

Research Article

Anonymous Fingerprinting with Robust QIM Watermarking Techniques

J. P. Prins, Z. Erkin, and R. L. Lagendijk

Information and Communication Theory Group, Faculty of Electrical Engineering, Mathematics, and Computer Science, Delft University of Technology, 2628 Delft, The Netherlands

Correspondence should be addressed to Z. Erkin, z.erkin@tudelft.nl

Received 20 March 2007; Revised 4 July 2007; Accepted 8 October 2007

Recommended by A. Piva

Fingerprinting is an essential tool to shun legal buyers of digital content from illegal redistribution. In fingerprinting schemes, the merchant embeds the buyer's identity as a watermark into the content so that the merchant can retrieve the buyer's identity when he encounters a redistributed copy. To prevent the merchant from dishonestly embedding the buyer's identity multiple times, it is essential for the fingerprinting scheme to be anonymous. Kuribayashi and Tanaka, 2005, proposed an anonymous fingerprinting scheme based on a homomorphic additive encryption scheme, which uses basic quantization index modulation (QIM) for embedding. In order, for this scheme, to provide sufficient security to the merchant, the buyer must be unable to remove the fingerprint without significantly degrading the purchased digital content. Unfortunately, QIM watermarks can be removed by simple attacks like amplitude scaling. Furthermore, the embedding positions can be retrieved by a single buyer, allowing for a locally targeted attack. In this paper, we use robust watermarking techniques within the anonymous fingerprinting approach proposed by Kuribayashi and Tanaka. We show that the properties of an additive homomorphic cryptosystem allow for creating anonymous fingerprinting schemes based on distortion compensated QIM (DC-QIM) and rational dither modulation (RDM), improving the robustness of the embedded fingerprints. We evaluate the performance of the proposed anonymous fingerprinting schemes under additive-noise and amplitude-scaling attacks.

Copyright © 2007 J. P. Prins et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

1. INTRODUCTION

Intellectual property protection is a severe problem in today's digital world due to the ease of illegal redistribution through the Internet. As a countermeasure to deter people from illegally redistributing digital content such as audio, images, and video, a fingerprinting scheme embeds specific information related to the identity of the buyer by using watermarking techniques. In conventional fingerprinting schemes, this identity information is embedded into the digital data by the merchant and the fingerprinted copy is given to the buyer. When the merchant encounters redistributed copies of this fingerprinted content, he can retrieve the identity information of the buyer who (illegally) redistributed his copy. From the buyer's point of view, however, this scenario is unattractive because during the embedding procedure, the merchant obtains the identity information of the buyer. This enables a cheating merchant to embed the identity information of the

buyer into any content without the buyer's consent and subsequently accuse the buyer of illegal redistribution.

To protect the identity of the buyer, anonymous fingerprinting schemes have been proposed [1, 2]. In [2], the buyer and the merchant follow an interactive embedding protocol, in which the identity information of the buyer remains unknown to the merchant. When the buyer wishes to purchase, for instance, an image, he registers himself to a registration centre and receives a proof of his identity with a signature of the registration centre. Then the buyer encrypts his identity and sends both encrypted identity and the proof of identity to the merchant. The merchant checks the validity of the signature by using the public key of the registration centre. After the buyer convinces the merchant, through the provided identity proof, that the encrypted identity indeed contains the identity information of the buyer, the merchant embeds the identity information of the buyer into the (encrypted) image data by exploiting the homomorphic

property of the cryptosystem. Then the encrypted fingerprinted image is sent to the buyer for decryption and future use.

In this scheme, the merchant can only retrieve the identity information of the buyer when it is detected in a copy of the fingerprinted image. This idea, first presented in [2], was constructed in [3, 4] using digital coins. In order to embed the identity information of the buyer, a single-bit commitment scheme with exclusive, or homomorphism, is used that allows for computing the encrypted XOR of two bits by multiplying their ciphertexts. In [5], Kuribayashi and Tanaka observe that this construction is not efficient because of the low enciphering rate. The single bit commitment scheme can only contain one bit of information for a $\log_2 n$ -bit ciphertext, where n is a product of two large primes.

In order to increase the enciphering rate, Kuribayashi and Tanaka suggested using a cryptosystem with a larger message space. They introduced an anonymous fingerprinting algorithm based on an additive homomorphic cryptosystem that allows for the addition of values in the plaintext domain by multiplying their corresponding ciphertexts. Consequently, Kuribayashi and Tanaka used a basic amplitude quantization-based scheme similar to the well-known quantization index-modulation (QIM) scheme as the underlying watermarking scheme. Since QIM essentially modulates (integer-valued) quantization levels to embed information bits into a signal, QIM can elegantly be implemented in an additive homomorphic cryptosystem. However, QIM is a basic watermarking scheme that has limited robustness compared to other watermarking schemes. The embedding positions can easily be retrieved from an individual fingerprinted copy and are thus vulnerable to local attacks. Such attacks result in minimal overall signal degradation, while completely removing the fingerprint. Furthermore, QIM is vulnerable to simple, either malevolent or unintentional, global attacks such as randomization of the least significant bits, addition of noise, compression, and amplitude scaling.

In this paper, we use the ideas in [5] to build anonymous versions of state-of-the-art watermarking schemes, namely, distortion-compensated QIM (DC-QIM) [6] and rational dither modulation (RDM) [7]. By adapting these watermarking schemes to the anonymous fingerprinting protocol of Kuribayashi and Tanaka, we improve the robustness of the embedded fingerprints and, as a consequence, the merchant's security. As DC-QIM and RDM are based on subtractive-dither QIM (SD-QIM), they both hide the embedding locations from the buyer more effectively, preventing local, targeted attacks on the fingerprint. With respect to global attacks, like additive noise and amplitude scaling, RDM is provably equivalent in robustness, while DC-QIM is provably better in robustness against additive noise attacks. Furthermore, RDM improves the QIM scheme so that the fingerprint becomes robust to amplitude-scaling attacks.

The outline of this paper is as follows. In Section 2, we introduce the basic QIM watermarking scheme, as well as the additive homomorphic cryptosystem of Okamoto-Uchiyama [8], on which the approach in [5] is based. In Section 3, we review the anonymous fingerprinting scheme by Kurib-

TABLE 1: Table of symbols.

A.1. Cryptosystems	
Symbol	Usage
p, q	Large primes of size k
n	Modulus
g	Generator
m	Message
c	Cipher-text
$r, s \in_R \mathbb{Z}_n^*$	r and s are random blinding factors from \mathbb{Z}_n^*
$E(m)$	Encryption (and integer rounding) of m
$D(c)$	Decryption of ciphertext c
A.2. Watermarking and fingerprinting	
Symbol	Usage
x/X	Original sample/original signal
y/Y	Watermarked sample/watermarked signal
z/Z	Received sample/received signal
w/W	Individual watermark bit/total watermark
d	Dither
Δ	Quantization step size
$Q_\Delta(\cdot)$	Uniform quantizer with step size Δ
α	DC-QIM factor
ρ	Gain factor
c	Scaling factor used for rounding/reducing quantization step size
$\nu(\cdot)$	Function to normalize coefficients for RDM.
id	Buyer identity

ayashi and Tanaka. In Section 4, we describe the proposed anonymous fingerprinting schemes using the subtractive dither QIM, DC-QIM, and RDM watermarking schemes. Section 5 describes the experiments that evaluate the robustness of the proposed schemes compared to the original watermarking schemes. Section 6 discusses the security benefits of using specially constructed buyer ids. Conclusions are given in Section 7. A list of used symbols is provided in Table 1.

2. WATERMARKING AND ENCRYPTION PRELIMINARIES

2.1. Basic quantization-index modulation

Quantization-index modulation (QIM) is a relatively recent watermarking technique [6]. It has become popular because of the high watermarking capacity and the ease of implementation. The basic quantization-index modulation algorithm embeds a watermark bit w by quantizing a single-signal sample x by choosing between a quantizer with even or odd values, depending on the binary value of w . These quantizers with a step size $\Delta \in \mathbb{N}$ are denoted by $Q_{\Delta\text{-even}}(\cdot)$ and $Q_{\Delta\text{-odd}}(\cdot)$, respectively.

Figure 1 shows the input and output characteristics of the quantizer, where $w \in \{0, 1\}$ denotes the message bit that is

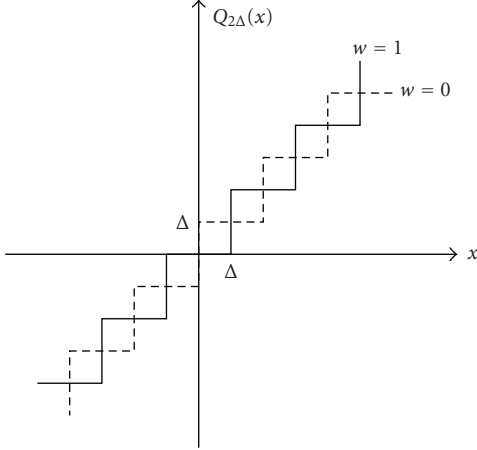


FIGURE 1: Quantizer input-output characteristics.

embedded into the host data. The watermarked signal sample y then is

$$y = \begin{cases} Q_{\Delta\text{-even}}(x), & \text{if } w = 0, \\ Q_{\Delta\text{-odd}}(x), & \text{if } w = 1. \end{cases} \quad (1)$$

The quantizers $Q_{\Delta\text{-even}}(\cdot)$ and $Q_{\Delta\text{-odd}}(\cdot)$ are designed such that they avoid biasing the values of y , that is, the expected (average) value of x and y are identical. The trade-off between embedding distortion and robustness of QIM against additive noise attacks is controlled by the value of Δ . The detection algorithm requantizes the received signal sample z with both $Q_{\Delta\text{-even}}(\cdot)$ and $Q_{\Delta\text{-odd}}(\cdot)$. The detected bit $\hat{w} = \{0, 1\}$ is determined by the quantized value $Q_{\Delta\text{-even}}(z)$ or $Q_{\Delta\text{-odd}}(z)$ with the smallest distance to the received sample z .

This scheme of even and odd quantizers can also be implemented by using a single quantizer with a step-size of 2Δ and subtracting/adding Δ when $w = 1$. Implementing the quantizer in this way allows for the implementation of the scheme in the encrypted domain as was shown in [5].

A serious drawback of basic QIM watermarking is its sensitivity to amplitude-scaling attacks [7], in which signal samples are multiplied by a gain factor ρ . If the gain factor ρ is constant for all samples, the attack is called a fixed-gain attack (FGA). In amplitude-scaling attacks, the detector does not possess the factor ρ , which causes a mismatch between embedder and decoder quantization lattices, affecting the QIM-detector performance dramatically.

Another drawback of basic QIM is that the embedding positions can be retrieved from a single copy. The embedding positions are those signal values x_i that have been (heavily) quantized to $Q_{\Delta\text{-even}}(x_i)$ and $Q_{\Delta\text{-odd}}(x_i)$, and have a constant difference value equal to Δ , that is, the quantizer coarseness parameter. By constructing a high-resolution histogram, the buyer can easily observe the even-spaced spikes of signal intensity values and identify, and thus attack the embedding positions locally. This results in the removal of the fingerprint with little degradation to the overall signal.

2.2. Homomorphic encryption schemes

The idea of processing encrypted data was first suggested by Ahituv et al. in [9]. In their paper, the problem of decrypting data before applying arithmetic operations is addressed and a new approach is described as processing data without decrypting it first.

Succeeding works showed that some asymmetric cryptosystems preserve structure, which allows for arithmetic operations to be performed on encrypted data. This structure preserving property, called homomorphism, comes in two main types, namely, additive and multiplicative homomorphism. Using additive homomorphic cryptosystems, performing a particular operation (e.g., multiplication) with encrypted data, results in the addition of the plaintexts. Similarly, using a multiplicatively homomorphic cryptosystem, multiplying ciphertexts, results in the multiplication of the plaintexts. Paillier [10], Okamoto-Uchiyama [8], and Goldwasser-Micali [11] are additively homomorphic cryptosystems while RSA [12] and ElGamal [13] are multiplicatively homomorphic cryptosystems.

The anonymous fingerprinting scheme proposed in [5] is based on the addition of the fingerprint to the digital data, and hence, an additive cryptosystem is used. Among the candidates, the Okamoto-Uchiyama cryptosystem is chosen for efficiency considerations [5]. In the next section, the Okamoto-Uchiyama cryptosystem is described. We observe, however, that the anonymous fingerprinting schemes, proposed in this paper, can easily be implemented by using other additively homomorphic cryptosystems. It is, however, required to have a sufficiently large message space to represent the signal samples. Further, the underlying security protocols, such as the proof protocol for validating the buyer identity, must be suitable for the chosen cryptosystem.

A requirement for the cryptosystem is that it is probabilistic in order to withstand chosen plaintext attacks. Such attacks are easily performed in our scheme because individual signal samples are usually limited in value (e.g., 8 bit). If we were to use a nonprobabilistic cryptosystem, this would enable the buyer to construct a codebook of ciphertexts for all possible messages (in total, $2^8 = 256$) using the public key and decrypt through this codebook. Fortunately probabilistic cryptosystems were introduced in [11], which enable the encryption of a single plaintext to n ciphertexts, where n is a security parameter related to the size of the key. To which ciphertext the plaintext is encrypted is dependent on a blinding factor r , which is usually taken at random. Selecting different r 's does not affect the decrypted plaintext. By having a multitude of ciphertexts for a single plaintext, the size of a codebook will become $2^8 \cdot 2^n$, and thus impractically large, preventing such attacks. All the above-mentioned additive homomorphic-encryption schemes (Paillier, Okamoto-Uchiyama, and Goldwasser-Micali) are probabilistic, and hence withstand chosen plaintext attacks.

From Section 3 onwards, we compactly denote the encryption and the decryption of a message with $E(m)$ and $D(c)$, respectively, omitting the dependency on the random factor r . In the scope of this paper, an additive homomorphic cryptosystem will be used for encrypting signal samples

which do not necessarily need to be integer values. In this case, rounding to the nearest integer value precedes the encryption, and thus, in this paper, $E(\cdot)$ denotes both rounding and encryption.

2.2.1. Okamoto-Uchiyama cryptosystem

Okamoto and Uchiyama [8] proposed a semantically secure and probabilistic public key cryptosystem based on composite numbers. Let $n = p^2q$, where p and q are two prime numbers of length k bits, and let g be a generator such that the order of $g^{p-1} \bmod p^2$ is p . Another generator is defined as $h = g^n$. In this scheme, the public key is $pk = (n, g, h, k)$ and the secret key is $sk = (p, q)$.

Encryption.

A message m ($0 < m < 2^{k-1}$) is encrypted as follows:

$$c = E(m, r) = g^m h^r \bmod n, \quad (2)$$

where r is a random number in \mathbb{Z}_n^* .

Decryption.

Decoding the cipher-text is defined as

$$m = D(c) = \frac{L(c^{p-1} \bmod n)}{L(g^{p-1} \bmod n)} \bmod p, \quad (3)$$

where the function $L(\cdot)$ is

$$L(u) = \frac{u-1}{p}. \quad (4)$$

The Okamoto-Uchiyama cryptosystem has the additive homomorphic property such that, given two encrypted messages $E(m_1, r_1)$ and $E(m_2, r_2)$, the following equality holds:

$$\begin{aligned} E(m_1, r_1) \times E(m_2, r_2) &= g^{m_1} h^{r_1} \times g^{m_2} h^{r_2} \bmod n \\ &= g^{m_1+m_2} h^{r_1+r_2} \bmod n \\ &= E(m_1 + m_2, r_1 + r_2). \end{aligned} \quad (5)$$

Here, \times denotes integer-modulo- n multiplication.

3. KURIBAYASHI AND TANAKA ANONYMOUS FINGERPRINTING PROTOCOL

The fingerprinting scheme in [5] is carried out between buyer and merchant, and has, as objective to anonymously embed, the buyer's identity information into the merchant's data (e.g., audio, image, or video signal). The buyer decomposes his l -bit identity W into bits as $W = (w_0, w_1, \dots, w_{l-1})$. For applications such as embedding identity information in multimedia data, the value of l is typically between 32 and 128 (bits), which is sufficiently large to prevent the merchant from guessing valid buyer ids. Where necessary, we assume that the probability $P[w_j = 0]$ and $P[w_j = 1]$ are equal. After decomposition of W into individual bits, the buyer encrypts each bit with his public key using the Okamoto-Uchiyama

cryptosystem, so that $E(W) = (E(w_0), E(w_1), \dots, E(w_{l-1}))$. These encrypted values are sent to the merchant.

The merchant first quantizes the samples of the (audio, image, and video) signal that the buyer wishes to obtain, using a quantizer with coarseness 2Δ , that is, $x' = Q_{2\Delta}(x)$. Here, the quantizer step size Δ is a positive integer to ensure that the quantized value can be encrypted. He then encrypts all quantized signal samples x' with the public key of the buyer, yielding $E(x')$. The merchant selects watermark embedding positions by using a unique secret key that will be used to extract the watermark from the redistributed copies. In order to embed a single bit of information w_j into one of the quantized and encrypted value $E(x')$ at a particular watermark embedding position, the merchant performs the following operation:

$$\begin{aligned} E(y) &= E(x') \times E(w_j)^\Delta \\ &= E(x' + w_j \Delta). \end{aligned} \quad (6)$$

The result is an encrypted and watermarked signal value y , as can be readily seen by the following relation:

$$\begin{aligned} D(E(y)) &= x' + w_j \Delta, \\ y &= \begin{cases} Q_{2\Delta}(x), & \text{if } w_j = 0, \\ Q_{2\Delta}(x) + \Delta, & \text{if } w_j = 1. \end{cases} \end{aligned} \quad (7)$$

The encrypted signal, with the buyer's identity information embedded into it in the form of a watermark, is finally sent to the buyer. Obviously, only the buyer can decrypt the watermarked signal values.

In order for the system to be robust against local attacks, the relation between the buyer's identity-information bits w_j and the signal values y (audio samples, image, or video pixels), into which the information bits are embedded, should be kept secret from the buyer. Note that, as a consequence, all signal values x will have to be encrypted, also the ones that do not carry a bit w_j of the buyer's identity information, as so to hide these embedding positions.

Compared to the original QIM scheme in (1), the above watermarking scheme introduces a bias, as the expected (average) value of y is $\Delta/2$ larger than that of x . This bias is introduced because Δw_j is always added to the quantized signal value x' and never subtracted. In order to avoid this undesirable side effect, either the even or odd quantizer should be selected depending on the watermark bit w_j as in (1). However, the merchant has only the encrypted version of each watermark bit w_j , which prevents him from deciding between the two quantizers. To overcome this problem, the merchant compares the signal values x and x' , and depending on the result, the encrypted value of Δw_j can be added or subtracted [5]. When x' is smaller than x , Δw_j is added, otherwise, it is subtracted. This procedure now is equivalent to (1) and thus effectively removes the bias. As the decision is not dependent on the value of w_j , no information is leaked about the value of w_j . The resulting embedding procedure for identity-information bit w_j then becomes

$$E(y) = \begin{cases} E(x') \times E(w_j)^\Delta, & \text{if } x \geq Q_{2\Delta}(x), \\ E(x') \times (E(w_j)^\Delta)^{-1}, & \text{if } x < Q_{2\Delta}(x), \end{cases} \quad (8)$$

where $()^{-1}$ denotes modular inverse in the cyclic group defined by the encryption scheme. When the buyer decrypts the received encrypted and watermarked signal values, he obtains the following result for the watermark embedding positions:

$$y = \begin{cases} x' + w_j \Delta, & \text{if } x \geq Q_{2\Delta}(x), \\ x' - w_j \Delta, & \text{if } x < Q_{2\Delta}(x). \end{cases} \quad (9)$$

For all other positions, the unwatermarked and unchanged, but encrypted and therefore rounded, signal values x are transmitted.

In the above embedding protocol, we have assumed that the buyer provides encrypted values of a valid binary decomposition $(w_0, w_1, \dots, w_{l-1})$ of his identity information W to the merchant. Since, however, the decomposed bits of the identity information of the buyer are encrypted, the merchant cannot easily check this assumption. In the original work by Kuribayashi and Tanaka [5], a registration center is used, which assures the legitimacy of the buyer. During the purchase, the merchant first confirms the identity of the buyer, and then the buyer proves the validity of the decomposed bits of his identity information by using zero-knowledge proof protocols. Since this procedure is entirely independent of the watermarking scheme, we refer, for details on the identity and decomposition validation and the security of this procedure, to [5], where it is given for the Okamoto-Uchiyama encryption scheme. The focus of this paper is on the application of the homomorphic embedding procedure described above to the more robust watermarking schemes of [6, 7].

4. ANONYMOUS FINGERPRINTING USING ADVANCED WATERMARKING SCHEMES

From the perspective of the merchant, the embedding of the buyer's identification information must be as robust as possible in order to both withstand malicious and benign signal-processing operations on the fingerprinted signal. If the buyer id-embedding procedure is not robust, the buyer could remove the fingerprint either intentionally or unintentionally, and as a consequence, the merchant would lose his ability to trace illegally redistributed copies. The fingerprints embedded in the Kuribayashi and Tanaka (KT) anonymous fingerprinting protocol, described in Section 3, are known to be sensitive to a number of signal-processing operations, and are, in fact, relatively easy to remove through attacks mentioned in Section 2.1. We propose to increase the robustness of the Kuribayashi and Tanaka anonymous fingerprinting protocol, as perceived by the merchant, by applying their approach to two advanced quantization-based watermarking schemes, namely, DC-QIM and RDM.

So far, we have embedded the bits of the identity information into signal values without specifying what these signal values actually are. In the rest of this paper, we will use block-DCT transform coefficients of images to embed the identity bits into. A particular block-DCT coefficient, into which, we embed an information bit w_j , will be abstractly denoted by x_i . Of course, in actual images, x_i may be a particular DCT coefficient of a particular DCT block in the image.

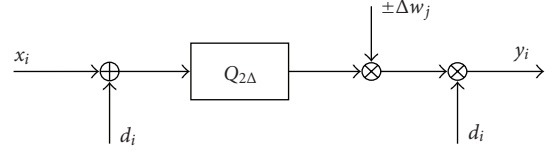


FIGURE 2: Subtractive dither QIM.

The relation between the bits w_j and watermark embedding positions x_i is determined by a key known only to the merchant. In practical cases of interest, the number of candidate embedding positions is in the same order as the number of signal samples, whereas the number of information bits is typically between 32 and 128. For instance, for a 1024×1024 pixels image, the maximum number of possible embedding combinations for 128 bits of information is $\binom{1024^2}{128}$, which provides enough security. In the case of embedding the bits w_j into DCT coefficients, the number of possible embedding combinations will be smaller depending on the DCT block size and the number of DCT coefficients in one block that are (perceptually and qualitatively) suitable for embedding a watermark bit into.

It is important to note that the goal for each watermarking scheme within the Kuribayashi-Tanaka protocol is to compute the encryption of watermarked coefficients y_i , while only having available original signal values x_i , the encrypted bits $E(w_j)$ of the buyer's decomposed identity, and the public key pk of the selected additively homomorphic encryption scheme. Once the buyer identification information is correctly embedded in the encrypted domain, the encrypted coefficients (i.e., encrypted digital content) will be sent to the buyer, who can decrypt these with his private key to obtain correctly watermarked data. Since the information bits are embedded in the DCT domain, a trivial inverse DCT on the decrypted data is necessary as the last step to obtain the purchased digital image. Because this is easiest performed in the plaintext domain, we leave it to the buyer to perform this inverse DCT after decryption, which is much like JPEG decompression.

4.1. Subtractive dither-quantization-index modulation

Fingerprints embedded by the basic QIM watermarking scheme used by Kuribayashi and Tanaka as described in Section 2.1 can be locally attacked because the buyer can find the embedding positions x_i without checking all possible (for instance $\binom{1024^2}{128}$) combinations. A common solution to this weakness of the basic QIM watermarking scheme is to add pseudorandom noise, usually called dither, to x_i before embedding an information bit w_j , and subtracting the dither after embedding. As a consequence, the quantization levels and their constant difference Δ can no longer be observed, making the separation between embedding positions x_i and nonembedding positions impossible. The resulting watermarking scheme, illustrated in Figure 2, is called subtractive dither QIM (SD-QIM).

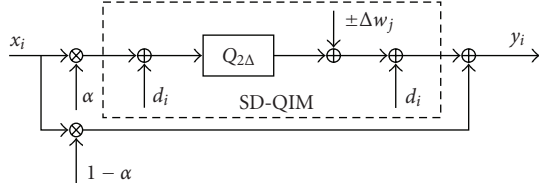


FIGURE 3: Distortion-compensated QIM.

In QIM terminology, a small amount of dither d_i is added prior to quantizing the signal amplitude x_i to an odd or even value depending on the information bit w_j . After quantization of $x_i + d_i$, the same amount of dither d_i is subtracted. It is desirable that the dither can be used in cooperation with the QIM uniform quantizers $Q_{\Delta\text{-odd}}(\cdot)$ and $Q_{\Delta\text{-even}}(\cdot)$, which use a quantization step size of 2Δ , as in the basic QIM. It has been shown [14] that a suitable choice for the PDF of the random dither d_i is a uniform distribution on $[-\Delta, \Delta]$.

In order to embed the buyer's identity information bit $E(w_j)$ into coefficient x_i using the Kuribayashi-Tanaka protocol in combination with subtractive dither, we carry out the following protocol.

- (i) Add random dither d_i to the signal sample or coefficient x_i .
- (ii) Quantize $x_i + d_i$ with a quantization coarseness of 2Δ , and encrypt the result using the buyer's public key, yielding $E(Q_{2\Delta}(x_i + d_i))$.
- (iii) Multiply by $E(w_j)^\Delta$ or its modular inverse depending on the value of $x_i + d_i$, in order to achieve the desired quantization level.
- (iv) Encrypt the dither d_i to obtain $E(d_i)$. Note that, since $d_i \in \mathbb{R}$, the encryption operation includes modulo n rounding to an integer. Multiply the result of the previous step with the modular inverse of $E(d_i)$ as to implement the subtraction of the dither d_i from $Q_{2\Delta}(x_i + d_i)$.

Summarizing the above protocol steps, we obtain

$$E(t_i) = \begin{cases} E(Q_{2\Delta}(x_i + d_i)) \times E(w_j)^\Delta, & \text{if } x_i \geq Q_{2\Delta}(x_i), \\ E(Q_{2\Delta}(x_i + d_i)) \times (E(w_j)^\Delta)^{-1}, & \text{if } x_i < Q_{2\Delta}(x_i), \end{cases} \quad (10)$$

$$E(y_i) = E(t_i) \times E(d_i)^{-1}.$$

After decryption, the buyer obtains the (DCT transformed) image, into which, his identity information is embedded in certain DCT coefficients y_i according to the following subtractive dither QIM scheme

$$y_i = \begin{cases} Q_{\Delta\text{-even}}(x_i + d_i) - d_i, & \text{if } w_j = 0, \\ Q_{\Delta\text{-odd}}(x_i + d_i) - d_i, & \text{if } w_j = 1. \end{cases} \quad (11)$$

The above embedding procedure demonstrates the usage of the Kuribayashi-Tanaka protocol to subtractive-dither QIM. The plaintext subtractive-dither QIM and the above Kuribayashi-Tanaka subtractive-dither QIM (KT SD-QIM)

are equivalent except for the rounding of the dither d_i to integers before encryption. How to limit the adverse effect of integer rounding will be addressed next.

Two improvements of (10) are desirable. In the first place, we can subtract d_i before encrypting $Q_{2\Delta}(x_i + d_i)$. This effectively removes the last protocol step, and hence eliminates an unnecessary encryption operation. The resulting scheme can then be rewritten as follows:

$$E(y_i) = \begin{cases} E(Q_{2\Delta}(x_i + d_i) - d_i) \times E(w_j)^\Delta, & \text{if } x_i \geq Q_{2\Delta}(x_i), \\ E(Q_{2\Delta}(x_i + d_i) - d_i) \times (E(w_j)^\Delta)^{-1}, & \text{if } x_i < Q_{2\Delta}(x_i). \end{cases} \quad (12)$$

The second improvement concerns the quantization operation. The quantizer not only rounds the signal amplitudes to predetermined (not necessarily integer) quantization levels, but it must also round signal values or DCT coefficients $x_i + d_i$ to integers because of the ensuing encryption operation. If the signal values of DCT coefficients x_i are sufficiently large, using integer-valued coefficients is not a restriction at all. For smaller values of x_i , however, using integer values may be too restrictive or may yield too large deviations between the results of (12) and (11).

We propose to circumvent this problem by scaling all coefficients x_i with a constant factor c before embedding. Scaling has little effect on the en-/decryption, as long as the samples are not scaled beyond the message group size of the encryption scheme used. The message group size is, however, usually very large because of encryption security requirements (typically $> 2^{512}$). As a consequence of scaling x_i , the dither d_i and all encrypted bits $E(w_j)$ of the decomposed identity of the buyer also have to be scaled by c . We note that scaling introduces extra computation. However, the dither can be scaled and subtracted before encryption, resulting in a very small increase in complexity. The scaling of the encrypted bits $E(w_j)$ of the decomposed identity of the buyer has to be taken into account in the protocol steps, which is relatively easy since the scaling can be combined with the multiplication of w_j with Δ . The resulting embedding equation can be summarized as follows:

$$E(y_i) = \begin{cases} E(c \cdot (Q_{2\Delta}(x_i + d_i) - d_i)) \times E(w_j)^\Delta, & \text{if } x_i \geq Q_{2\Delta}(x_i), \\ E(c \cdot (Q_{2\Delta}(x_i + d_i) - d_i)) \times (E(w_j)^\Delta)^{-1}, & \text{if } x_i < Q_{2\Delta}(x_i). \end{cases} \quad (13)$$

The scaling factor c has to be communicated to the buyer so that the buyer can rescale the entire image after decryption to the proper (original) intensity range.

4.2. Distortion-Compensated QIM

Distortion-compensated QIM (DC-QIM) [6] is an extension to the subtractive dither-QIM scheme described in the previous section. Rather than directly adding dither to and quantizing of x_i , a fraction $\alpha \cdot x_i$ is used in the SD-QIM procedure (see Figure 3). The information bits will be embedded only in the fraction $\alpha \cdot x_i$, where α lies within the range $[0, 1]$. The remaining fraction $(1 - \alpha) \cdot x_i$ is added back to the watermarked

signal component $\alpha \cdot x_i$ to form the final embedded coefficient y_i . The embedder chooses an appropriate value for α depending on the desired detection performance and robustness of DC-QIM; an often selected value is as in [15]:

$$\alpha = \frac{\sigma_w^2}{\sigma_w^2 + \sigma_n^2}, \quad (14)$$

where $\sigma_w^2 = \Delta^2/3$ is the variance of the watermark in the watermarked signal, and σ_n^2 is the variance of the noise or other degradation that an attacker applies in an attempt to render the watermark bits undetectable. Obviously, the standard SD-QIM scheme is optimal only if an attacker inserts little or no noise into the watermarked image since, for $\sigma_n^2 \rightarrow 0$, we find $\alpha \rightarrow 1$. The difference in robustness between SD-QIM and DC-QIM becomes especially relevant if the variance of the attacker becomes large relative to σ_w^2 , that is, $\sigma_n^2 \rightarrow \sigma_w^2$.

As the differences between the SD-QIM and DC-QIM watermarking scheme merely consist of plaintext multiplications and ciphertext additions, DC-QIM can also be achieved within the limitations of the homomorphic additive encryption scheme used by the Kuribayashi-Tanaka protocol. The basic embedding operations can now be written as follows:

$$E(t_i) = \begin{cases} E(Q_{2\Delta}(\alpha \cdot x_i + d_i) - d_i) \times E(w_j)^\Delta, & \text{if } \alpha \cdot x_i \geq Q_{2\Delta}(\alpha \cdot x_i), \\ E(Q_{2\Delta}(\alpha \cdot x_i + d_i) - d_i) \times (E(w_j)^\Delta)^{-1}, & \text{if } \alpha \cdot x_i < Q_{2\Delta}(\alpha \cdot x_i), \end{cases} \quad (15)$$

$$E(y_i) = E(t_i) \times E((1 - \alpha) \cdot x_i).$$

Equation (15) results in the following watermarked values y_i after decryption:

$$t_i = \begin{cases} Q_{2\Delta}(\alpha \cdot x_i + d_i) - d_i + w_j \cdot \Delta, & \text{if } \alpha \cdot x_i \geq Q_{2\Delta}(\alpha \cdot x_i), \\ Q_{2\Delta}(\alpha \cdot x_i + d_i) - d_i - w_j \cdot \Delta, & \text{if } \alpha \cdot x_i < Q_{2\Delta}(\alpha \cdot x_i), \end{cases} \quad (16)$$

$$y_i = t_i + (1 - \alpha) \cdot x_i.$$

The plaintext distortion-compensated QIM and the above Kuribayashi-Tanaka distortion-compensated QIM (KT DC-QIM) are equivalent, except again for the rounding of the real-valued dither d_i and $(1 - \alpha) \cdot x_i$ to integers before encryption.

Similar to the subtractive dither-QIM watermark algorithm, KT DC-QIM can be modified to subtract the dither before encryption, and to scale the signal values before encryption. Furthermore, the term $(1 - \alpha) \cdot x_i$ can be added before encryption, further reducing the number of encryptions needed. The resulting KT DC-QIM embedding equations then become:

$$E(t_i) = \begin{cases} E(c \cdot (Q_{2\Delta}(\alpha \cdot x_i + d_i) - d_i)) \times E(w_j)^\Delta, & \text{if } \alpha \cdot x_i \geq Q_{2\Delta}(\alpha \cdot x_i), \\ E(c \cdot (Q_{2\Delta}(\alpha \cdot x_i + d_i) - d_i)) \times (E(w_j)^\Delta)^{-1}, & \text{if } \alpha \cdot x_i < Q_{2\Delta}(\alpha \cdot x_i). \end{cases} \quad (17)$$

$$E(y_i) = E(t_i) \times E(c \cdot (1 - \alpha) \cdot x_i).$$

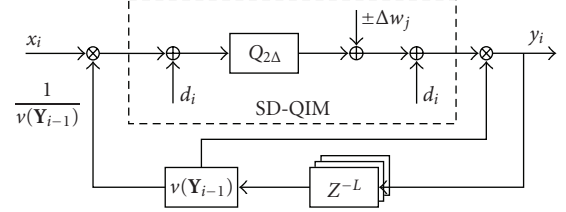


FIGURE 4: Rational dither modulation.

4.3. Rational dither modulation

DC-QIM provides a significant improvement in robustness compared to the basic QIM scheme. Nevertheless, the DC-QIM scheme is known to be very sensitive to gain or volumetric attacks, which is just simply scaling of the image intensities. Because of the use of the scaling factor c in SD-QIM and DC-QIM in order to reduce the sensitivity to integer-rounding before encryption, the buyer has an excellent opportunity to perform a gain attack on the watermarked signal. The gain effect causes the quantization levels used at the detector to be misaligned with those embedded in the purchased and illegally distributed digital data, effectively making the retrieval of the watermarked identity bits impossible [16].

Perez-Gonzalez et al. [7], proposed the usage of QIM on ratios between signal samples as so to make the watermarking system robust against fixed gain attacks. The resulting approach, known as rational dither modulation (RDM), is robust against both additive-noise and fixed-gain attacks. The RDM-embedding scheme is illustrated in Figure 4. The robustness against fixed gain attacks is achieved by normalizing the signal value (or DCT coefficient) x_i by $v(\mathbf{Y}_{i-1})$, which is a function that combines L previous watermarked signal values $\mathbf{Y}_{i-1} = (y_{i-1}, y_{i-2}, \dots, y_{i-L})$. An example for the function $v(\mathbf{Y}_{i-1})$ is the Hölder vector norm, as suggested in [7]:

$$v(\mathbf{Y}_{i-1}) = \left(\frac{1}{L} \sum_{m=i-L}^{i-1} |y_m|^p \right)^{1/p}. \quad (18)$$

The SD-QIM watermark embedding will then take place using the normalized signal values $x_i/v(\mathbf{Y}_{i-1})$, yielding

$$y_i = \begin{cases} v(\mathbf{Y}_{i-1}) \cdot \left(Q_{\Delta\text{-even}} \left(\frac{x_i}{v(\mathbf{Y}_{i-1})} + d_i \right) - d_i \right), & \text{if } w_j = 0, \\ v(\mathbf{Y}_{i-1}) \cdot \left(Q_{\Delta\text{-odd}} \left(\frac{x_i}{v(\mathbf{Y}_{i-1})} + d_i \right) - d_i \right), & \text{if } w_j = 1, \end{cases} \quad (19)$$

where the multiplication of the quantization results with $v(\mathbf{Y}_{i-1})$ is required to scale the coefficients to their original value range. Another way of viewing RDM is that it is equivalent to using SD-QIM with a signal amplitude-dependent quantization coarseness $v(\mathbf{Y}_{i-1}) \cdot \Delta$.

The normalization of x_i takes place on a function of $(y_{i-1}, y_{i-2}, \dots, y_{i-L})$ rather than of $(x_{i-1}, x_{i-2}, \dots, x_{i-L})$. The usage of $v(\mathbf{Y}_{i-1})$ is preferable because only the watermarked

values y_i are available during watermark detection. In the Kuribayashi-Tanaka protocol, the watermarked signal values or DCT coefficients y_i are only available to the merchant in an encrypted form $E(y_i)$. Unfortunately, the embedder cannot make use of $v(\mathbf{Y}_{i-1})$ as a normalization factor, primarily because the homomorphic division (and multiplication for that matter) is not defined for two encrypted values in a homomorphic additive-encryption scheme. Also the evaluation of the normalization function $v(\mathbf{Y}_{i-1})$ (e.g., (18)) may not be computable on encrypted values.

Consequently, we will have to use the original signal/coefficient values $(x_{i-1}, x_{i-2}, \dots, x_{i-L})$, which will have the same statistics as $(y_{i-1}, y_{i-2}, \dots, y_{i-L})$ for sufficiently large value of L . Experimental results have shown that an appropriate value of L is 25. For this value of L , the detection results, using normalization on $v(\mathbf{X}_{i-1})$, are sufficiently close to the results based on normalization using $v(\mathbf{Y}_{i-1})$.

Since RDM applies QIM on the ratio $x_i/v(\mathbf{X}_{i-1})$, attention should be paid to the integer rounding process. Since $x_i/v(\mathbf{X}_{i-1})$ will usually be around (the real number) 1.0, the rounding to an integer will almost always yield (the integer) 1, introducing unacceptably large watermarking distortions. Therefore, the scaling of the ratio with a factor c becomes essential in RDM. Furthermore, after quantization of the ratio $x_i/v(\mathbf{X}_{i-1})$, the result needs to be multiplied with $v(\mathbf{X}_{i-1})$. Thanks to the homomorphic property, this can be carried out by an exponentiation in modulo arithmetic with $v(\mathbf{X}_{i-1})$ in the encrypted domain. To this end, obviously $v(\mathbf{X}_{i-1})$ has to be an integer, requiring another rounding step. In case this rounding effect is severe, another scaling can be carried out on $v(\mathbf{X}_{i-1})$. Since, in our experiments, this effect showed to be negligible, we do not consider scaling of $v(\mathbf{X}_{i-1})$ itself. We denote the rounded value of $v(\mathbf{X}_{i-1})$ by $v_{\text{int}}(\mathbf{X}_{i-1})$.

Using again the notation d_i for the uniformly distributed dither, the RDM-embedding equations become

$$E(t_i) = \begin{cases} E\left(c \cdot \left(Q_{2\Delta}\left(\frac{x_i}{v_{\text{int}}(\mathbf{X}_{i-1})} + d_i\right) - d_i\right)\right) \times E(w_j)^\Delta, \\ \quad \text{if } \left(\frac{c \cdot x_i}{v_{\text{int}}(\mathbf{X}_{i-1})}\right) \geq Q_{2\Delta}\left(\frac{c \cdot x_i}{v_{\text{int}}(\mathbf{X}_{i-1})}\right), \\ E\left(c \cdot \left(Q_{2\Delta}\left(\frac{x_i}{v_{\text{int}}(\mathbf{X}_{i-1})} + d_i\right) - d_i\right)\right) \times \left(E(w_j)^\Delta\right)^{-1}, \\ \quad \text{if } \left(\frac{c \cdot x_i}{v_{\text{int}}(\mathbf{X}_{i-1})}\right) < Q_{2\Delta}\left(\frac{c \cdot x_i}{v_{\text{int}}(\mathbf{X}_{i-1})}\right), \end{cases}$$

$$E(y_i) = E(t_i)^{v_{\text{int}}(\mathbf{X}_{i-1})}. \quad (20)$$

With the above scheme, we have succeeded in adapting the RDM watermarking scheme, one of the most recent QIM watermarking approaches, to the constraints set by the Kuribayashi-Tanaka protocol.

5. EXPERIMENTAL VALIDATION

In this section, we experimentally compare the plaintext versions of the SD-QIM, DC-QIM, and RDM watermarking schemes with the proposed version based on the Kuribayashi-Tanaka fingerprinting protocol. The buyer's

TABLE 2: Table of parameters.

Algorithm	Scaling factor	Quantization step size	Noise
SD-QIM	$c = 1, 2, 5, 10, 100$	$\Delta = k$ for $k, 1 \leq k \leq 20$	
DC-QIM	$c = 1, 10, 100$	$\Delta = 5k$ for $k, 1 \leq k \leq 20$	$\sigma_n = 15$
	$c = 10$	$\Delta = k$ for $k, 1 \leq k \leq 20$	$\sigma_n = 15$
RDM	$c = 100$	$\Delta = k$ for $k, 1 \leq k \leq 20$	
	$c = 1000$	$\Delta = 8k$ for $k, 1 \leq k \leq 20$	
	$c = 10.000$	$\Delta = 75k$ for $k, 1 \leq k \leq 20$	

identity information will be embedded into the DC DCT coefficients of 8×8 blocks. Per image, we embed 64 bits of identity information into 64 DC DCT coefficients that are pseudorandomly selected based on a secret key only known to the merchant. In all experiments, we use the 256×256 pixels gray-valued Lena and Baboon images. Because of runtime efficiency and the availability of the necessary proofs, we selected the Okamoto-Uchiyama cryptosystem for all experiments as in [5]. The Okamoto-Uchiyama cryptosystem has a smaller encryption rate compared to (generalized versions of) Paillier because of a smaller message space for the same security level. However, as signal values are usually sampled with 8 bit precision, a smaller message space is not a problem for our application, while the ciphertext size is reduced with the Okamoto-Uchiyama cryptosystem, resulting in lower overall computational complexity.

We not only compare the performance of the plaintext and ciphertext versions of the SD-QIM, DC-QIM, and RDM watermarking schemes, but we also evaluate the effect of integer rounding and the scaling parameter c on the performance. In our graphs, each point shown is based on 100 measurements, and each measurement is a complete, new iteration of the Kuribayashi-Tanaka protocol. A table of parameters¹ for algorithms can be found in Table 2.

5.1. Subtractive dither QIM

An important performance measure of a watermarking scheme is the bit-error rate (BER) of the watermark detector as a function of the strength of embedding the watermark. The BER is a measure that quantifies the probability P_e of incorrectly detecting a single bit of information. Usually, the buyer's identity information contains some form of channel coding so that the buyer's identity can still be retrieved even if a few bits are incorrectly detected from the fingerprinted image, this is further discussed in Section 6.

In order to measure the distortion that the watermark introduces into the host signal, we use the document-to-watermark ratio (DWR):

$$\text{DWR} = 10 \log_{10} \left(\frac{\sigma_x^2}{\sigma_w^2} \right) \text{ (dB)}. \quad (21)$$

¹ The codes for the implementation can be found in <http://ict.ewi.tudelft.nl>.

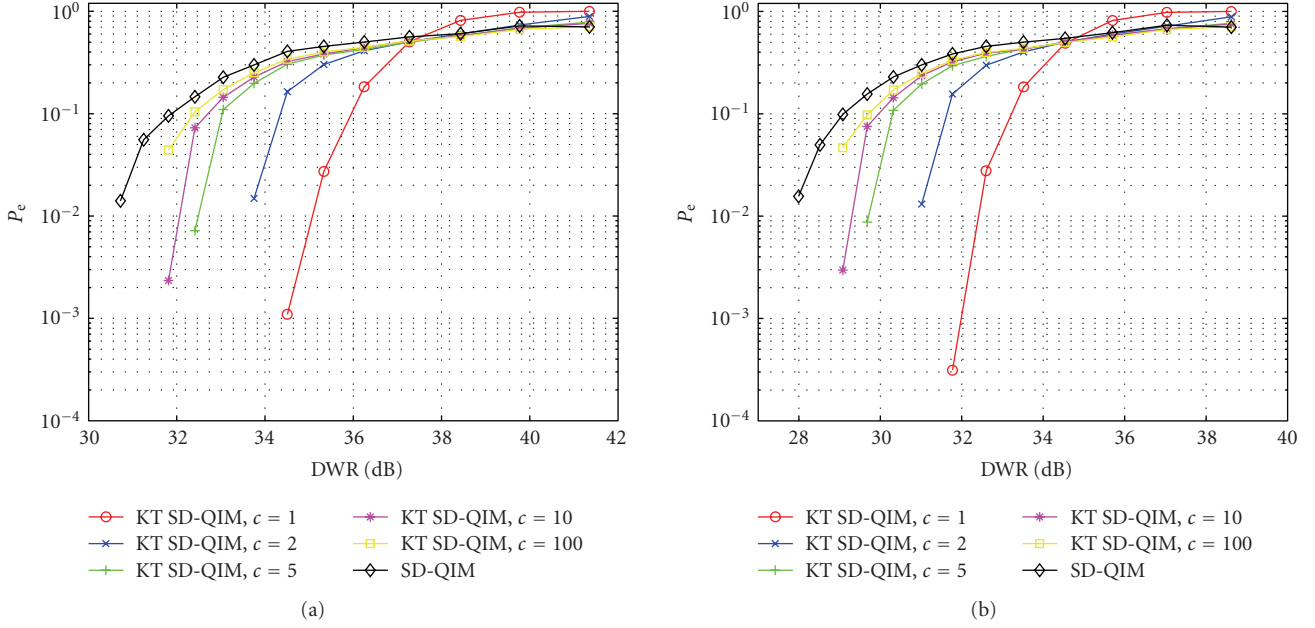


FIGURE 5: SD-QIM bit error rate (BER) P_e as a function of the document-to-watermark ratio (DWR) for the original SD-QIM scheme and KT SD-QIM with different scaling factors $c = 1, 2, 5, 10,$ and 100 for (a) Lena and (b) Baboon images.

Here, σ_x^2 is the variance of the data, into which the watermark is embedded, which, in our case, are the DC DCT coefficients of 8×8 blocks. Further, σ_w^2 is the variance of the distortion caused by the embedded watermark. Following [6], we equate $\sigma_w^2 = \Delta^2/3$. The objective, a watermarking scheme, is to have a low BER with a high DWR. The proper values for the DWR and thus Δ is application and data dependent. In this paper, we are not concerned with selecting a suitable value of Δ . We rather study the behavior of the BER as a function of the DWR for the plaintext and Kuribayashi-Tanaka versions of the SD-QIM watermarking scheme.

Figure 5 shows the BER-DWR relation for the two versions of the SD-QIM algorithm. The performance of the Kuribayashi-Tanaka version of the SD-QIM (KT SD-QIM) watermarking scheme is shown for several values of the scaling factor c . Although there is no deliberate attack performed on the watermark, the inverse DCT transform, and consequential rounding to 8 bit pixel values introduces a distortion into the fingerprinted signal. The robustness of the watermarking scheme is sufficient, however, to result in no-bit errors at a DWR of 31–34 dB. A peculiar effect is the increased robustness of the heavily rounded (i.e., scaling factor $c = 1$) KT SD-QIM compared to the original watermarking scheme. We believe that this behavior is caused by the distorting effect of the (inverse) DCT transform. By increasing the scaling factor c , we can approximate the performance of the original SD-QIM. The performance is already closely approximated with $c = 100$ in this instance, but in general, the application, the data, and the implementation of the DCT will determine which value of c is required to approximate the performance of the plaintext SD-QIM scheme.

5.2. Distortion-Compensated QIM

Figure 5 showed the BER in a scenario without any explicit attacks on the watermark. Distortion-compensated QIM can be used to provide optimal robustness against additive noise attacks. Therefore, we will show the performance of the Kuribayashi-Tanaka adaptation of DC-QIM and compare it with the original DC-QIM and the previously discussed SD-QIM. A measure of the amount of noise introduced relative to the strength of the watermark is the watermark-to-noise ratio (WNR):

$$\text{WNR} = 10 \log_{10} \left(\frac{\sigma_w^2}{\sigma_n^2} \right) \text{ (dB)}. \quad (22)$$

Here, σ_n^2 is the variance of the additive zero-mean Gaussian noise that the attacker adds to the fingerprinted content. The value of α is chosen according to (14) so that the DC-QIM scheme is tuned for a specific additive noise-variance level. In all our experiments, we use $\sigma_n = 15$ and change the value of $\Delta = \sqrt{3}\sigma_w$ as so to obtain a varying WNR.

Figure 6 shows the BER-WNR relation for SD-QIM and DC-QIM. We choose to fix the amount of additive noise instead of the DWR because we are interested in the effect the scaling factor c has on the required embedding strength (i.e., value of Δ and thus the watermark power) and not a variable amount of additive noise. Therefore, Figure 6 cannot be easily compared to other literature on watermark robustness. As in our previous experiment, the watermark distortion is calculated using the expression $\sigma_w^2 = \Delta^2/3$ [6].

As can be observed, the performance of the DC-QIM is better than SD-QIM with additive noise, which is in accordance with [6]. We are mostly concerned with the comparison of the original version of the DC-QIM scheme and the

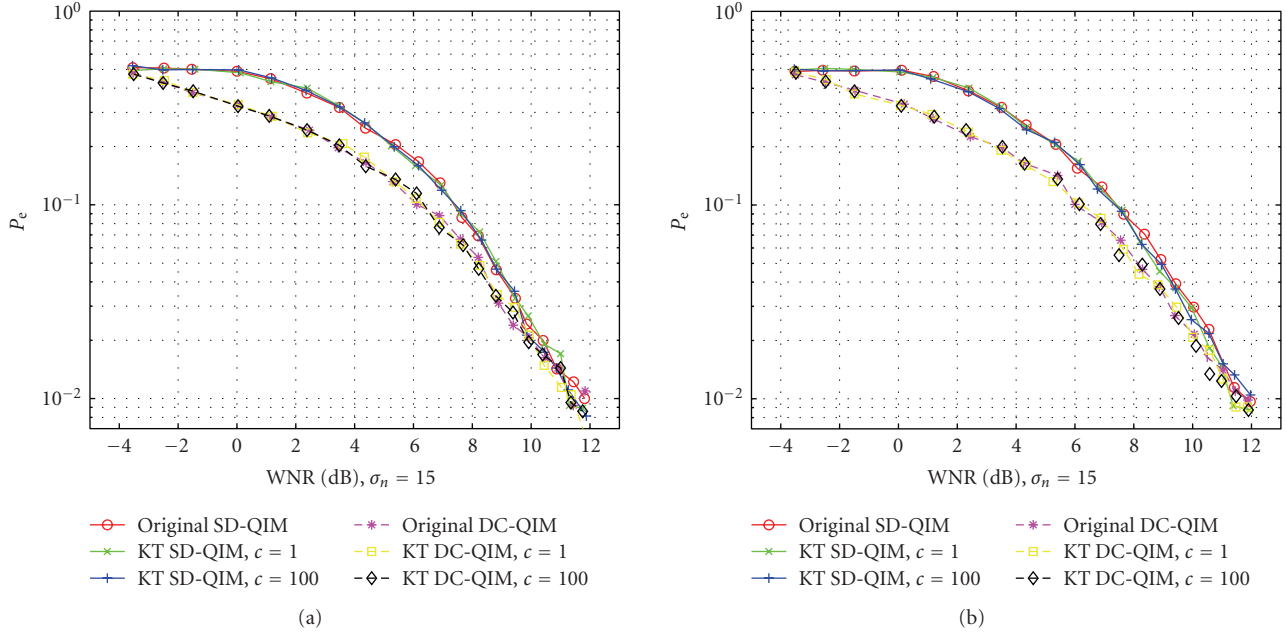


FIGURE 6: SD-QIM and DC-QIM bit error rate (BER) as a function of the watermark-to-noise ratio (WNR) with additive noise ($\sigma_n = 15$) for the original SD-QIM and DC-QIM schemes and the KT SD-QIM and DC-QIM schemes with different scaling factors c for (a) Lena and (b) Baboon images.

Kuribayashi-Tanaka adaptation of DC-QIM. As expected, the performance of the original DC-QIM scheme and the Kuribayashi-Tanaka adaptation of DC-QIM (KT DC-QIM) differ very little. Also the scaling factor c has little effect on the BER. This can be explained by the fact that the additive noise dominates the errors caused by the integer rounding.

5.3. Rational dither modulation

Unlike the previous two watermarking schemes, rational dither modulation (RDM) depends on a sufficiently large scaling factor c in order to achieve a quantization coarseness Δ lower than 1. The scaling factor c determines the possible resolution of Δ . We are interested to see which resolution is required in order to achieve good performance. Although the results depend on the data and the strength of the added noise, the trend of these results will be observed for other cases and data as well because the signal coefficients x_i are normalized before embedding.

Figure 7 shows the bit error rate (BER) performance of RDM as a function of the watermark-to-noise ratio (WNR) for the plain text and Kuribayashi-Tanaka versions of RDM. The different curves reflect different values for the scaling factor c . Because of the complexity of the analytical expression of the watermark distortion σ_w^2 in [7], we measured the watermark distortion directly from the data.

Figure 7 shows that the value of the scaling factor c determines the points of the P_e -WNR curve, which are attainable by the Kuribayashi-Tanaka RDM scheme. With a scaling factor $c = 10$, only WNRs with 12 dB or higher are reachable (see “KT RDM, $c = 10$ ” curve in Figure 7, which starts at 12 dB), allowing for very little flexibility in choosing the op-

timal embedding strength for a specific application. A scaling factor of 100 performs much better, but 1000 approximates the original RDM closely.

Besides the equivalent robustness to additive-noise attacks of RDM compared to SD-QIM, RDM is robust against amplitude-scaling attacks. Figure 8 shows the robustness of SD-QIM, DC-QIM, and RDM to a performed amplitude-scaling attack. SD-QIM and DC-QIM, show a high vulnerability against amplitude-scaling attacks. At a small gain factor ρ of 1.05, approximately 50 percent of the buyer’s identifying information cannot be retrieved correctly, while RDM is robust throughout the whole range for the gain factor. Although theoretically RDM should not be at all affected by an amplitude-scaling attack, some bit errors start to show up at gain factors larger than 1.06. These are inherent to the 8 bit data-representation format, which easily overflows for large gain factors.

6. SECURITY ASPECTS OF BUYER IDENTITY

As fingerprint detection is a signal processing operation, detected fingerprints will usually be distorted even without attacks on the fingerprint by a malicious buyer, as discussed in Section 4. The fingerprint can, for instance, be distorted by perfectly legitimate signal-processing operations such as compression, the obligatory inverse DCT, and consequential rounding. In this scenario, the merchant would normally not be able to present a perfectly retrieved buyer id. The registration center could accept merchant buyer id submissions, which are similar to a correct buyer id. However, the security of the buyer depends on the inability of the merchant to guess a correct buyer id. To allow the merchant to submit similar

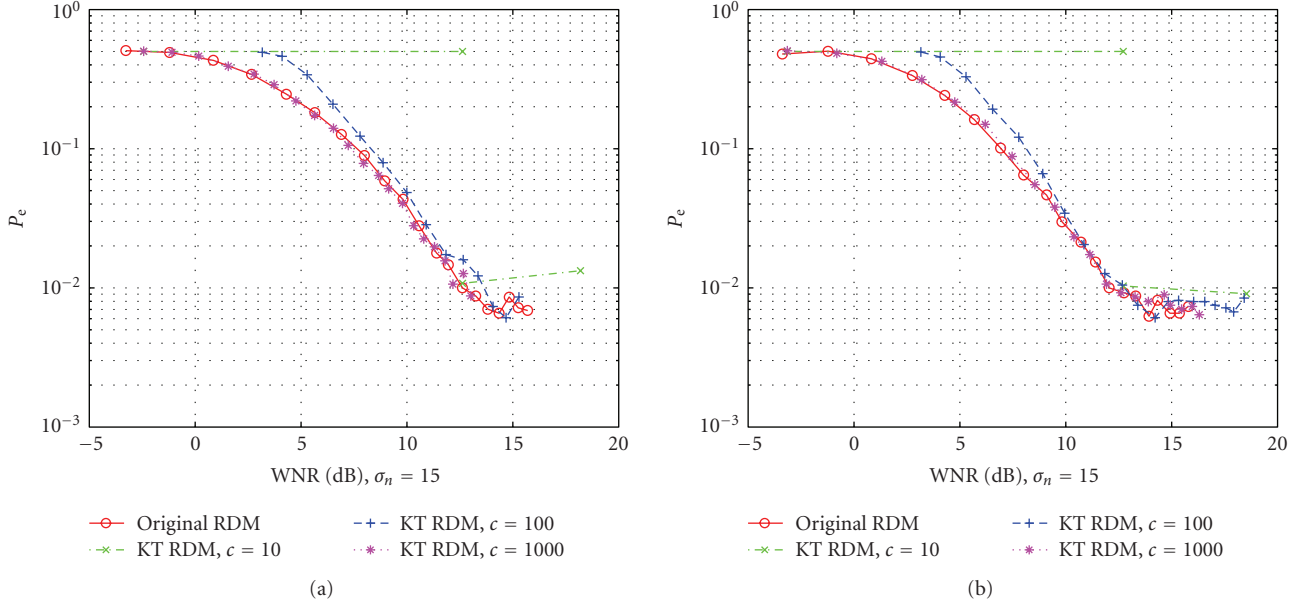


FIGURE 7: RDM bit error rate (BER) as a function of the watermark-to-noise ratio (WNR) with additive noise ($\sigma_n = 15$) for the original RDM scheme and KT RDM scheme with different scaling factors c for (a) Lena and (b) Baboon images.

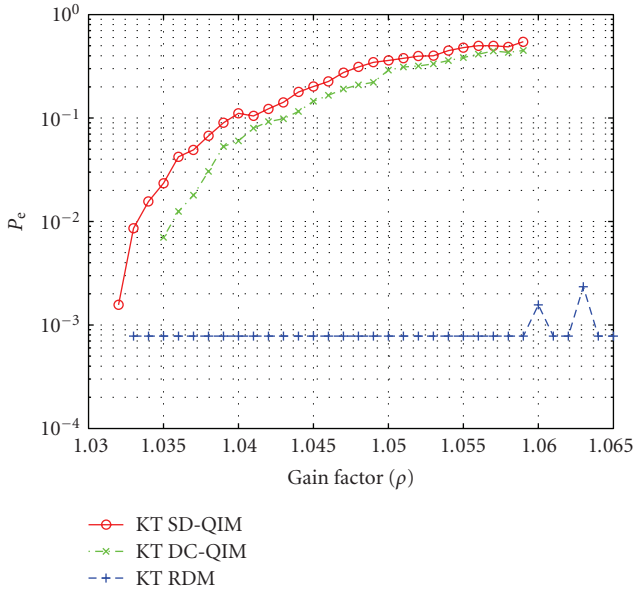


FIGURE 8: KT bit error rate (BER) as a function of the gain factor (ρ) for KT SD-QIM, KT DC-QIM and KT RDM schemes with $c = 1000$. The DWR is fixed to 7.1 dB. Datapoints below a BER of 10^{-3} are plotted for visualization, but in reality 0.

buyer ids and for the registration center to accept these would thus harm the buyer’s security.

By letting the registration center extend the buyer id with a forward-error-correcting scheme, the merchant can compensate for a small and fixed maximum number of bit errors in the buyer id. This is of course equivalent to increasing the size of the buyer id and allowing for a small number of bit er-

rors at the registration center. This approach has the advantage that it moves the computational complexity of the error correction from the registration center to the merchant.

There is a choice to be made concerning the locations of the embedding positions for each buyer. The embedding positions can be changed for each buyer, but this would not provide any real benefits to the robustness of the total fingerprinting scheme other than that colluding buyers would have to compare their individual fingerprinted version with a number of other versions in order to detect the embedding locations. If the embedding locations are identical for each fingerprinted copy, buyers who have located these embedding positions could publish these, and all buyers could then remove the fingerprint from their copy. Using unique embedding positions for each buyer has, however, a big disadvantage upon detection. As with any fingerprinting scheme, the merchant cannot know the used embedding positions before detection, as the detection procedure is the sole method to discriminate between copies. The unavailability of the embedding positions prevents the merchant from detecting the buyer id, resulting in a deadlock. In order to break this deadlock, the merchant could estimate the embedding positions by using a nonblind detection procedure (e.g., subtract the original image from the encountered image and thus find the most likely candidate embedding locations, as they will show up to have a high difference to the original signal) or by embedding a pilot signal to identify the used embedding positions. However, this would be ineffective for heavily attacked copies, which are heavily distorted by attacks. Another way to retrieve the correct buyer id is to let the merchant detect for all possible embedding locations and use a (soft) error-correction scheme to determine the most likely buyer id, based on the distance, the detected id is from a valid code-word in the used error-correction scheme. This, however,

makes the detection procedure linear in complexity related to the number of buyers as it has to be performed for each used combination of embedding positions.

Although dithering prevents an individual buyer to detect the embedding positions, a coalition of buyers can collude to find them. By comparing different fingerprinted copies, the coalition can locate the differing samples and coefficients and, as the fingerprint embedding is the predominant cause of these differing samples, consequently, the embedding positions. This vulnerability can be eliminated by constructing the buyer's *ids* through the scheme of Boneh and Shaw [17], making them collusion secure. The collusion security of the scheme of Boneh and Shaw depends on generating buyer *ids* such that they have a number of identical bits w_j for any colluding coalition of c buyers. Because these buyer *id* bits are identical, the coalition is not able to detect these embedded bits by comparing their individually fingerprinted copies. This does, however, require that the embedding positions are identical for each fingerprinted copy. Because the embedding positions for these bits cannot be determined, they are safe from targeted attacks and can therefore be detected correctly by the merchant even after the attack by the colluding buyer coalition. Constructing such a collusion-secure code for a large coalition constitutes a large increase in the buyer *id* length. As shown in [17], the length is equal to $O(c^4 \log(N/e) \log(1/e))$, where c is the number of colluding buyers, N is the total number of buyers, and e is the probability that the cheating buyer cannot be retrieved after a collusion attack. Because of the anonymity of the embedding procedure, the registration center will have to generate the collusion-secure buyer *ids* as this will be the only person the merchant trusts to generate a valid buyer *id*.

7. CONCLUSION

In conventional fingerprinting schemes, the buyer's identity is known to the merchant during embedding. This knowledge can be easily abused by a malicious merchant by creating fingerprinted copies containing this identity information without the buyer's consent. After distribution, the merchant can claim a license violation for this specific buyer. To deal with this problem, Kuribayashi and Tanaka proposed a reasonably efficient solution in [5] based on embedding the buyer identification information using additive homomorphic encryption schemes. The problem of the proposed protocol in [5] is the vulnerability of the underlying basic QIM watermarking scheme, which is fragile to simple attacks like amplitude scaling and allows for the detection of the embedding positions. Therefore, we have proposed to adapt DC-QIM and RDM techniques to the anonymous fingerprinting scheme of Kuribayashi and Tanaka.

We have adapted DC-QIM and RDM techniques, which hide the embedding locations, unlike basic QIM, because they are based on SD-QIM. They perform provably equivalent (RDM) or better (DC-QIM) than the watermarking scheme in the original work against additive-noise attacks. Furthermore, RDM provides robustness to amplitude-scaling attacks which is a major drawback of the basic QIM scheme used in [5].

Although rounding errors can be made arbitrarily small through the use of scaling factors, the practical need, as shown in the experiments, is small. As integer quantization step sizes have to be used because of the homomorphic encryption scheme, the distortion introduced by the fingerprint embedding is usually larger than the distortion introduced by integer rounding. As a consequence, rounding with a scaling factor of one (i.e., no scaling) already has acceptable performance. The scaling factor has its use, however, in increasing the effective quantizer resolution. Although this is of limited use for signals with a relatively large value range, it is essential for signals with a small value range, as is the case for RDM after normalization.

Due to attacks on the digital content or transmission errors, the identity information of the buyer can be extracted with bit errors. In that case, using error-correction codes can improve the abilities of the merchant to recover the identity information. By letting the registration center select the buyer identity information, we can incorporate these error-correction capabilities or even provide a collusion-secure fingerprinting scheme. This greatly increases the embedded buyer's identification information and the complexity of constructing a valid identity at the registration center. Although this might not be practical in real applications, it provides a theoretical solution to the problem of collusion.

By adapting the DC-QIM and RDM watermarking schemes to the anonymous fingerprinting protocol of Kuribayashi and Tanaka, we increased the robustness of the embedded fingerprints, while preserving the anonymity of the fingerprinting protocol. Consequently, the buyer's ability to successfully attack embedded fingerprints is reduced, increasing the deterrence to the illegal redistribution of digital content.

ACKNOWLEDGMENTS

The work described in this paper has been supported in part by the European Commission through the IST Programme under Contract no. 034238-SPEED. The information in this document reflects only the authors' views, is provided as is and no guarantee or warranty is given that the information is fit for any particular purpose. The user thereof uses the information at its sole risk and liability.

REFERENCES

- [1] N. Memon and P. Wong, "A buyer-seller watermarking protocol," *IEEE Transactions on Image Processing*, vol. 10, no. 4, pp. 643–649, 2001.
- [2] B. Pfitzmann and M. Waidner, "Anonymous fingerprinting," in *International Conference on the Theory and Application of Cryptographic Techniques (EUROCRYPT '97)*, vol. 1233, pp. 88–102, Konstanz, Germany, May 1997.
- [3] B. Pfitzmann and A.-R. Sadeghi, "Coin-based anonymous fingerprinting," in *International Conference on the Theory and Application of Cryptographic Techniques (EUROCRYPT '99)*, vol. 1592, pp. 150–164, Prague, Czech Republic, May 1999.

- [4] B. Pfitzmann and A.-R. Sadeghi, "Anonymous fingerprinting with direct non-repudiation," in *Proceedings of the 6th International Conference on the Theory and Application of Cryptology and Information Security (ASIACRYPT '00)*, vol. 1976, pp. 401–414, Kyoto, Japan, December 2000.
- [5] M. Kuribayashi and H. Tanaka, "Fingerprinting protocol for images based on additive homomorphic property," *IEEE Transactions on Image Processing*, vol. 14, no. 12, pp. 2129–2139, 2005.
- [6] B. Chen and G. W. Wornell, "Quantization index modulation: a class of provably good methods for digital watermarking and information embedding," *IEEE Transactions on Information Theory*, vol. 47, no. 4, pp. 1423–1443, 2001.
- [7] F. Perez-Gonzalez, C. Mosquera, M. Barni, and A. Abrardo, "Rational dither modulation: a high-rate data-hiding method invariant to gain attacks," *IEEE Transactions on Signal Processing*, vol. 53, no. 10, part 2, pp. 3960–3975, 2005.
- [8] T. Okamoto and S. Uchiyama, "A new public-key cryptosystem as secure as factoring," in *International Conference on the Theory and Application of Cryptographic Techniques (EUROCRYPT '98)*, vol. 1403, pp. 308–318, Espoo, Finland, June 1998.
- [9] N. Ahituv, Y. Lapid, and S. Neumann, "Processing encrypted data," *Communications of the ACM*, vol. 30, no. 9, pp. 777–780, 1987.
- [10] P. Paillier, "Public-Key Cryptosystems Based on Composite Degree Residuosity Classes," in *International Conference on the Theory and Application of Cryptographic Techniques (EUROCRYPT '99)*, vol. 1592 of *Lecture Notes in Computer Science*, pp. 223–238, Springer, Prague, Czech Republic, May 1999.
- [11] S. Goldwasser and S. Micali, "Probabilistic encryption," *Journal of Computer and System Sciences*, vol. 28, no. 2, pp. 270–299, 1984.
- [12] R. L. Rivest, A. Shamir, and L. Adleman, "A method for obtaining digital signatures and public-key cryptosystems," *Communications of the ACM*, vol. 21, no. 2, pp. 120–126, 1978.
- [13] T. ElGamal, "A public key cryptosystem and a signature scheme based on discrete logarithms," *IEEE Transactions on Information Theory*, vol. 31, no. 4, pp. 469–472, 1986.
- [14] I. D. Shterev and R. L. Lagendijk, "Amplitude scale estimation for quantization-based watermarking," *IEEE Transactions on Signal Processing*, vol. 54, no. 11, pp. 4146–4155, 2006.
- [15] M. Costa, "Writing on dirty paper," *IEEE Transactions on Information Theory*, vol. 29, no. 3, pp. 439–441, 1983.
- [16] F. Bartolini, M. Barni, and A. Piva, "Performance analysis of ST-DM watermarking in presence of nonadditive attacks," *IEEE Transactions on Signal Processing*, vol. 52, no. 10, pp. 2965–2974, 2004.
- [17] D. Boneh and J. Shaw, "Collusion-secure fingerprinting for digital data," *IEEE Transactions on Information Theory*, vol. 44, no. 5, pp. 1897–1905, 1998.