A Symbolic Transform Attack on Lightweight Encryption Based on Wavelet Filter Parameterization[*]

Dominik Engel, Rade Kutil and Andreas Uhl
Department of Computer Sciences
University of Salzburg, Austria
{dengel,rkutil,uhl}@cosy.sbg.ac.at

ABSTRACT

We present a family of attacks on lightweight encryption schemes for visual data that rely on wavelet filter parameterizations to provide security. All of the attacks construct a symbolic representation of the inverse wavelet transform. We show that this representation can be used in ciphertext-only attacks, known-plaintext attacks and in attacks in which some information on the plaintext is available. We investigate the success and feasibility of each of these attacks, and conclude that the presented type of attacks poses a principal problem for lightweight encryption schemes that rely on the parameterization of a (linear) transform.

Categories and Subject Descriptors

I.4.2 [**Image Processing and Computer Vision**]: Compression (Coding); E.3 [**Data**]: Data Encryption

General Terms

Security

Keywords

attack, ciphertext-only, JPEG2000, known-plaintext, secret frequency domain, secret parameterized wavelet filters

1. INTRODUCTION

For securing multimedia data – like any other type of data – full encryption with a traditional cipher, such as AES, remains the most secure option. However, in the area of multimedia, many applications do not require the level of security this option provides, and seek a trade-off in security to enable other requirements, including low processing demands, retaining bitstream compliance and scalability, and the support for increased functionality, such as transparent encryption [10]. Lightweight encryption aims at striking a balance between security and these other requirements.

Some recent propositions for lightweight encryption make use of parameterized wavelet transforms to provide security. In this paper, we present a family of attacks on such security schemes. All of the discussed attacks rely on a symbolic representation of the inverse wavelet transform that is constructed for each pixel of the reconstructed image. Depending on the information available to the attacker, this representation can be used in a variety of attacks. We discuss these attacks in the context of parameterized wavelet lifting, as this seems to be the most promising parameterization technique from an encryption point of view, but the principle is applicable to any security scheme that uses wavelet parameterization, or a parameterization of any (linear) transform, to provide lightweight security. In this respect, our main goal here is to discuss the basic feasibility of the symbolic inverse wavelet transform for attacks – sophisticated implementations of the attacks against individual parameterization schemes are not our focus here.

2. RELATED WORK

Recently, there have been a number of propositions for lightweight encryption of wavelet-coded data. One group of these propositions operate in a bitstream-oriented manner, i.e. they selectively encrypt parts of the final bitstream to implement access control. For JPEG2000 bitstreams, for example, [6] propose scrambling the coefficient signs in codeblocks.

Another group of contributions aim at using the degrees of freedom in the wavelet transform to construct a unique frequency domain for the transformation step. By keeping the frequency domain secret, these approaches provide lightweight security. This procedure can be seen as a form of header encryption, as only the information pertaining to the frequency domain needs to be encrypted, the rest of the data remains in plaintext. For wavelet-coded data, there are two groups of approaches: one group uses parameterized wavelet filters to construct the frequency domain, another group uses secret wavelet packet decompositions. In this work we focus on the security of the former group.

Different parameterization schemes have been employed for various security techniques in a number of wavelet-based codecs. Apart from lightweight encryption, which we discuss below, wavelet parameterizations are also used in other areas of multimedia security: For increasing the security of watermarking schemes, different parameterizations have been investigated by [11, 8], whereas [12] aim at increasing security for visual hashes.

[*]Funding by the Austrian Science Fund (FWF), project number P15170, and European Network of Excellence ECRYPT is gratefully acknowledged.

MM&Sec'06, September 26–27, 2006, Geneva, Switzerland.
Copyright 2006 ACM 1-59593-493-6/06/0009 ...$5.00.

Some propositions that employ techniques which are in spirit related to parameterized wavelet filters should also be mentioned. In [18] the encryption of the filter choice used for wavelet decomposition is proposed. However, this suggestion remains vague and is not supported by any experiments. [5] introduce the concept of key-dependent basis functions to protect a watermark from hostile attacks. This approach suffers from significant computational complexity. There are also some propositions that use secret Fourier transforms: The embedding of watermarks in an unknown domain is discussed by [2], and [17] suggest to use this technique for encryption of visual data.

3. LIGHTWEIGHT ENCRYPTION WITH PARAMETERIZED WAVELETS

Three wavelet parameterization schemes have been investigated in the context of lightweight encryption: the parameterization for a family of orthogonal wavelets proposed by [15], the parameterization for even and odd length biorthogonal filters proposed by [7], and the lifting parameterization of the CDF 9/7 wavelet proposed by [19]. For the latter parameterization, the original goal was to create a version of the CDF 9/7 wavelet with simpler filter taps, but it turns out that a whole family of wavelets can be constructed. The parameterization is based on a combination of the lifting steps of the CDF 9/7 wavelet [1] with the construction theorem and the perfect reconstruction conditions. The factors in the original lifting scheme are all expressed as functions of one parameter α.

The utility of orthogonal and biorthogonal wavelet parameterization schemes is compared in [16]. The compression performance of both, the investigated orthogonal [15] and biorthogonal [7] wavelet parameterization, is reported as unreliable. Similar observations for the orthogonal parameterization are made by [9] in the context of JPEG2000. The usability of high-dimensional wavelet parameterization, based on the parameterization of orthogonal wavelet filters [15], is investigated by [14]. The authors report that although security is improved by longer filters, the poor compression quality of the orthogonal filters remains a drawback.

The authors of [3] use the biorthogonal lifting parameterization presented by [19] with JPEG2000 and report compression performance that is superior to the other parameterization schemes. (This is not surprising as the family of parameterized wavelet filters is based on the CDF 9/7 biorthogonal wavelet, the standard wavelet for JPEG2000.) Furthermore, the proposed encryption scheme introduces only minimal computational overhead and the only information to be encrypted in the header is the parameter α. A shortcoming of the scheme is that due to the limited range of α in which the filters exhibit sufficient variance, the keyspace is relatively small. To overcome this, the used filters are varied for each level of the wavelet transform (non-stationary variation) and for each direction of the wavelet transform (inhomogeneous variation). In [4], the authors further increase keyspace size by a combination of parameterized wavelet filters with the randomized wavelet packet decompositions, at the cost of introducing computational complexity in the transform step. Most of the attacks discussed here can theoretically be applied to any of these extensions, but will increase in computational demands and decrease in precision.

4. ATTACKS

Security schemes that rely on wavelet parameterizations use the degrees of freedom that the wavelet transform provides to produce filters that are suitable for both, providing security and achieving good image compression. Thereby the fact that the wavelet transform is a linear transformation poses a threat for security. Linear transforms are in principle not well suited for keeping information secret. Note that the symbolic attack does not presume a linear transform. It would also work with non-linear transforms that continuously depend on a finite number of parameters. However, with a linear transform it is more likely that the symbolic expressions can be contracted and therefore evaluated faster than when the complete transform has to be performed.

Note that an attacker does not have to obtain the exact parameter value, an approximation is sufficient to yield an image with little distortion. How close the attack value has to be to the encoding value depends on the used parameterization and the used discretization.

The attacks discussed here are based on the symbolical computation of the inverse wavelet transform. Let I be a grayscale image of size $n \times n$ pixels, with luminance values represented as a vector with elements $I_i, i = 0, \ldots, n^2 - 1$, where I_i is the luminance value of the pixel at position (i mod $n, \lfloor \frac{i}{n} \rfloor$). Assume I is decomposed with a parameterized wavelet transformation that depends on m parameters $\alpha_j, j = 0, \ldots, m-1$. The inverse transformation with the correct set of parameters will reconstruct the original image I (assuming, for sake of simplicity, a reversible transform and lossless coding).

An attacker, who does not know the values of α_j, can build a symbolic expression for each pixel value in the reconstructed image containing the necessary operations for the inverse transformation. The resulting term will depend on the values of some of the transform coefficients $C_i, i = 0, \ldots, n^2 - 1$, all of which are known to the attacker. The only unknowns are formed by the parameters of the wavelet transformation, α_j. By performing a full symbolic inverse wavelet transformation, the attacker can construct a complete symbolic description of the operations necessary to reconstruct I. We illustrate this procedure with an example.

For the parameterization of the CDF 9/7 wavelet, the terms of the symbolic attack become relatively complex. For illustration we therefore construct a parameterized version of the simple Haar wavelet. The lifting steps for the forward transform with the Haar wavelet can be written as follows [1]:

$$s_l^{(0)} = x_{2l} \tag{1}$$

$$d_l^{(0)} = x_{2l+1} \tag{2}$$

$$d_l = d_l^{(0)} - s_l^{(0)} \tag{3}$$

$$s_l = s_l^{(0)} + \frac{1}{2} d_l \tag{4}$$

with the inverse transform written as:

$$s_l^{(0)} = s_l - \frac{1}{2} d_l \tag{5}$$

$$d_l^{(0)} = d_l + s_l^{(0)} \tag{6}$$

$$x_{2l+1} = d_l^{(0)} \tag{7}$$

$$x_{2l} = s_l^{(0)}. \tag{8}$$

To construct a simple parameterized version of the Haar wavelet, we change the forward prediction step to

$$d_l \quad = \quad d_l^{(0)} - \alpha s_l^{(0)}. \tag{9}$$

Accordingly, the forward update step is given by:

$$s_l \quad = \quad \frac{1}{2}\left((1+\alpha)s_l^{(0)} + d_l\right). \tag{10}$$

The prediction and update steps of the inverse transform are given by:

$$s_l^{(0)} \quad = \quad \frac{1}{1+\alpha}(2s_l - d_l) \tag{11}$$

$$d_l^{(0)} \quad = \quad d_l + \alpha s_l^{(0)}. \tag{12}$$

For $\alpha = 1$, the original Haar wavelet is obtained. We denote a horizontal transformation that is followed by vertical transformation by two letters. For example, ds refers to the subband that contains the lowpass transform coefficients for horizontal decomposition and the highpass subbands of the subsequent vertical decomposition.

Imagine that this parameterization is used in a lightweight encryption scheme for an image I of size n^2. For this purpose, α_e, the parameter value used for encryption, is chosen randomly from the range of admissible values. This range would first have to be determined, based on compression performance. (Of course, the Haar filter is not well suited for image compression, and such a parameterization even less.) For our example we assume $\alpha \in [0.5, 3]$. After transformation with α_e, reconstruction with a wrong α_d, the parameter value used for decryption, will yield a distorted image. An example for the Haar parameterization is given in Fig. 1(a), for $\alpha_e = 1.1$ and $\alpha_d = 1.5$.

For the proposed attack, the attacker symbolically computes the inverse wavelet transform. Let S be an $n \times n$ matrix to hold the symbolic expressions. Initially each entry $S_{i,j}$ of this matrix is filled with the representation of the corresponding transform coefficient $C_{i,j}$. Then the operations for each step of the wavelet reconstruction that pertain to a certain position (i, j) are recorded symbolically in $S_{i,j}$. After the full inverse wavelet transformation the complete reconstruction of the whole image is described symbolically by S. Each entry $S_{i,j}$ represents the necessary operations to reconstruct the pixel value at position (i, j).

As an example, consider an image of size 4×4 pixels, for which an attacker wants to construct $S_{2,2}$, the symbolic representation of the pixel at position $(2, 2)$. We assume that a one-level wavelet analysis was done in-place. We denote the vertical and horizontal wavelet transformations by operators F_v and F_h, respectively. The following matrix shows which subband coefficients the entries in the symbolic matrix correspond to for the one-level wavelet transformation:

$$C = F_v F_h S = \begin{pmatrix} ss_{0,0} & sd_{0,0} & ss_{0,1} & sd_{0,1} \\ ds_{0,0} & dd_{0,0} & ds_{0,1} & dd_{0,1} \\ ss_{1,0} & sd_{1,0} & ss_{1,1} & sd_{1,1} \\ ds_{1,0} & dd_{1,0} & ds_{1,1} & dd_{1,1} \end{pmatrix}. \tag{13}$$

Initially $F_v F_h S_{2,2}$ contains the representation of the last coefficient in the LL-subband $C_{2,2}$, or $ss_{1,1}$ in the notation used above. After the first vertical analysis transform, $F_h S_{2,2}$ contains the operations necessary to obtain $ss_{1,1}^{(0)}$:

$$ss_{1,1}^{(0)} = \frac{1}{1+\alpha}(2 \cdot ss_{1,1} - ds_{1,1}), \tag{14}$$

which translates to

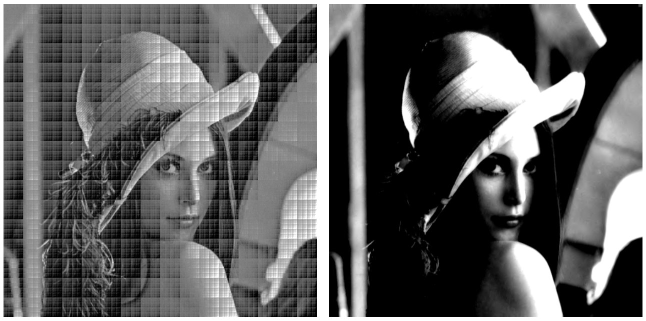

(a) Haar parameterization, $\alpha_e = 1.1$, $\alpha_d = 1.5$

(b) CDF 9/7 param., $\alpha_e = -2.5$, $\alpha_d = -6$

Figure 1: Reconstructions with wrong parameters

$$F_h S_{2,2} = \frac{1}{1+\alpha}(2 \cdot C_{2,2} - C_{3,2}). \tag{15}$$

The reversal of the splitting step makes $F_h S_{2,2}$ contain $s_{2,1}$. The symbolic synthesis step of the horizontal decomposition yields $s_{2,1}^{(0)}$:

$$s_{2,1}^{(0)} = \frac{1}{1+\alpha}(2 \cdot s_{2,1} - d_{2,1}), \tag{16}$$

where $d_{2,1}$ has been constructed in position $F_h S_{2,3}$ in the same way as $F_h S_{2,2}$, and is given by

$$F_h S_{2,3} \quad = \quad \frac{1}{1+\alpha}(2 \cdot sd_{1,1} - dd_{1,1}) \tag{17}$$

$$\quad = \quad \frac{1}{1+\alpha}(2 \cdot C_{2,3} - C_{3,3}). \tag{18}$$

Putting it all together, we obtain the final $S_{2,2}$:

$$S_{2,2} = \frac{1}{(1+\alpha)^2}\left(2 \cdot (2 \cdot C_{2,2} - C_{3,2}) - (2 \cdot C_{2,3} - C_{3,3})\right). \tag{19}$$

Assume that like the image, the matrix is represented as a one-dimensional vector that contains the concatenated lines of the matrix. Let the elements of this vector be denoted by S_i with $i = 0, \ldots, n^2 - 1$. For example, for $n = 32$, the Haar parameterization and a level one horizontal and vertical transformation, the entry S_{24}, is then given by:

$$S_{24} = \frac{2(\frac{\alpha(2C_{12}-C_{29})}{1+\alpha} + C_{28}) - \frac{\alpha(2C_{524}-C_{540})}{1+\alpha} - C_{540}}{1+\alpha} \tag{20}$$

If a 32×32 version of the Lena image has been transformed with $\alpha_e = 1.4$, then by inserting this parameter along with the transform coefficient values into the equation above, S_{24} takes the value of $I_{24} = 129$, the correct pixel value in this position.

As the transform coefficients are known to the attacker, the parameters of the wavelet parameterization are the only unknown part of the symbolic expressions. If no information on the reconstructed image is available, trying to derive the correct settings of the parameters corresponds to a ciphertext-only attack. A situation in which the reconstructed image is fully or partly available corresponds to a full or partial known-plaintext attack.

4.1 Ciphertext-Only

One possible attack on security schemes that are based on parameterized wavelet filters is the use of a function correlated to image quality. By applying such a measure to the symbolic equations, a single term with the filter parameters as the only unknowns can be be obtained. Their values can then be determined by minimizing (or maximizing, depending on the used measure) this term, e.g. by analyzing the first derivative or by employing numerical methods. If there is more than one solution, the possible solutions can be tested very quickly for their usability.

No function exists that can accurately determine if a given image is a natural image. However, as natural images tend to exhibit a certain amount of smoothness, possible indicator functions for image quality could be (inverse) measures of smoothness. A first approach is to use the sum of absolute difference between neighboring pixels and minimize this term. For this purpose the symbolic inverse wavelet transform S is computed for the desired wavelet transformation, as described above. Then the pixel difference p is given by

$$p = \sum_{i=0}^{n^2-2} |S_i - S_{i+1}|. \tag{21}$$

The utility of the absolute difference of neighboring pixel values, as a first candidate for such a function, depends very much on the distortion introduced by the used parameterization. If high frequency noise is introduced for wrong parameter settings, then the pixel difference should yield a good indicator for image quality. For the Haar parameterization, the introduced distortions are indeed situated in the higher frequencies. As an example we transform a version of Lena of size 32 × 32 pixels with $\alpha_e = 1.1$ at a nearly lossless bitrate. Minimizing the symbolic representation of the pixel difference p, we obtain a value of $\alpha_d = 1.17$. A reconstruction with this parameter value yields a PSNR of 40dB.

In the case of biorthogonal lifting parameterization the distortions introduced for wrong parameter settings are of a different kind, as illustrated in Figure 1(b). Rather than introducing noise, wrong parameter values produce more of a blurring and smoothing effect (which is one of the reasons why this particular parameterization is well suited for transparent encryption). Thus, the pixel difference attack fails for this parameterization.

The sample variance as an inverse measure of smoothness is another candidate for an image quality indicator. For a symbolic inverse wavelet transform S, the sample variance s^2 is given by:

$$s^2 = \frac{1}{n^2-1} \sum_{i=0}^{n^2-1} (S_i - m)^2, \tag{22}$$

where m is the mean pixel value.

For the Haar parameterization, the success of this attack is mediocre. For example, we take a transformation with $\alpha_e = 1.1$ for a 32 × 32 pixel version of Lena at nearly lossless coding. The symbolic representation of the variance is computed and by minimizing this term we obtain a value of $\alpha_d = 1.27$. While a reconstruction with this parameter value still yields a PSNR of 32.2 dB for nearly lossless coding, the result is relatively far from the expected target. For the parameterization techniques used in [13], the sample variance as an inverse measure of smoothness yields more successful attacks.

For the biorthogonal lifting parameterization, variance does not provide a strong indicator for image quality either. Again, this is due to the nature of the parameterization: rather than introducing high frequency artefacts for wrong keys, the parameterization exhibits a reduction of energy in the highpass subbands. As pointed out in [3], the correlation is too weak to substantially decrease search complexity for the full quality image. However, even if the correct parameter cannot be determined precisely, a minimization of the variance can lead the attacker in the right direction. As an example, we transformed an 8 × 8 pixel test image with a one level wavelet transformation with $\alpha = -1.6$ and then computed the symbolic equations. A minimization of the variance of the equations for $\alpha \in [-6, -1.2]$ yields the solution $\alpha = -1.48$.

As a possible counter-measure for the ciphertext-only attack, the use of different parameters on different resolutions and decomposition directions, i.e. non-stationary and inhomogeneous variation, can be employed [3]. For security schemes relying on the biorthogonal lifting parameterization, this increases the number of keys to $2l$ where l is the wavelet decomposition depth. Finding the sequence of parameters from a single equation is not possible, and the attack can only serve to limit the range of possible parameters.

For a single parameter, a problem for this attack is that there is no function that is correlated to image quality for all parameterization techniques. However, the fact that so far no suitable predictor for image quality could be found that works for all parameterizations, does not rule out the existence of a measure that can achieve a general correlation to image quality, which would make this attack a more serious threat. In any case, this kind of attack can be used to provide a starting point in a brute-force search and possibly narrow down the range of parameters to be tested.

4.2 Full/Partial Known-Plaintext

The previous version of a symbolic attack is a ciphertext-only attack, and depends on the existence of a predictor function for image quality. Other versions of the attack do not depend on the existence of such a function. They assume that (parts of) the plaintext, or at least some information on the plaintext is available, which is a realistic assumption with lightweight encryption in general, and with transparent encryption in particular.

Full Known-Plaintext. If the full reconstructed plaintext image is available to the attacker, then the attacker can easily determine the used wavelet parameters by solving the equations for the pixels of the reconstructed image. A situation in which the full plaintext is available is rather unlikely, but a scenario could be conceived in which the attacker has obtained the ciphertext of a set of images (all encrypted with the same α_e) and has received the full plaintext of a single image from this set in full resolution and quality as an incentive for buying the whole set.

In the case of the biorthogonal lifting parameterization, the set of equations provides the attacker with ample information to deduce the correct value of α. If inhomogeneous and non-stationary variation as discussed above was used for transformation, instead of one parameter α, the attacker has to derive a sequence of parameters $\alpha_0, \ldots, \alpha_{2l-1}$, where l is the wavelet decomposition depth. Even in this case, for images of sufficient size, the attacker will have enough information to derive the correct sequence. The same is true for parameterizations that depend on more than one parameter.

Note that if the full plaintext image is available, other attacks become feasible as well. In many of the parameterizations, the PSNR for different admissible parameter values is a monotonous function that reaches its peak for the correct parameter value. All the attacker needs to do is to approach the correct reconstruction parameters by iteratively reconstructing the ciphertext coefficients and comparing the PSNR of the resulting image to the PSNR of the available plaintext image.

Partial Plaintext Information. In many cases it is not necessary for the attacker to have access to the full plaintext. Symbolic attacks can be conceived that utilize only minimal information about the plaintext image. This assumption is realistic, as many of the proposed schemes support transparent or sufficient encryption. For sufficient encryption, the scheme tolerates image reconstructions with the wrong parameters which yield a discernible version of the original visual data. The only assertion of the scheme is that these reconstructed versions do not exceed a certain quality threshold. For transparent encryption [10], the scheme does not only tolerate image reconstructions of reduced quality, but uses them as a preview image. Such a preview can be of advantage in "try-and-buy"-scenarios, where they serve as an incentive to acquire the correct key to obtain the full quality version of the visual data. In both cases, a potential attacker can make use of the information provided by the reconstructions obtained with wrong parameters.

Plaintext Pixel Samples. For this attack, the attacker is assumed to have obtained individual pixel samples from the reconstructed image. This can for example be achieved by access to a preview image: the attacker selects homogeneous regions or edges that are likely to have the same luminance values in the full reconstructed image. Equating these values with the symbolic representation of the appropriate pixel position, the attacker can construct a linear system of equations, with the parameters of the wavelet transform as the only unknowns. In the pixel sample attack, the actual or approximate value I_i is known for a number of symbolic terms S_i.

In the case of the Haar parameterization, a single correct pixel value is sufficient to determine the value of α used for encoding. Also for the biorthogonal lifting parameterization, a single correct pixel is sufficient to produce the correct value of α. As an example, we used a 64×64 pixel version of the Lena image. This image was transformed in JPEG2000 at a nearly lossless bitrate of 5 bpp using one level of wavelet decomposition, with the biorthogonal lifting parameterization, $\alpha = -1.6$. Solving the symbolic equation for, e.g. I_{26} with the correct luminance value yields 7 solutions, 6 of which lie in the complex space. The remaining solution is the correct parameter $\alpha = -1.6$.

With inhomogeneous and non-stationary variation and for higher-dimensional parameterizations, more pixel values are needed for an accurate attack, at least as many as the number of filters involved in the parameterization. Depending on the amount and accuracy of information regarding plaintext pixels, the attacker can obtain a more or less accurate solution for the used wavelet parameters.

Average Luminance Value. For this attack, the attacker obtains the average luminance value of the reconstructed plaintext image. A good approximation can usually be obtain if a preview image is available. From the symbolic plaintext equations, a symbolic representation of the average luminance value is constructed, i.e.

$$L = \frac{1}{n^2} \sum_{i=0}^{n^2-1} S_i. \tag{23}$$

Then the attacker equates this expression with the obtained value of the average luminance of the reconstructed plaintext image to obtain the parameter value of the transformation. The accuracy of the derived parameter values depends the accuracy of the obtained average luminance value.

As an example we tested this attack with the 64×64 pixel version of the Lena image, using the biorthogonal lifting parameterization with nearly lossless settings and $\alpha = -1.6$. The reconstructed image has a PSNR of 50.2 dB. The mean pixel value of the reconstructed image before quantization to an integer value is 99.0217. If an attacker obtains the average luminance value of 100 from a preview image, $\alpha = -1.55099$ can be derived. This leads to a reconstructed image quality of 42.2dB. (Note that the fact that this attack works shows that the used parameterization does not strictly conform to the construction theorem, because if it did, the average luminance value would be preserved even for wrong reconstruction filters.)

As the average luminance value only produces a single equation, it cannot be used for parameterizations that use more than one parameter. Inhomogeneous and non-stationary variation therefore make the attack unusable. However, if the average luminance value is available for higher-dimensional parameterization, it can be used in conjunction with the other attacks to reduce their complexity.

4.3 Computational Complexity

In this work we focus on proving that attacks that construct a symbolic inverse transform are successful against encryption schemes that employ wavelet parameterizations to provide lightweight security. To prove this point, the use of a symbolic computation without any optimization at all, i.e. each step of the lifting is computed individually over and over again, is sufficient. We use a straightforward symbolic representation of the inverse wavelet transform. An adapted version of JJ2000 computes the equations of the inverse parameterized wavelet transform and provides output that can be read into Mathematica®, where the equations for the reconstructed pixels are stored in a matrix. Each entry of the matrix corresponds to a pixel in the image and holds the symbolic expression for the reconstruction of this pixel from the transformed image, i.e. the ciphertext.

For an efficient attack, the lifting steps themselves should be represented symbolically and processing should be optimized. That being said, it should be noted that the construction of the symbolic matrix will remain a computationally demanding task, as for each lifting step the symbolic representations have to be handled. However, for specific parameterization, image size and decomposition depth, this matrix has to be computed only once.

The attacks themselves vary in computational demands. For testing we used Mathematica® 5.0 on an AMD Athlon® CPU at 1.66 GHz and 2 GB of RAM. All the timing results pertain to the biorthogonal lifting parameterization with one level of wavelet decomposition. The pixel sample attack is the least demanding and could be performed in 2.7 seconds. The average luminance value attack for a 64×64 pixel image took approximately 1875 seconds. The ciphertext only attack with variance as (weak) image quality indicator is the most demanding attack. We used the `NMinimize` function in Mathematica to minimize the variance of the symbolic matrix for $\alpha \in [-6, -1.2]$, the negative range given in [3], for an 8×8 pixel test image. Even for such a small image, the calculation time for the correct parameter value is 390 seconds. However, these long calculation times even for small images should not deter from the basic applicability of symbolic computation attacks, as they mainly result from the simple implementation that lacks optimization.

5. CONCLUSION

The goal of this work is the investigation of a novel kind of attacks on encryption schemes that use parameterized wavelet transforms to provide lightweight security. We have shown for small images that these attacks are feasible and principally present a threat to such security schemes. A ciphertext-only attack that uses an image quality indicator could potentially be very successful, however, no proper function exists so far that can produce a good indicator for image quality in all investigated parameterizations by only working on the symbolic representations of the reconstructed pixels. A full plaintext attack is very successful on any parameterization scheme as it provides ample information to deduce the parameters used for transformation. It turns out that a much lower fraction of the plaintext or limited information on the plaintext also yields expedient attacks. The success of these depends on the one hand on the complexity and number of parameters of the wavelet parameterization and on the other hand on the accuracy of the available plaintext information. Inhomogeneous and non-stationary variation of the wavelet parameters increase the number of parameters and therefore make the attack more difficult.

Although we have presented some timing results to assess computational demands of these attacks, with the simple implementation and missing optimization in our tests, the presented timing values are not very expressive. There is a lot of room for optimization and it is to be expected that these attacks can be scaled to perform well for larger images. This is the subject of further research.

On a principal note, the attacks presented here show a general problem of lightweight encryption schemes that rely on linear transforms for providing security. Even if these schemes only claim to provide lightweight security, attacks of the style presented here are a potential threat and should be taken into account.

6. REFERENCES

[1] I. Daubechies and W. Sweldens. Factoring wavelet transforms into lifting steps. *J. of Fourier Analysis Applications*, 4(3):245–267, 1998.

[2] I. Djurovic, S. Stankovic, and I. Pitas. Digital watermarking in the fractional fourier transformation domain. *J. of Network and Computer App.*, 24:167–173, 2001.

[3] D. Engel and A. Uhl. Parameterized biorthogonal wavelet lifting for lightweight JPEG 2000 transparent encryption. In *Proc. of ACM Multimedia and Security Workshop, MM-SEC '05*, pages 63–70, New York, NY, USA, Aug. 2005.

[4] D. Engel and A. Uhl. Security enhancement for lightweight JPEG 2000 transparent encryption. In *Proc. of* 5^{th} *Int. Conf. on Information, Communication and Signal Processing, ICICS '05*, pages 1102–1106, Bangkok, Thailand, Dec. 2005.

[5] J. Fridrich, A. C. Baldoza, and R. J. Simard. Robust digital watermarking based on key-dependent basis functions. In D. Aucsmith, editor, *Information hiding: second international workshop*, volume 1525 of *LNCS*, pages 143–157, Portland, OR, USA, Apr. 1998. Springer Verlag, Berlin, Germany.

[6] R. Grosbois, P. Gerbelot, and T. Ebrahimi. Authentication and access control in the JPEG 2000 compressed domain. In A. Tescher, editor, *Applications of Digital Image Processing XXIV*, volume 4472 of *Proc. of SPIE*, pages 95–104, San Diego, CA, USA, July 2001.

[7] F. Hartenstein. Parametrization of discrete finite biorthogonal wavelets with linear phase. In *Proc. of Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP'97)*, Apr. 1997.

[8] J. Huang, J. Hu, D. Huang, and Y. Q. Shi. Improve security of fragile watermarking via parameterized wavelet. In *Proc. of the IEEE Int. Conf. on Image Processing (ICIP'04)*, Singapore, Oct. 2004. IEEE Signal Processing Society.

[9] T. Köckerbauer, M. Kumar, and A. Uhl. Lightweight JPEG 2000 confidentiality for mobile environments. In *Proc. of the IEEE Int. Conf. on Multimedia and Expo, ICME '04*, Taipei, Taiwan, June 2004.

[10] B. M. Macq and J.-J. Quisquater. Cryptology for digital TV broadcasting. *Proc. of the IEEE*, 83(6):944–957, June 1995.

[11] P. Meerwald and A. Uhl. Watermark security via wavelet filter parametrization. In *Proc. of the IEEE Int. Conf. on Image Processing (ICIP'01)*, volume 3, pages 1027–1030, Thessaloniki, Greece, Oct. 2001.

[12] A. Meixner and A. Uhl. Security enhancement of visual hashes through key dependent wavelet transformations. In F. Roli and S. Vitulano, editors, *Image Analysis and Processing - ICIAP 2005*, volume 3617 of *Lecture Notes on Computer Science*, pages 543–550, Cagliari, Italy, Sept. 2005. Springer-Verlag.

[13] A. Pommer and A. Uhl. Selective encryption of wavelet-packet encoded image data — efficiency and security. *ACM Multimedia Systems (Special issue on Multimedia Security)*, 9(3):279–287, 2003.

[14] A. Pommer and A. Uhl. Lightweight protection of visual data using high-dimensional wavelet parametrization. In F. Roli and S. Vitulano, editors, *Image Analysis and Processing - ICIAP 2005*, volume 3617 of *Lecture Notes on Computer Science*, pages 645–652, Cagliari, Italy, Sept. 2005. Springer-Verlag.

[15] J. Schneid and S. Pittner. On the parametrization of the coefficients of dilation equations for compactly supported wavelets. *Computing*, 51:165–173, May 1993.

[16] A. Uhl and A. Pommer. Are parameterised biorthogonal wavelet filters suited (better) for selective encryption? In J. Dittmann and J. Fridrich, editors, *Multimedia and Security Workshop 2004*, pages 100–106, Magdeburg, Germany, Sept. 2004.

[17] G. Unnikrishnan and K. Singh. Double random fractional fourier-domain encoding for optical security. *Optical Engineering*, 39(11):2853–2859, Nov. 2000.

[18] L. Vorwerk, T. Engel, and C. Meinel. A proposal for a combination of compression and encryption. In *Visual Communications and Image Processing 2000*, volume 4067 of *Proc. of SPIE*, pages 694–702, Perth, Australia, June 2000.

[19] G. Zhong, L. Cheng, and H. Chen. A simple 9/7-tap wavelet filter based on lifting scheme. In *Proc. of the IEEE Int. Conf. on Image Processing (ICIP'01)*, pages 249–252, October 2001.