced data labeling or watermarking technique is being considered as a viable alternative [1], [2], [13], [26], [29]. By embedding an invisible and robust watermark into the image or video data, unauthorized copies can be traced [6], [8], [15], [17], [28] and copy protection schemes can be implemented [11], [12], [14]. There are several approaches to embed a watermark into an image or video frame. First generation watermarking techniques typically embed a secret message or label bit string into an image via characteristic pseudorandom noise patterns. These noise patterns can be generated and added in the either spatial [4], [20], [26], Fourier [25], DCT [6], [7], [21], or wavelet domain [2],

[10]. Commonly the amplitudes of the noise patterns are made dependent on the local image content as so trade-off the perceptual image degradation due to the noise and the robustness of the embedded information against image processing-based attacks. Other approaches use a particular order of discrete cosine transform (DCT) coefficients to embed the watermark [9]. More recent approaches use salient geometric image properties such as isolated corner points [24] or the correspondence map between domain and range blocks in fractal compression algorithm [3], [22] to embed the watermark.

In this paper we propose a watermarking algorithm that is suitable for—but not limited to—real-time watermarking of JPEG or MPEG streams because it operates directly on DCT blocks. The advantage of our technique is that it avoids the need for decoding JPEG or MPEG encoded information yielding a lightweight watermarking process that is well suited for implementation in consumer products. The proposed technique is robust against attempts to remove the watermark by re-encoding the JPEG or MPEG encoded bit streams. Removal of the watermark can only be done by image-based processing operations, which requires full decoding and re-encoding of the watermarked image or video streams.

The proposed method is based on selectively discarding high frequency DCT coefficients in the compressed data stream. The information bits of the data identifier (label) are encoded in the pattern of DCT blocks in which high frequency DCT coefficients are removed, i.e., in a pattern of energy differences between DCT blocks. For this reason, we call our technique a differential energy watermark (DEW).

The performance of the proposed technique depends on three parameters. The first parameter is the number of 8×8 DCT blocks n that is used to embed a single information bit of the data identifier. The larger n is chosen, the more robust the watermark becomes against watermark-removal attacks, but the fewer information bits can be embedded into an image or a single frame of a video sequence.

The second parameter controls the robustness of the watermark against re-encoding attacks. In a re-encoding attack the watermarked image or video is partially or fully decoded and subsequently re-encoded at a lower bit rate. Our method anticipates the re-encoding at lower bit rates up to a certain minimal rate. Without loss of generality we will elaborate on the re-encoding of *JPEG* compressed images, in which case the anticipated re-encoding bit rate can be expressed by the JPEG quality factor setting Q_{jpeg}. The smaller Q_{jpeg} is the more robust the watermark becomes against re-encoding attacks. However, for decreasing Q_{jpeg} increasingly more (high to middle frequency)

DCT coefficients have to be removed upon embedding of the watermark, which leads to an increasing probability for artifacts to become visible due to the presence of the watermark.

The third parameter is the so-called minimal *cutoff* index c_{min}. This value represents the smallest index—in zigzag scanned fashion—of the DCT coefficient that is allowed to be removed from the image data upon embedding the watermark. The smaller c_{min} is chosen, the more robust the watermark becomes but at the same time, image degradations due to removing high frequency DCT coefficients may become apparent. For a given c_{min}, there is a certain probability that a label bit cannot be embedded. Consequently, sometimes a *random* information bit will be recovered upon watermark detection, which is denoted as a *label bit error* in this paper. Clearly, the objective is to make the probability for label bit errors as small as possible.

In order to optimize the performance of the proposed watermark technique, the above mentioned parameters have to be determined. In an earlier paper [12], we used experimentally determined settings for these parameters. For a given image and watermark this is, however, an elaborate process. In this paper, we will show that it is possible to derive an expression for the label bit error probability P_{be} as a function of the parameters Q_{jpeg} and n. The relations that we derive analytically describe the behavior of the watermarking algorithm, and they make it possible to select suitable values for the three parameters (n, Q_{jpeg}, c_{min}), as well as suitable error correcting codes for dealing with label bit errors.

In Section II, we first describe the basic concept of the DEW algorithm. Then, in Section III, we derive an analytical expression for the probability mass function (PMF) of the cutoff indices. In Section IV, this PMF is verified with real-world data. After deriving and validating the obtained PMF, we use the PMF to find the probability that a label string cannot be recovered correctly (Section V) and the optimal parameter settings (n, Q_{jpeg}, c_{min}) (Section VI). Subsequently, in Section VII, we experimentally validate the results from Section VI. The paper concludes with a discussion on the proposed watermarking technique and its optimization in Section VIII.

II. DEW ALGORITHM

The information that we wish to embed into the image or video frame is represented by the label bit string L consisting of label bits L_j ($j = 0, 2, \cdots, l - 1$). This label bit string is embedded bit-by-bit in a set of n 8×8 DCT blocks taken from a JPEG compressed still image or from an I-frame of an MPEG compressed video stream. For the purpose of simplicity of the discussion, we will refer to still images and MPEG I-frames as "image." In this paper we will assume that the image is already in compressed format, so that operating on 8×8 DCT blocks is a natural choice. In case the images are not DCT compressed, the DEW algorithm requires a block-based DCT transformation of the image data as a preprocessing step.

In order to obtain sufficient robustness, typically n takes on values between 16 and 64, which means that a single label bit is embedded in a region of the image. However, before the label bits are embedded, the positions of the 8×8 DCT blocks in the
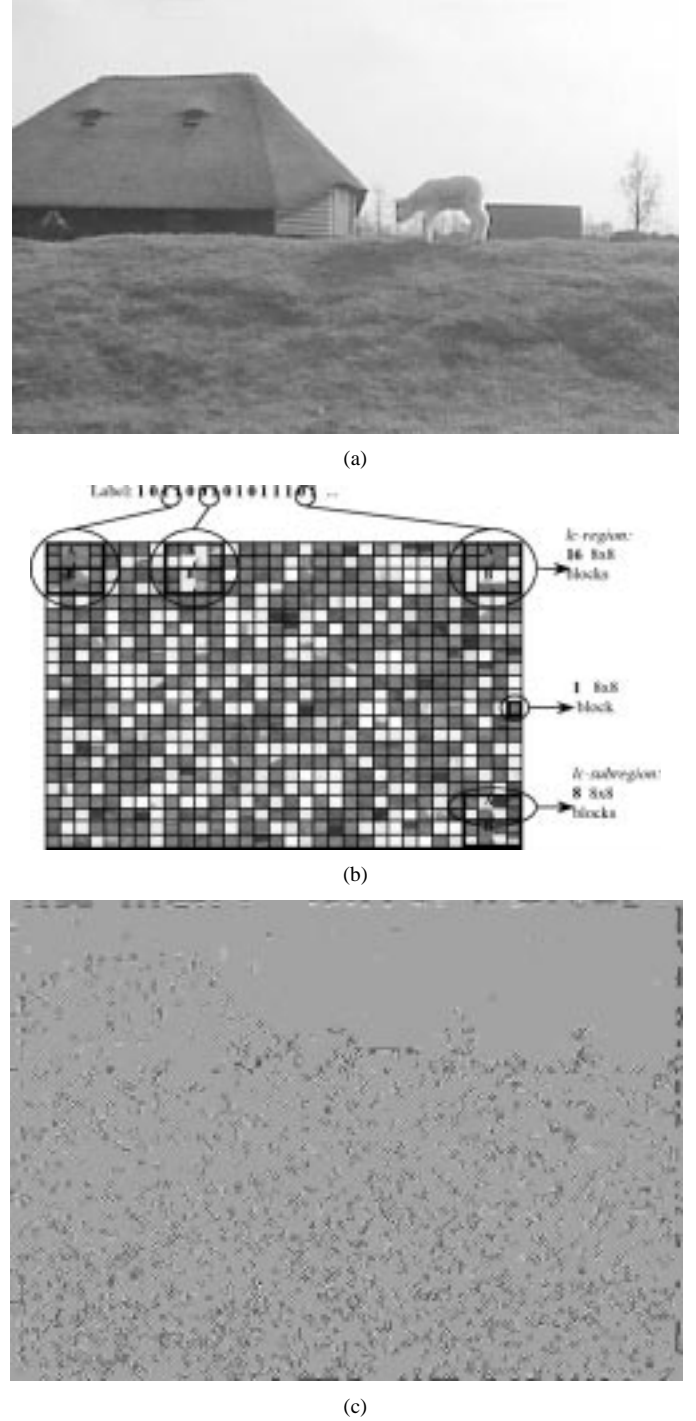

Fig. 1. (a) Sample I-frame, (b) block-based randomly shuffled I-frame showing the label-carrying (lc) regions and lc-subregions, and (c) difference between the original and watermarked image showing that the DEW algorithm put the watermark in regions with a lot of spatial details.

image are shuffled randomly as illustrated in Fig. 1. This shuffling operation on the one hand forms the secret key of the labeling algorithm, while on the other hand it spatially randomizes the statistics of DCT blocks. The latter observation implies that for the embedding process and for our analysis in the following sections, the shuffled image can be regarded approximately spatially stationary.

Each bit of the label bit string is embedded in its private label bit-carrying-region, or *lc-region* for short, in a shuffled image.

For instance, in Fig. 1 the first bit is located in the top-left-corner of the image in an lc-region of $n = 16$ DCT blocks. The value of the label bit is encoded by introducing an energy difference between the high frequency DCT-coefficients of the top half of the lc-region (denoted by *lc-subregion* A) containing in this case $n/2 = 8$ DCT blocks, and the bottom half (denoted by *lc-subregion* B) also containing $n/2 = 8$ DCT blocks. If the *lc-subregion* A contains more high frequency energy than the *lc-subregion* B, the label bit value "0" has been embedded into the data, and vice versa.

In order to make the determination of "high-frequency" energy easy for images or video frames that are JPEG or MPEG compressed, we compute energies over a subset of zigzag scanned DCT-coefficients indicated by $S(c)$

$$S(c) = \{i \in \{0, 63\}|(i > c)\}. \tag{1}$$

The zigzag scanned DCT coefficients are numbered according to Fig. 2. The index $i = 0$ refers to the DC-coefficient of a DCT block. The subset of DCT coefficients $S(c)$ over which energies are computed, is defined by the *cutoff index c*. The selection of a suitable cutoff index c for an lc-region is essential for the robustness and the visibility of the label bit. The larger the cutoff index is chosen, the less degradation the label embedding will introduce. For the moment we assume that we have available a suitable cutoff index c for each lc-region. Note that different lc-regions may have different cutoff indices depending on their spatial contents.

The (DCT high-frequency) *energy* E_A in lc-subregion A is now defined as follows:

$$E_A(c, n, Q_{jpeg}) = \sum_{b=0}^{n/2-1} \sum_{i \in S(c)} \left([\theta_{i,b}]_{Q_{jpeg}}\right)^2. \tag{2}$$

Here, $\theta_{i,b}$ denotes the nonweighted DCT coefficient with index i in the bth DCT block of the lc-subregion A under consideration. The notation $[]_{Q_{jpeg}}$ indicates that, prior to the calculation of E_A, the DCT-coefficients are re- or pre-quantized, in our case using the standard JPEG quantization procedure [18] with quality factor Q_{jpeg}. For embedding label bits into MPEG compressed I-frames a similar approach can be followed, but in this paper, we confine ourselves to the JPEG notation without loss of generality. The prequantization is done only in determining the cutoff indices and the calculation of (2), but is *not* applied to the actual image data upon embedding the label. The energy in lc-subregion B, denoted by E_B, is defined similarly.

We now define the energy difference D between the lc-subregions A and B as follows:

$$D(c, n, Q_{jpeg}) = E_A(c, n, Q_{jpeg}) - E_B(c, n, Q_{jpeg}). \tag{3}$$

The value of a label bit is encoded as the sign of the energy difference D. Label bit "0" is defined as $D > 0$ and label bit "1" as $D < 0$. The label *embedding* procedure must therefore adapt E_A and E_B to manipulate the energy difference D. If label bit "0," must be embedded, all energy after the cutoff index c in
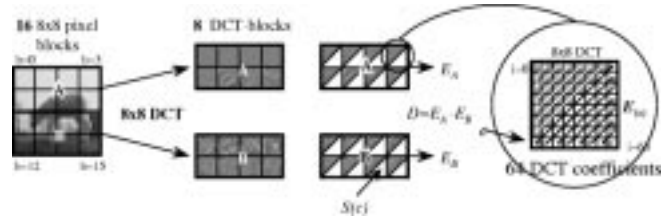

Fig. 2. Illustration of the calculation of the energy carried by high-frequency DCT coefficients by (2).

the DCT-blocks of lc-subregion B is eliminated by setting the corresponding DCT-coefficients to zero, yielding

$$D = E_A - E_B = E_A - 0 = +E_A. \tag{4}$$

If label bit "1" must be embedded, all energy after the cutoff index c in the DCT-blocks of lc-subregion A is eliminated, yielding $D = -E_B$. In case the watermark is embedded into a compressed image or video frame, i.e., the watermark is embedded in the compressed bit stream, the DCT coefficients can easily be forced to zero without re-encoding the bit stream by shifting the *end of block marker* (EOB) of 8×8 DCT blocks in one of the two lc-subregions toward the DC-coefficient, up to the selected cutoff index.

In Fig. 2, the complete procedure to calculate the energy difference D in an lc-region is illustrated for $n = 16$ nonshuffled 8×8 DCT blocks. The white triangularly shaped areas illustrate the subsets over which the energies are calculated for a particular choice of the cutoff index $c = 27$. At the right a blowup of one 8×8 DCT block is presented. In Fig. 1(c), the difference between the original and watermarked image is shown, illustrating that the DEW algorithm embeds information bits in those regions of the image that contain many details. Because of the pre-quantization with (JPEG) quality Q_{jpeg} in the calculation of the energy of the (high-frequency) DCT coefficients in (2), the DEW algorithm effectively embeds the label bits in perceptually important image details that are not significantly affected by JPEG/MPEG compression. Consequently, removing the DEW watermark is not possible without strongly affecting the perceptual image quality.

The selection of the cutoff index depends on the desired energy difference between the two lc-subregions. To determine the cutoff index c for an lc-region given a desired energy difference D, we first calculate the energies $E_A(c, n, Q_{jpeg})$ and $E_B(c, n, Q_{jpeg})$ for all possible cutoff indices $c = 0 \cdots 63$. Since the resulting optimal cutoff index varies per label bit that we wish to embed, it can be interpreted as a stochastic variable that depends on n, Q_{jpeg}, D, and $c_{\min}$, i.e., $C(n, Q_{jpeg}, D, c_{\min})$. If D is the energy difference that is needed to represent a label bit in an lc-region, the cutoff index is found as the *largest* index of the DCT coefficients for which (2) gives an energy *larger* than the required difference D in both subregions A and B.

Since the parameter D directly determines the number of DCT-coefficients that are discarded during labeling, it also determines the visibility and robustness of the label. In controlling the visual quality of the watermarked image, we wish to avoid the situation that the important low frequency DCT coefficients

are discarded. To this end, we require the selected cutoff index to always be larger than a certain minimum $c_{\min}$. Mathematically, this gives the following expression for determining C:

$$
\begin{aligned}
C&(n, Q_{jpeg}, D, c_{\min}) \\
&= \max\{c_{\min}, \max\{i \in \{0, 63\} | (E_A(i, n, Q_{jpeg}) > D) \\
&\qquad\qquad \cap (E_B(i, n, Q_{jpeg}) > D)\}\}. \quad (5)
\end{aligned}
$$

To extract a label bit from an lc-region we have to recover the cutoff index that was used for that lc-region during the embedding process. Upon label bit extraction, first the energies $E_A(c, n, Q_{jpeg})$ and $E_B(c, n, Q_{jpeg})$ are calculated for all possible cutoff indices $c = 0 \cdots 63$. Since either in lc-subregion A or lc-subregion B several DCT-coefficients have been eliminated during the watermark embedding, we first find the *smallest* index of the DCT coefficients for which (2) gives an energy smaller than a threshold D' in either of the two lc-subregions. The actually used cutoff index is then found as the maximum of these two numbers

$$
\begin{aligned}
C^{(\text{extract})}&(n, Q'_{jpeg}, D') \\
&= \max\{\min\{i \in \{0, 63\} | E_A(i, n, Q'_{jpeg}) < D'\}, \\
&\qquad\qquad \min\{i \in \{0, 63\} | E_B(i, n, Q'_{jpeg}) < D'\}\}.
\end{aligned}
$$
$$(6)$$

In the above procedure, the parameters Q'_{jpeg} and D' can be chosen equal to the parameters Q_{jpeg} and D, which are used in the embedding phase. The requantization step can also be omitted ($Q'_{jpeg} = 100$) without significantly influencing the reliability of the label bit extraction. Since Q_{jpeg} and D are not fixed parameters but may vary per image, the label extraction procedure must be able to determine suitable values for Q'_{jpeg} and D' itself. The most reliable way for doing this is to start the label bit string with several fixed label bits, so that during the label extraction those values for Q'_{jpeg} and D' can be chosen that result in the fewest errors in the known label bits.

III. MODELING THE DEW ALGORITHM

When operating the DEW algorithm, different values for the cutoff index are obtained. Insight in the actually selected cutoff indices is important since the cutoff indices used determine the quality and robustness of the DEW. Therefore, in this section we will derive the probability mass function (PMF) for the cutoff index based on a stochastic model for DCT coefficients. This PMF depends only on the parameters Q_{jpeg} and n. The model will be validated in Section IV, while in Section V we will use this PMF to obtain a function for the label bit error probability.

A. PMF of the Cutoff Index

In order to be able to compute the PMF of the cutoff index, we first assume that the energy difference D in (4) is chosen in the range $[1, D_{\max}(Q_{jpeg})]$. Here $D_{\max}(Q_{jpeg})$ indicates the maximum of the range of energies defined by (2) that *does not* occur in quantized DCT blocks because of the JPEG or MPEG compression process. Fig. 3 illustrates this effect by showing an histogram of the energy $E(c, n, Q_{jpeg})$ for a wide range of values of c, n, and Q_{jpeg}. We notice a clear "gap" in the

histogram for smaller energies, because DCT blocks with that small amount of energy can no longer exist after compression.

In general, the maximum $D_{\max}(Q_{jpeg})$ depends on how heavy the image has been compressed, i.e. it depends on Q_{jpeg}. The smaller Q_{jpeg} is, the larger $D_{\max}(Q_{jpeg})$ will be. Mathematically, this relation is given by

$$
D_{\max}(Q_{jpeg}) = \left(F(Q_{jpeg}) \min_i(W_i)\right)^2
$$
$$
F(Q_{jpeg}) = \begin{cases} 50/Q_{jpeg} & Q_{jpeg} < 50 \\ (100 - Q_{jpeg})/50 & Q_{jpeg} \geq 50 \end{cases} \quad (7)
$$

where $F(Q_{jpeg})$ denotes the coarseness of the quantizer used, and W_i is the ith element ($i \in [c_{\min}, 63]$) of the zigzag scanned standard JPEG luminance quantization table [18].

Theorem 1: If the enforced energy difference D is chosen in the range $[1, D_{\max}(Q_{jpeg})]$, where $D_{\max}(Q_{jpeg})$ is defined by (7), and if we do not constrain the cutoff index by $c_{\min}$, the PMF of the cutoff index is given by

$$
\begin{aligned}
P&[C(n, Q_{jpeg}) = c] \\
&= P[E(c, n, Q_{jpeg}) \neq 0]^2 - P[E(c + 1, n, Q_{jpeg}) \neq 0]^2
\end{aligned}
$$
$$(8)$$

where $E(c, n, Q_{jpeg})$ is defined in (2). Observe that in this theorem $C(n, Q_{jpeg})$—besides being not constrained by $c_{\min}$—is no longer dependent on D due to the wide range of values in which D can be selected.

Proof: We first define the set

$$
\begin{aligned}
\Psi&(c, n, Q_{jpeg}) \\
&= \{A, B | E_A(c, n, Q_{jpeg}) > D \cap E_B(c, n, Q_{jpeg}) > D\}.
\end{aligned}
$$
$$(9)$$

As for fixed n and Q_{jpeg} the following relation holds:

$$
E(c, n, Q_{jpeg}) \geq E(c+1, n, Q_{jpeg}) \geq E(c+2, n, Q_{jpeg}) \geq \cdots
$$
$$(10)$$

we clearly have

$$
\Psi(0, n, Q_{jpeg}) \supseteq \Psi(1, n, Q_{jpeg}) \supseteq \cdots \supseteq \Psi(63, n, Q_{jpeg}).
$$
$$(11)$$

Furthermore, we easily see that for fixed n and Q_{jpeg}

$$
\{A, B | C(n, Q_{jpeg}) = c\} = \Psi(c, n, Q_{jpeg}) \backslash \Psi(c+1, n, Q_{jpeg}).
$$
$$(12)$$

Therefore, we can calculate the following probability:

$$
\begin{aligned}
P&[C(n, Q_{jpeg}) = c] \\
&= P[\Psi(c, n, Q_{jpeg})] - P[\Psi(c + 1, n, Q_{jpeg})] \\
&= P[E_A(c, n, Q_{jpeg}) > D] . P[E_B(c, n, Q_{jpeg}) > D] \\
&\quad - P[E_A(c + 1, n, Q_{jpeg}) > D] \\
&\quad \cdot P[E_B(c + 1, n, Q_{jpeg}) > D].
\end{aligned}
$$
$$(13)$$

Here, we have to assume that $E_A(c, n, Q_{jpeg})$ and $E_B(c, n, Q_{jpeg})$ are mutually independent due the random shuffling of the positions of the DCT blocks. Furthermore, since the lc-subregions are both built-up from block-shuffled

image data, we can assume that the probabilities in (13) do not depend on the actual lc-subregion for which they are calculated, yielding

$$
\begin{aligned}
P[C(n,\, &Q_{jpeg}) = c] \\
&= P[E(c,\, n,\, Q_{jpeg}) > D]^2 - P[E(c+1,\, n,\, Q_{jpeg}) > D]^2.
\end{aligned}
$$
(14)

To calculate (14), we need to have an expression for probabilities of the form $P[E(c,\, n,\, Q_{jpeg}) > D]$. As illustrated by Fig. 3, the histogram of $E(c,\, n,\, Q_{jpeg})$ is zero for small $E(c,\, n,\, Q_{jpeg})$'s because the quantization process maps many small DCT coefficients to zero. As a consequence, the energy defined in (2) is either equal to 0 (for instance for large values of c), or the energy has a value larger than the smallest nonzero squared *quantized* DCT coefficient in the lc-subregion under consideration. This value has been defined as $D_{\max}(Q_{jpeg})$ in (7). Since we always choose the value of D smaller than $D_{\max}(Q_{jpeg})$, probabilities of the form $P[E(c,\, n,\, Q_{jpeg}) > D]$ can be simplified to $P[E(c,\, n,\, Q_{jpeg}) > D] = P[E(c,\, n,\, Q_{jpeg}) \neq 0]$. Substitution of this relation into (14) yields (9).

B. Model for the DCT-Based Energies

Theorem 2: If the probability density function (PDF) of the DCT coefficients is modeled as a generalized Gaussian distribution with shape parameter γ, then the probability that the energy $E_A(c,\, n,\, Q_{jpeg})$ is not equal to zero is given by

$$
\begin{aligned}
P[E(c,\, &n,\, Q_{jpeg}) \neq 0] \\
&= 1 - \left[\prod_{i=c}^{63} \left\{ 1 - e^{-(b_i Q_i)^\gamma} \cdot \left(\sum_{h=0}^{\gamma^{-1}-1} \frac{(b_i Q_i)^{h \cdot \gamma}}{h!} \right) \right\} \right]^{(n/2)}
\end{aligned}
$$
(15)

where

$$
\gamma^{-1} = 1,\, 2,\, 3,\, \cdots
$$
(16a)

$$
b_i Q_i = \frac{W_i F(Q_{jpeg})}{2\sigma_i} \sqrt{\frac{(3 \cdot \gamma^{-1} - 1)!}{(\gamma^{-1} - 1)!}}.
$$
(16b)

Further, $F(Q_{jpeg})$ denotes the coarseness of the quantizer as defined in (7), σ_i^2 represents the variance of the ith DCT-coefficient (in zigzag scanned fashion), and W_i represents the corresponding element of standard JPEG luminance quantization table.

Proof: The expression for $P[E_A(c) \neq 0]$ can be derived using (2). To this end, we first need a probability model for the DCT coefficients θ_i. Following literature at this point, we use the generalized Gaussian distribution [16], [27] with shape parameter γ

$$
p[\theta_i] = a_i e^{-|b_i \cdot \theta_i|^\gamma}
$$
(17a)

where
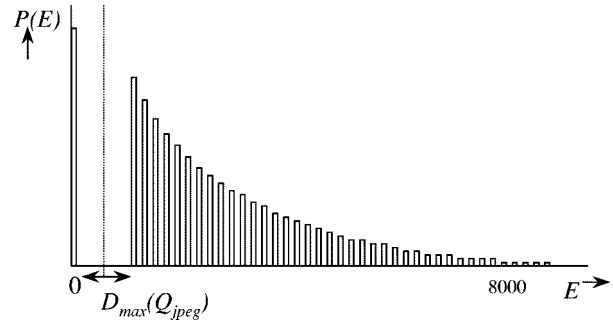
$$
a_i = \frac{b_i \cdot \gamma}{2(\gamma^{-1} - 1)!}
$$

Fig. 3. Histogram of the energy carried by high-frequency DCT coefficients—calculated using (2)—for a wide range of parameters $(c,\, n,\, Q_{jpeg})$.

and

$$
b_i = \frac{1}{\sigma_i} \sqrt{\frac{(3 \cdot \gamma^{-1} - 1)!}{(\gamma^{-1} - 1)!}} \quad \text{for} \quad \gamma^{-1} = 1,\, 2,\, 3,\, \cdots.
$$
(17b)

This PDF has zero-mean and variance σ_i^2. Typically, the shape parameter γ takes on values between 0.10 and 0.50. In a more complicated model, the shape parameter could be made dependent on the index of the DCT coefficient. We will, however, use a constant shape parameter for all DCT coefficients. Using (17), we can now calculate the probability that a DCT coefficient is quantized as zero

$$
\begin{aligned}
P[\hat{\theta}_i = 0] &= \int_{-Q_i}^{Q_i} a_i \cdot e^{-|b_i \cdot \theta_i|^\gamma} \, d\theta_i \\
&= 1 - e^{-(b_i Q_i)^\gamma} \cdot \left(\sum_{h=0}^{\gamma^{-1}-1} \frac{(b_i Q_i)^{h \cdot \gamma}}{h!} \right)
\end{aligned}
$$
(18)

where Q_i is the coarseness of the quantizer applied to the DCT coefficients. The probability that $E_A(c,\, n,\, Q_{jpeg})$ is equal to zero is now given by the probability that all quantized DCT coefficients with index larger than c in all $n/2$ DCT blocks are equal to zero

$$
P[E(c) = 0] = \left[\prod_{i=c}^{63} P(\hat{\theta}_i = 0) \right]^{n/2}.
$$
(19)

Equations (18) and (19) use the quantizer parameter Q_i. In JPEG, this parameter is determined by the parameter W_i and the function $F(.)$ that depends on the user parameter Q_{jpeg} via (7). Taking into account that JPEG implements quantization through rounding operations yields

$$
Q_i = 1/2 \, W_i \, F(Q_{jpeg}).
$$
(20)

Combining (17)–(20) yields (15).

IV. MODEL VALIDATION WITH REAL-WORLD DATA

We validate Theorem 1 as follows. From a wide range of differently textured images we calculated the normalized histogram of $P[E(c,\, n,\, Q_{jpeg}) \neq 0]$ as a function of c. As an example we show here the situation of $Q_{jpeg} = 50$ and $n = 16$. Using this histogram, (8) is evaluated to get an estimate of the
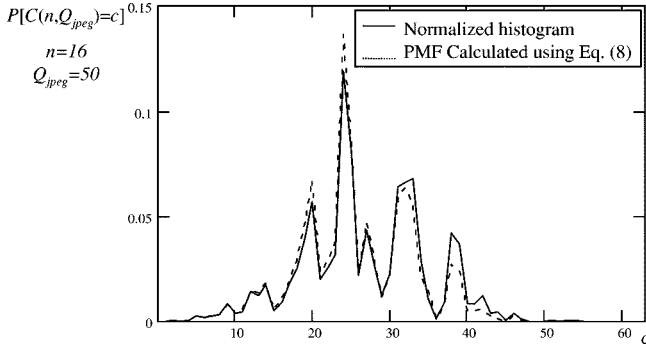
Fig. 4. Probability mass function of the cutoff index $P[C(n, Q_{jpeg}) = c]$ as a function of c, calculated as a normalized histogram directly from watermarked images (solid line), and calculated using the derived in Theorem 1 (8) (dotted line).
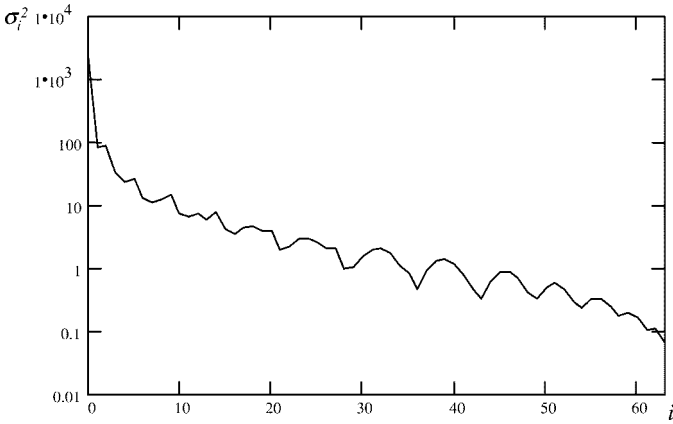

Fig. 5. Measured variances of the (unquantized) DCT-coefficients as a function of the coefficient number along the zigzag scan.
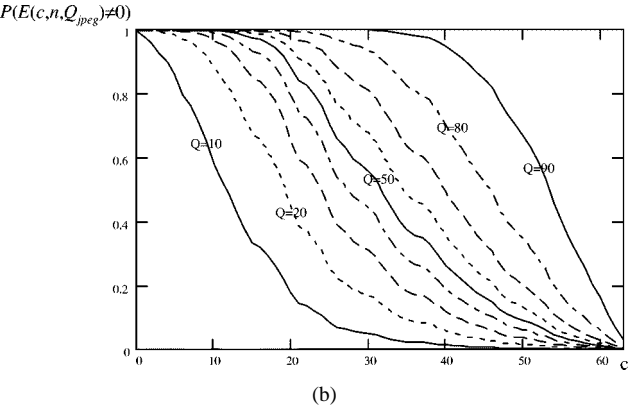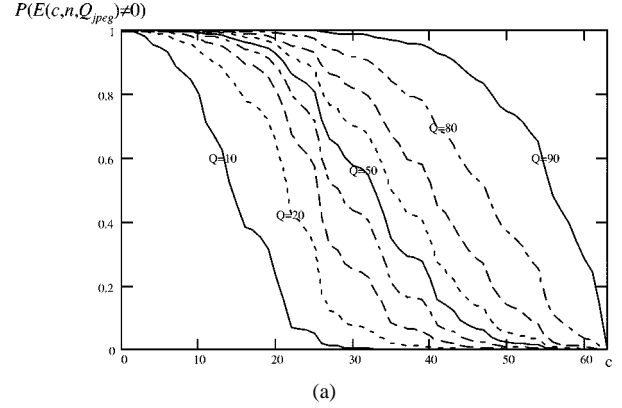

(a)

(b)

Fig. 6. Probability $P(E(c, n, Q_{jpeg}) \neq 0)$ as a function of c (a) calculated as a normalized histogram directly from watermarked images ($n = 16$) and (b) calculated using Theorem 2 ($n = 16$).

PMF $P[C(n, Q_{jpeg}) = c]$. The resulting PMF is shown in Fig. 4 as the dotted line. Using the same test data, we then directly calculated the histogram of $P[C(n, Q_{jpeg}) = c]$ as a function of c. The resulting (normalized) histogram is shown in Fig. 4 as the solid line. It shows that both curves fit well, which validates the correctness of the assumptions made in the derivation of Theorem 1.

For the validation of Theorem 2, we first need a reasonable estimate of the shape parameter γ and the variance σ_i^2 of the DCT coefficients. The shape parameter can be estimated in two different ways, namely 1) *a priori*, by statistical testing using the Kolmogorov–Smirnov test and 2) *a posteriori*, by fitting the theoretically calculated probabilities with the measured curves using experimental data. In fitting the PDF of the DCT coefficient we concentrated on obtaining a correct fit for the more important low frequency DCT coefficients, and obtained $\gamma = 1/7$. The variances of the DCT coefficients were measured over a large set of images, yielding Fig. 5. For the time being, we will use these experimentally determined variances, but later we will replace these with a fitted polynomial function.

In Fig. 6(a), normalized histograms of the energy $E(c, n, Q_{jpeg}) \neq 0$ are plotted for $n = 16$ and several values of Q_{jpeg} as a function of c. In Fig. 6(b) the probabilities $P[E(c, n, Q_{jpeg}) \neq 0]$ are shown as calculated with

(15) from Theorem 2 using the measured variances of the DCT-coefficients. Comparing the Fig. 6(a) and (b), we see that the estimated and calculated probabilities match quite well. There are some minor deviations for very small values of $Q_{jpeg}(Q_{jpeg} < 15)$, which is the result of the imperfect model for the DCT coefficients of real image data. We consider these deviations insignificant since they occur only at very high image compression factors. We conclude that the models underlying Theorem 2 give results for $P[E(c, n, Q_{jpeg}) \neq 0]$ that are sufficiently close to the actually observed data.

By combining Theorems 1 and 2, we can derive PMFs of the cutoff index as a function of the parameters n and Q_{jpeg} based merely on the variances of the DCT coefficients. To validate the combined theorems we compared the PMFs calculated using (8) and (15) with the normalized histograms directly calculated on a wide range of differently textured images. In Fig. 7, two examples of the PMFs are plotted. In these examples, the solid lines represent the normalized histograms of $C(n, Q_{jpeg})$ calculated from watermarked image data, while the dotted lines represent the PMF $P[C(n, Q_{jpeg}) = c]$ calculated using (8) and (15). The highly varying behavior of these curves as a function of c is mainly due to the zigzag scanning order of the DCT coefficients. We observe that an acceptable fit between the two curves is obtained with some deviations for higher cutoff indices. Since the PMF $P[C(n, Q_{jpeg}) = c]$ will be used for calculating the probability of a label bit error, i.e., the probability that the watermarking procedure attempts to select a cutoff index smaller than the minimum allowed values $c_{\min}$, slight deviations at higher
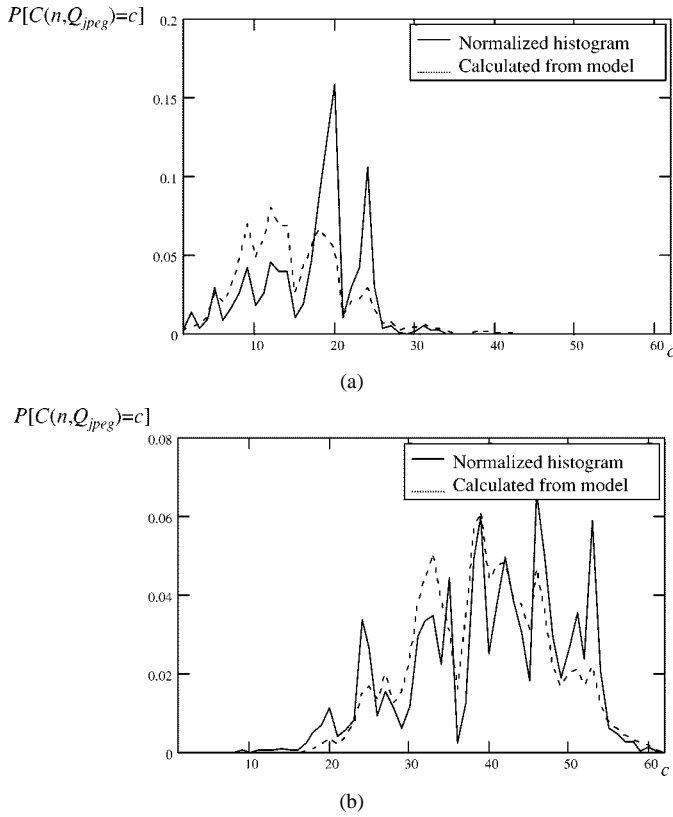
(a)

(b)

Fig. 7. Probability mass function of $C(n, Q_{jpeg})$, calculated as the normalized histogram directly from watermarked image data (solid line), and calculated using (8) and (15). The watermarking parameters in (a) are $n = 16$ and $Q_{jpeg} = 20$ and in (b) are $n = 16$ and $Q_{jpeg} = 80$.
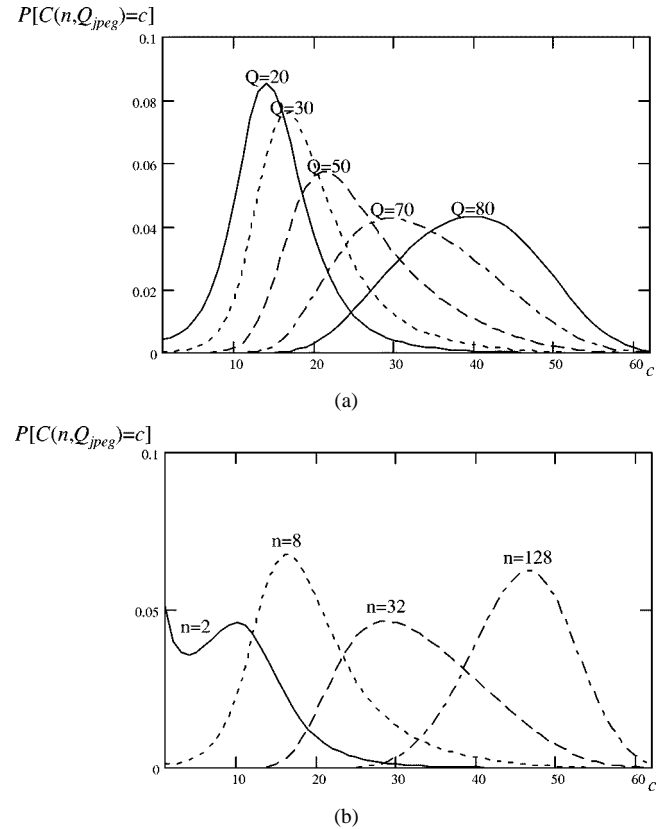

(a)

(b)

Fig. 8. (a) Analytically calculated PMF $P[C(n, Q_{jpeg}) = c]$ using Theorems 1 and 2 for various values of Q_{jpeg} and $n = 16$ and (b) analytically calculated PMF $P[C(n, Q_{jpeg}) = c]$ using Theorems 1 and 2 for various values of n and $Q_{jpeg} = 50$. These curves are computed using only the shape parameter γ and the fitting parameters of the DCT variances as input.

values for the cutoff index are not relevant to the objectives of this paper.

The final step is to use the relation (8) and (15) to *analytically* estimate the PMF $P[C(n, Q_{jpeg}) = c]$ of the cutoff index for different values of the parameters Q_{jpeg} and n. In this final step we rid ourselves of the erratic behavior of the curves in Figs. 5 and 7 due to the zigzag scan order of the DCT coefficients by approximating the variances of the DCT coefficients in Fig. 5 by a second order polynomial function. The overall effect of using a polynomial function for the DCT coefficients is the smoothing of the PMF $P[C(n, Q_{jpeg}) = c]$.

In Fig. 8, the analytically calculated PMF's are shown. These curves are computed using Theorems 1 and 2 with only the shape parameter γ and the fitting parameters of the DCT variances as input. In Fig. 8(a), $P[C(n, Q_{jpeg}) = c]$ is shown as a function of Q_{jpeg} keeping n constant, and in Fig. 8(b) $P[C(n, Q_{jpeg}) = c]$ is shown as a function of n keeping Q_{jpeg} constant. It can clearly be seen that decreasing n or Q_{jpeg} leads to an increased probability of lower cutoff indices. This complies with our earlier experiments in [12], which showed that watermarks embedded with small values for n or Q_{jpeg} yields visible artifacts due to the removal of high-frequency DCT coefficients.

V. LABEL ERROR PROBABILITY

In the analysis of the DEW algorithm, we have seen that depending on the parameter settings (n, Q_{jpeg}) certain cutoff in-

dices are more likely than others. In this analysis, however, the selection of the cutoff index by the watermarking algorithm has been carried out irrespective of the visual impact on the image data. In order for the watermark to remain invisible, the cutoff indices are constrained to be larger than a certain minimum $c_{\min}$. Consequently, it may happen in certain lc-regions that a label bit cannot be embedded. This random event is typically the case in lc-(sub)regions that contain insufficient high-frequency details.

Using Theorems 1 and 2, we are able to derive the probability that this undesirable situation occurs, and obtain an expression for the label bit error probability P_{be} that depends on Q_{jpeg}, n and $c_{\min}$. If a label bit cannot be embedded because of the minimally required value of the cutoff index $c_{\min}$, there is a probability of 0.5 that during the extraction phase a random bit is extracted which equals the original label bit. We assume that due to the random shuffling of DCT blocks, the occurrence of a label bit error can be considered as a random event, independent of other label bit errors. The probability that a random error occurs in a label bit, can therefore be computed as follows:

$$P_{be}(n, Q_{jpeg}, c_{\min})$$
$$= 0.5\, P[C(n, Q_{jpeg}) < c_{\min}] = 0.5 \sum_{c=0}^{c_{\min}} P[C(n, Q_{jpeg}) = c].$$

$$(21)$$

Using this relation, we can calculate the label bit error probability for each value of $c_{\min}$ as a function of Q_{jpeg} and n.

Using the label *bit* error probability in (21), we can now derive the *label* error probability P_e, which is here defined as the probability that one or more label *bit* errors occur in the embedded information bit string. Assuming image dimensions of $N_1 \times N_2$, the number of information bits l that the image can contain is given by

$$l(N_1, N_2, n) = \left\lfloor \frac{N_1 \cdot N_2}{64 \cdot n} \right\rfloor \qquad (22a)$$

with which the label error probability can be calculated as

$$P_e(n, Q_{jpeg}, c_{\min}, N_1, N_2) = 1 - (1 - P_{be})^{l(N_1, N_2, n)}. \qquad (22b)$$

Let us consider one particular numerical example. If, for instance in a broadcast scenario, one incorrect label is accepted per month in a continuous 10 Mbit/s video stream, the label bit error rate should be smaller than 10^{-7}. To select the optimal setting for Q_{jpeg} and n that comply with this label bit error rate, Fig. 9 shows curves of the combinations Q_{jpeg} and n for which P_e equals 10^{-7}. Different curves refer to different values of $c_{\min}$. Further we have assumed the image dimensions $N_1 \times N_2 = 1024 \times 768$.

VI. OPTIMAL PARAMETER SETTING

Using results such as the ones shown in Fig. 9, we can now select optimal settings for Q_{jpeg} and n for specific situations. We consider three different cases:

- optimization for re-encoding robustness, number of information bits l, and watermark invisibility;
- optimization for number of information bits l, and watermark invisibility;
- optimization for watermark invisibility.

In all cases, the parameter D must be chosen in the range $[1, D_{\max}(Q_{jpeg})]$ in order for the models in Theorems 1 and 2 and the analytical results obtained from these results, to be valid.

If we tune the DEW watermark such that it tradesoff the reencoding robustness, number of information bits l, and watermark invisibility, typical choices are to anticipate reencoding up to JPEG quality factor of $Q_{jpeg} = 25$, and to allow a minimal cutoff index of $c_{\min} = 3$. In this case—using Fig. 9—we need at least $n = 54$ DCT blocks per label bit (which directly determines the number of information bits that can be stored in an image) to achieve the required label error probability of 10^{-7}.

If we require a large label but robustness against re-encoding attacks is not an issue, we can store more than three times as many bits in a label with the same label error probability of 10^{-7}. A typical parameter setting would for instance be $Q_{jpeg} = 75$, $n = 16$ and $c_{\min} = 3$, as can be seen from Fig. 9.

If visual quality is the most important factor, we need to take the minimal cutoff index sufficiently large. For instance we choose $c_{\min} = 15$. Clearly, to obtain the same label bit error probability more DCT blocks per label bit are required since the allowed minimal cutoff index is larger than in the previous
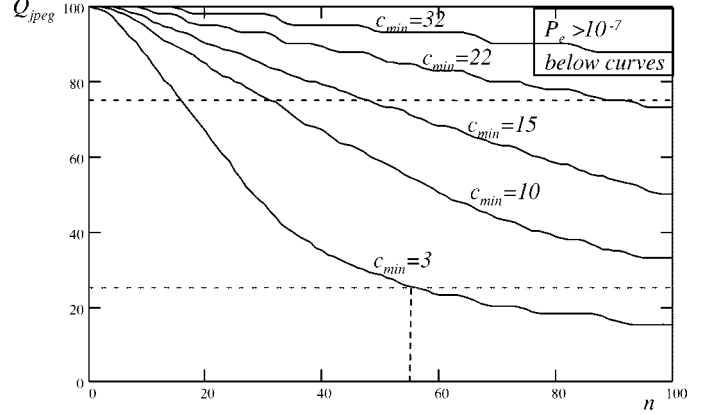

Fig. 9. Combinations of Q_{jpeg} and n for which $P_e = 10^{-7}$.

example. Using Fig. 9, we find as optimal settings in this case $Q_{jpeg} = 75$ and $n = 48$.

The performance of any watermarking system can be improved by applying error-correcting codes (ECCs). Since we know that the label bit errors occur randomly and independently of other label bit errors, we can compute the probability for *label error* in case an ECC is used that can correct up to M label bit errors, namely

$$P_e^{\text{ECC}(M)}(n, Q_{jpeg}, c_{\min}, N_1, N_2)$$
$$= 1 - \sum_{j=0}^{M} \binom{l(N_1, N_2, n)}{j} P_{be}^j (1 - P_{be})^{l(N_1, N_2, n) - j} \qquad (23)$$

with the label bit error probability P_{be} given by (21).

In Fig. 10, the label error probability $P_e^{ECC(M)}$ is shown as a function of the number of DCT blocks used to embed a single label bit (n) for $M = 0, 1, 2$, $Q_{jpeg} = 25$ and $c_{\min} = 3$. We had already found that for a watermark optimized for robustness without error correcting codes, the optimal value of $n = 54$ for a required bit error probability of $P_e < 10^{-7}$. From Fig. 10 we see that the same label error probability can be obtained using smaller values of n if we apply error correcting codes For instance, by using an ECC that can correct one error, n can be decreased from 54 to 33. Obviously the use of ECC's introduces some redundant bits. This overhead is however small compared to the increase in capacity due to the use of a smaller value of n. Table I gives some examples of the effective length of labels that can be embedded for $N_1 \times N_2 = 1024 \times 768$. In this table, standard BCH codes [23] are used that can correct one or two errors.

VII. EXPERIMENTAL RESULTS

In our earlier work on the DEW algorithm [12], we heuristically chose the following suitable parameters: $c_{\min} = 3$, $n = 16$, $Q_{jpeg} = 75$ and $D = 40$. Through the modeling and analysis described in this paper, we can now conclude that these settings are optimal for maximum label size and not for re-encoding robustness. For that reason, we will here compare the robustness of labels embedded using these settings with labels em-
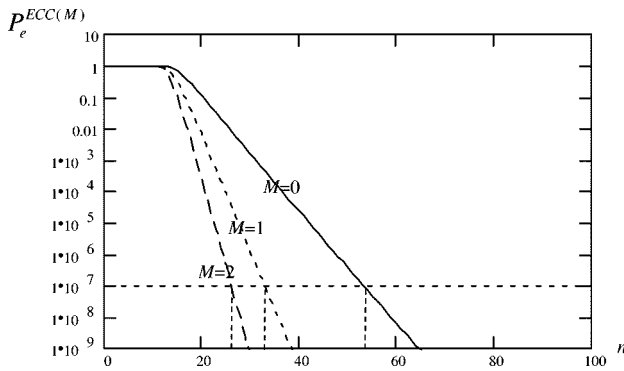
Fig. 10.   Label error probability with and without error correcting codes for $Q_{jpeg} = 25$ and $c_{\min} = 3$.

TABLE I
EFFECTIVE NUMBER OF BITS PER LABEL
THAT CAN BE EMBEDDED INTO AN IMAGE OF SIZE $N_1 \times N_2 = 1024 \times 768$,
WITH REQUIRED PERFORMANCE PARAMETERS $c_{\min} = 3$, $Q_{jpeg} = 25$
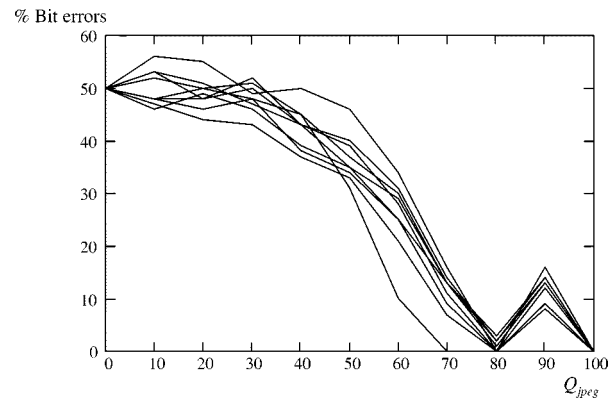AND $P_e^{\mathrm{ECC}(M)} < 10^{-7}$

ECC-Type	M	n	Parity-check bits ECC	Label size corrected for extra parity-check bits
no-ECC	0	54	0	227
BCH (511,502)	1	33	9	363
BCH (511,493)	2	27	18	437

(a)

(b)

Fig. 11.   Percentage bit errors after re-encoding (a) using parameter settings optimized for label size and (b) parameter settings optimized for robustness.

bedded using settings optimized for robustness, namely $c_{\min} = 3$, $n = 64$,[1] $Q_{jpeg} = 25$, and $D = 500$.
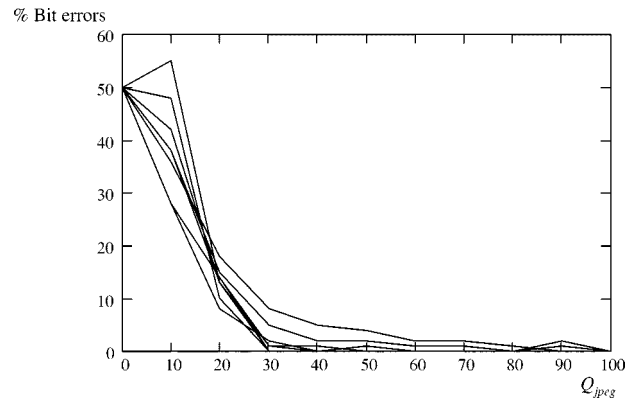
We will first check the robustness against re-encoding. Images are JPEG compressed with quality factor of 100. From these JPEG compressed images two watermarked version are produced, one for each parameter setting. Next, the images are reencoded using a lower JPEG quality factor. The quality factor of the reencoding process is made variable. Finally, the watermark is extracted from the reencoded images and bit-by-bit compared against the originally watermark. From this experiment, we find the percentages of label bit errors due to re-encoding as a function of the re-encoding quality factor. In Fig. 11, the resulting label bit error curves are shown for nine different images.

Comparing Fig. 11(a) (parameter setting optimized for label length using $c_{\min} = 3$, $n = 16$, $Q_{jpeg} = 75$, and $D = 40$) and Fig. 11(b) (parameter setting optimized for label robustness using $c_{\min} = 3$, $n = 64$, $Q_{jpeg} = 25$, and $D = 500$), we see an enormous gain in robustness. In Fig. 11(b), we see a breakpoint around $Q_{jpeg} = 25$. For higher re-encoding qualities, the percentage label bit errors is below 10%.

In [12] we noticed that the DEW watermarking technique is slightly resistant to line shifting. To investigate the effect of the

parameter settings optimized for robustness on the resistance to line shifting, we carry out the following experiment. Images are JPEG compressed with a quality factor of 85. These JPEG images are watermarked using the parameter settings optimized for label size or optimized for robustness. Next the images are decompressed, shifted to the right over R pixels and re-encoded using the same JPEG quality factor. Finally, a watermark is extracted from these re-encoded images and bit-by-bit compared with the originally embedded watermark. Consequently, we find the percentages bit errors due to line shifting. In Fig. 12, the bit error curves are shown for nine different images. As in the previous experiment, we see an improvement in robustness between Fig. 12(a) and (b). Using the parameter settings optimized for robustness, the DEW watermark becomes resistant to line shifts up to three pixels.

A more thorough evaluation of the DEW watermarking technique is given in [14]. In [14], we describe the benchmarking of the DEW and other algorithms in detail. In addition, from that reference, we give here Table II, which shows the visual quality rating of the DEW algorithm and the label bit error rate when the StirMark attack is used [19]. These results illustrate the robustness of the DEW algorithm as well as the invisibility of the watermark.

VIII. DISCUSSION

In this paper, we have derived, experimentally validated, and exploited a statistical model for our DCT -based DEW water-

---

[1]Our software implementation choices require that $n = 16 \cdot k^2$, where $k = 1, 2, 3 \cdots$. We therefore selected $n = 64$ instead of the optimal value $n = 54$.
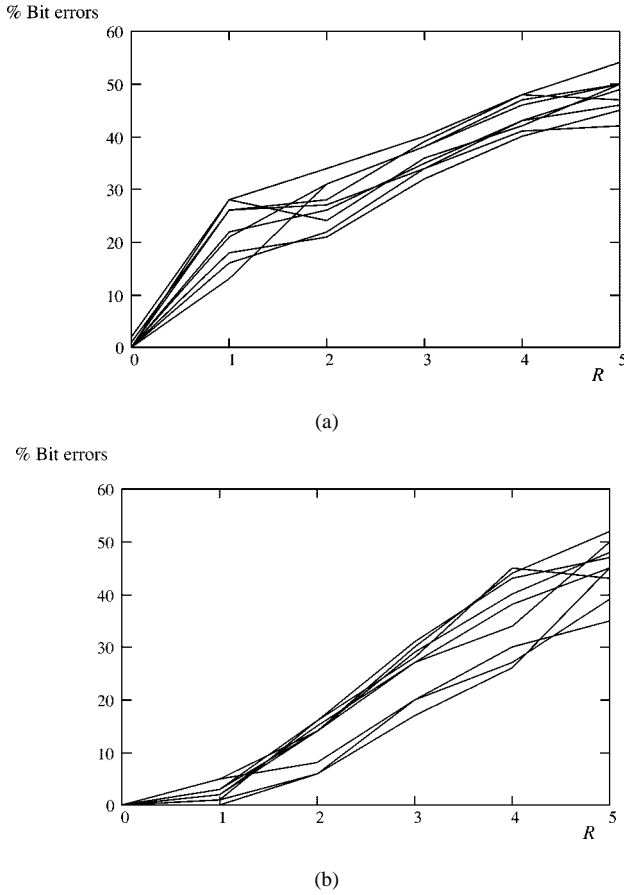
% Bit errors

(a)

% Bit errors

(b)

Fig. 12.   Percentage bit errors after shifting over R pixels using (a) parameter settings optimized for label size and (b) parameter settings optimized for robustness.

TABLE II
ITU–R REC. 500 QUALITY RATINGS AND PERCENTAGES LABEL BIT
ERRORS FOR THE DEW ALGORITHM AFTER APPLYING THE STIRMARK
ATTACK BASED ON GEOMETRICAL DISTORTIONS

Image name	Size	ITU-R Rec. 500 rating	% Label bit errors StirMark Attack[Pet98b]
Bike	720x512	4.3	34%
Bridge	720x512	4.5	16%
Butterfly	720x512	4.6	11%
Flower	720x512	4.5	15%
Grand Canyon	720x512	4.4	24%
Lena	512x512	4.6	17%
Parrot	720x512	4.7	28%
Rafter	720x512	4.3	24%
Red Square	720x512	4.6	15%
Sea	720x512	4.4	15%
Temple	720x512	4.6	17%
Tree	720x512	4.3	9%

marking algorithm. The performance of the DEW algorithm has been defined as its robustness against re-encoding attacks, the label size, and the visual impact. We have analytically shown how the performance is controlled by three parameters, namely Q_{jpeg}, n and $c_{\min}$. The derived statistical gives us an expression for the label bit error probability as a function of the three parameters Q_{jpeg}, n and $c_{\min}$. Using this expression, we can optimize a watermark for robustness, size, or visibility and add adequate error correcting codes.

The obtained expressions for the probability mass function of the cutoff indices can also be used for other purposes. For instance, with this PMF an estimate can be made for the variance of the watermarking "noise" that is added to an image by the DEW algorithm. This measure, possibly adapted to the human visual perception, can be used to carry out an overall optimization of the watermark embedding procedure using the (perceptually weighted) signal-to-noise-ratio as optimization criterion.

REFERENCES

[1] R. J. Anderson and F. A. P. Petitcolas, "On the limits of steganography," *IEEE J. Select. Areas Commun.*, vol. 16, pp. 474–481, May 1998.
[2] M. Barni, F. Bartolini, V. Cappellini, A. Lippi, and A. Piva, "A DWT-based technique for spatio-frequency masking of digital signatures," in *Proc. SPIE/IS&T Int. Conf. Security Watermarking Multimedia Contents*, vol. 3657, San Jose, CA, Jan. 25–27, 1999.
[3] P. Bas, J.-M. Chassery, and F. Davoine, "Self-similarity based image watermarking," in *Proc. IX Eur. Signal Processing Conf.*, Island of Rhodes, Greece, Sept. 1998, pp. 8–11.
[4] W. Bender, D. Gruhl, and N. Morimoto, "Techniques for data hiding," *Proc. SPIE*, vol. 2420, pp. 165–173, Feb. 9–10, 1995.
[5] G. W. Braudaway, "Protecting publicly-available images with an invisible watermark," in *Proc. IEEE Int. Conf. Image Processing*, Santa Barbara, CA, Oct. 1997.
[6] I. J. Cox, J. Kilian, F. T. Leighton, and T. Shamoon, "Secure spread spectrum watermarking for multimedia," *IEEE Trans. Image Processing*, vol. 6, Dec. 1997.
[7] F. Hartung and B. Girod, "Digital watermarking of MPEG-2 coded video in the bitstream domain," in *Proc. Int. Conf. Acoustics, Speech, Signal Processing '97*, vol. 4, Munich, Germany, Apr. 21–24, 1997, pp. 2621–2624.
[8] T. Kalker, G. Depovere, J. Haitsma, and M. Maes, "A video watermarking system for broadcast monitoring," in *Proc. SPIE Electronic Imaging '99, Security Watermarking Multimedia Contents*, San Jose, CA, Jan. 1999.
[9] E. Koch, J. Rindfrey, and J. Zhao, "Copyright protection for multimedia data," in *Proc. Int. Conf. Digital Media Electronic Publishing*, Leeds, U.K., Dec. 6–8, 1994.
[10] D. Kundur and D. Hatzinakos, "Digital watermarking using multi resolution wavelet decomposition," in *Proc. IEEE Int. Conf. Acoustics, Speech, Signal Processing*, vol. 5, Seattle, WA, May 1998, pp. 2969–2972.
[11] G. C. Langelaar, J. C. A. van der Lubbe, and R. L. Lagendijk, "Robust labeling methods for copy protection of images," in *Proc. SPIE Electronic Imaging '97, Storage, Retrieval Image Video Databases V*, San Jose, CA, Feb. 1997.
[12] G. C. Langelaar, R. L. Lagendijk, and J. Biemond, "Real-time labeling of MPEG-2 compressed video," *J. Vis. Commun. Image Represent.*, vol. 9, no. 4, pp. 256–270, 1998.
[13] G. C. Langelaar, "Conditional access to television service," in *Wireless Commun. Interactive Multimedia CD-ROM*, 3rd ed. Amsterdam, The Netherlands: Baltzer.
[14] G. C. Langelaar, "Real-time watermarking techniques for compressed video data," Ph.D. dissertation, [Online] Available: http://www-ict.its.tudelft.nl/~inald/pubs/Watermarking/WatermarkingofImagesandVideo.pdf, Delft Univ. Technol., Delft, The Netherlands, Jan. 2000.
[15] B. M. Macq and J. J. Quisquater, "Cryptology for digital TV broadcasting," *Proc. IEEE*, vol. 83, June 1995.

[16] F. Muller, "Distribution shape of two-dimensional DCT coefficients of natural images," *Electron. Lett.*, vol. 29, no. 22, pp. 1935–1936, 1993.

[17] N. Nikolaidis and I. Pitas, "Copyright protection of images using robust digital signatures," in *Proc. IEEE Int. Conf. Acoustics, Speech, Signal Processing '96*, vol. 4, Atlanta, GA, May 1996, pp. 2168–2171.

[18] W. B. Pennebaker and J. L. Mitchell, *The JPEG Still Image Data Compression Standard*. New York: Van Nostrand Reinhold, 1993.

[19] F. A. P. Petitcolas, R. J. Anderson, and M. G. Kuhn, "Weakness of copyright marking systems," in *Multimedia and Security Workshop at ACM Multimedia 1998*, Bristol, U.K., Sept. 1998.

[20] I. Pitas and T. H. Kaskalis, "Applying signatures on digital images," in *Proc. IEEE Workshop Nonlinear Signal Image Processing*, Neos Marmaras, Greece, June 20–22, 1995, pp. 460–463.

[21] A. Piva, M. Barni, F. Bartolini, and V. Cappellini, "DCT-based watermark recovering without resorting to the uncorrupted original image," in *Proc. Int. Conf. Image Processing '97*, Santa Barbara, CA, Oct. 1997.

[22] J. Puate and F. Jordan, "Using fractal compression scheme to embed a digital signature into an image," in *Proc. SPIE Photonics East Symp.*, Boston, MA, Nov. 18–22, 1996.

[23] M. Y. Rhee, *Error Correcting Coding Theory*. New York: McGraw-Hill, 1989.

[24] P. M. J. Rongen, M. J. J. J. B. Maes, and C. W. A. M. van Overveld, "Digital image watermarking by salient point modification," in *Proc. SPIE Electronic Imaging '99, Security Watermarking Multimedia Contents*, San Jose, CA, Jan. 1999.

[25] J. J. K. O Ruanaidh, W. J. Dowling, and F. M. Boland, "Phase watermarking of digital images," in *Proc. IEEE Int. Conf. Image Processing*, vol. 3, Lausanne, Switzerland, Sept. 16–19, 1996, pp. 239–242.

[26] J. R. Smith and B. O. Comiskey, "Modulation and information hiding in images," in *Proc. Information Hiding, Isaac Newton Institute Workshop*. Cambridge, U.K., May 1996.

[27] M. K. Varansai and B. Aazhang, "Parametric generalized Gaussian density estimation," *J. Acoust. Soc. Amer.*, vol. 86, pp. 1404–1415, Oct. 1989.

[28] R. B. Wolfgang, C. I. Podilchuk, and E. J. Delp, "Perceptual watermarks for digital images and video," *Proc. IEEE*, vol. 87, pp. 1108–1126, May 1999.

[29] J. Zhao and E. Koch, "Embedding robust labels into images for copyright protection," in *Proc. Int. Congr. Intellectual Property Rights Specialized Information, Knowledge, New Technologies*, Vienna, Austria, Aug. 21–25, 1995.

Gerrit C. Langelaar was born in Leersum, The Netherlands, in 1971. He received the M.Sc. degree from Technical University of Delft (TU-Delft), Delft, The Netherlands, in September 1995.

From September 1995 to October 1999, he was with the Information and Coommunication Theory Group, TU-Delft, working in in the European Union-sponsored SMASH project. He investigated and developed real-time watermarking techniques for compressed image and video data in copy protection systems. Since November 1999, he has been with the Philips Advanced Systems and Applications Laboratory, Eindhoven, The Netherlands, where he continues working in the field of copy protection systems.

Reginald L. Lagendijk received the M.Sc. and Ph.D. degrees in electrical engineering from the Technical University of Delft (TU-Delft), Delft, The Netherlands, in 1985 and 1990, respectively.

He became Assistant Professor and Associate Professor with TU Delft in 1987 and 1993, respectively. He was a Visiting Scientist with the Electronic Image Processing Laboratories, Eastman Kodak Research, Rochester, NY, in 1991. Since 1999, he has been Full Professor with the Information and Communication Theory Group, TU-Delft. He is author of the book *Iterative Identification and Restoration of Images* (Norwell, MA: Kluwer, 1991), and coauthor of the books *Motion Analysis and Image Sequence Processing* (Norwell, MA: Kluwer, 1993) and *Image and Video Databases: Restoration, Watermarking, and Retrieval* (Amsterdam, The Netherlands: Elsevier, 2000). He is currently Area Editor of *Signal Processing: Image Communication*. His research interests include signal processing and communication theory with emphasis on visual communications, compression, analysis, searching, and watermarking of image sequences. He has been involved in the European Research projects DART, SMASH, STORit, and DISTIMA. He is currently leading several projects in the field of wireless communications, among which the interdisciplinary research program Ubiquitous Communications at TU-Delft.

Dr. Lagendijk has served as Associate Editor of the IEEE TRANSACTIONS ON IMAGE PROCESSING and is a Member of the IEEE Signal Processing Society's Technical Committee on Image and Multidimensional Signal Processing.