

© SPIE and IS&T. This paper was published by SPIE/IS&T and is made available as an electronic reprint with permission of SPIE and IS&T. One print or electronic copy may be made for personal use only. Systematic or multiple reproduction, distribution to multiple locations via electronic or other means, duplication of any material in this paper for a fee or for commercial purposes, or modification of the content of the paper are prohibited.

An Evaluation of Lightweight JPEG2000 Encryption with Anisotropic Wavelet Packets

Dominik Engel and Andreas Uhl

Department of Computer Sciences, University of Salzburg, Austria

ABSTRACT

In this paper we evaluate a lightweight encryption scheme for JPEG2000 which relies on a secret transform domain constructed with anisotropic wavelet packets. The pseudo-random selection of the bases used for transformation takes compression performance into account, and discards a number of possible bases which lead to poor compression performance. Our main focus in this paper is to answer the important question of how many bases remain to construct the keyspace. In order to determine the trade-off between compression performance and keyspace size, we compare the approach to a method that selects bases from the whole set of anisotropic wavelet packet bases following a pseudo-random uniform distribution. The compression performance of both approaches is compared to get an estimate of the range of compression quality in the set of all bases. We then analytically investigate the number of bases that are discarded for the sake of retaining compression performance in the compression-oriented approach as compared to selection by uniform distribution. Finally, the question of keyspace quality is addressed, i.e. how much similarity between the basis used for analysis and the basis used for synthesis is tolerable from a security point of view and how this affects the lightweight encryption scheme.

Keywords: Anisotropic wavelet packets, JPEG2000, lightweight encryption, number of anisotropic bases, keyspace quality

1. INTRODUCTION

For providing security for visual data, full encryption of the media bitstream with a traditional cipher, such as AES, although it presents the most secure option, often is not the best choice in terms of functionality. Many applications aim at trading off security for increased functionality. This could be scalability, a decrease in computational complexity or features like the support for transparent encryption. Lightweight encryption aims at striking a balance between security and such other requirements.

There are two groups of approaches that aim at providing lightweight encryption. First there is a group that operates on the bitstream level. For JPEG2000, introducing pseudo-random noise in the high-frequency subbands by pseudo-randomly inverting the signs of the coefficients in the code-blocks of the high resolution levels has been proposed by Groisbois and coworkers.¹ Norcen and Uhl² propose random permutation of coefficients and code-blocks. Wee and Apostolopoulos³ investigate bitstream encryption in the context of motion JPEG2000 coding, integrated into their scalable streaming concept. Kiya and coworkers⁴ and Wu and Deng⁵ discuss the problem of marker emulation in the context of JPEG2000 scrambling and encryption and propose solutions that allow to retain standard bitstream compliance.

The second group of approaches uses a secret transform domain to provide lightweight security. These techniques apply encryption integrated with compression, or rather: part of the compression, the transform step, is also the encryption step. From an overcomplete library of bases a single specimen is selected to be used for compression. Without knowing the used basis, no reconstruction or only a construction of limited quality should be possible. Compression-integrated encryption has been proposed outside the domain of JPEG2000: the encryption of the filter choice used for wavelet decomposition is proposed by Vorwerk et al.⁶ However, this suggestion remains vague and is not supported by any experiments. Fridrich et al.⁷ introduce the concept of key-dependent basis functions to protect a watermark from hostile attacks. This approach suffers from significant

Further author information:

Dominik Engel: E-mail: dengel@cosy.sbg.ac.at

Andreas Uhl: E-mail: uhl@cosy.sbg.ac.at

computational complexity. Fridrich⁸ develops the idea further and proposes a faster method for the generation of key-dependent orthogonal patterns. There are also some propositions that use secret Fourier transforms: The embedding of watermarks in an unknown domain is discussed by Djurovic et al.,⁹ and Unnikrishnan and Singh¹⁰ suggest to use this technique for encryption of visual data. Pommer and Uhl¹¹ propose encrypting the filterbanks used for an NSMRA decomposition.

In the domain of JPEG2000, the two methods investigated for this purpose include key-dependent basis functions and secret wavelet packet decompositions. Parameterized wavelet filters have been employed for JPEG2000 lightweight security by Köckerbauer et al.¹² and Engel and Uhl.¹³ Here, we focus on the use of randomized wavelet packet bases. This approach has been discussed in the context of a significance-map-based compression framework by Pommer and Uhl,^{14,15} as well as in the context of JPEG2000 by Engel and Uhl.^{16,17}

The used codec influences the level of security of lightweight encryption schemes that rely on a secret transform domain. The conclusion for JPEG2000 is that the wavelet packet approach is not suitable for providing full confidentiality.^{16,17} However, wavelet packets work very well for “sufficient” and “transparent” encryption. For sufficient encryption, the scheme tolerates image reconstructions without key that yield a discernible version of the original visual data. The only assertion of the scheme is that the quality of the reconstructed versions does not exceed a certain threshold. For transparent encryption,¹⁸ the scheme does not only tolerate image reconstructions of reduced quality, but uses them as a preview image and guarantees a certain quality. Such preview images can be of advantage in “try-and-buy”-scenarios, where they serve as an incentive to acquire the correct key for obtaining the full quality version of the visual data.

Anisotropic wavelet packets (AWP) have been proposed for the compression of image^{19,20} and video²¹ data. In order to enhance the keyspace size of the lightweight encryption approach with wavelet packets we have suggested the use of AWP.¹⁷ We have observed that AWP enhance security while computational complexity is reduced and compression performance stays the same. These are promising results, but so far no details have been given on the exact size of the keyspace, and the method for generating randomized anisotropic wavelet packets has not been validated against a uniform distribution. In this work we analytically and empirically investigate this approach in detail. In order to assess the security it provides, we quantify the size of the keyspace in detail. Furthermore, we investigate how the way that bases are randomly chosen from the keyspace affects compression performance and security. We compare the uniform distribution to the compression-oriented distribution to answer two questions: (a) is compression quality of a uniform distribution too low for application and (b) is the decrease in keyspace size that is introduced by the compression-oriented distribution large enough to be a problem for security.

The paper is organized as follows: In Section 2 we introduce the basic concept of lightweight encryption with anisotropic wavelet packets. We show how a uniform distribution can be achieved in the selection process, and review the selection process that follows a compression-oriented distribution. In Section 3 we compare the compression quality of the two approaches to give an answer to question (a). Section 4 compares the size of the keyspace in both distributions and aims at answering question (b). In Section 5 we turn to assessing the quality of the keyspace: We investigate how much similarity between the basis used for analysis and the basis used for synthesis is tolerable from a security point of view and how this affects the lightweight encryption scheme.

2. RANDOMIZED ANISOTROPIC WAVELET PACKETS

The wavelet packet transformation presents an overcomplete library of bases suitable for energy compaction in the frequency domain for visual data. Other than in the case of the pyramidal wavelet transform, in which the decomposition step is recursively applied only to the approximation subband, in the wavelet packet transform also the detail subbands are subject to further decomposition. The AWP transform is a generalization of the isotropic case: whereas in the latter, horizontal and vertical wavelet decomposition are always applied in pairs for each subband to be decomposed, this restriction is lifted for AWP.

For lightweight encryption with AWP a random specimen of the set of admissible bases is selected for transformation and kept secret. The description of the used basis can be used as a separate secret key or be encrypted with a traditional cipher and inserted into the bitstream. Only a minimal amount of data needs to be encrypted. Not all of the possible bases exhibit the same compression performance, therefore a trade-off has to be made between keyspace size and compression performance.

2.1. Total Number of Anisotropic Bases

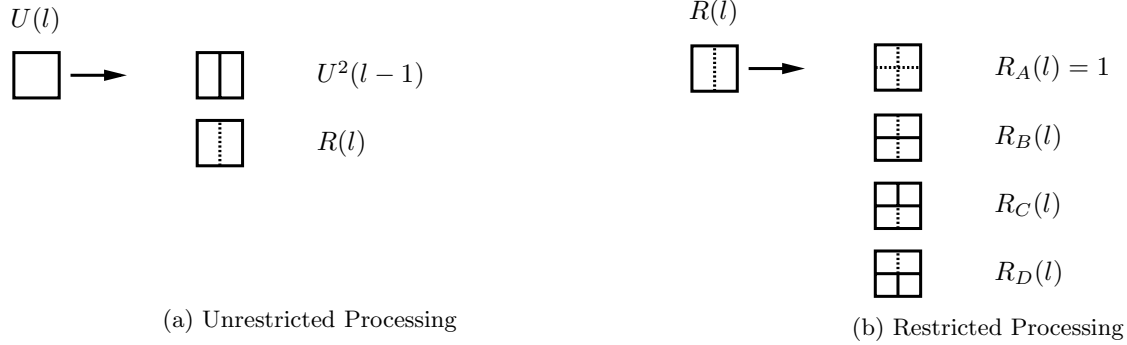


Figure 1. Case selection for uniform distribution of randomized AWP bases for joint maximum horizontal and vertical decomposition depth

For determining the number of AWP bases, the fact has to be taken into account that a horizontal decomposition followed by two vertical decompositions, or a vertical decomposition followed by two horizontal decompositions, lead to an equivalent result (“isotropic decomposition”), which must be counted only once. Xu and Do²⁰ suggest a way to determine the number of possible AWP bases for separate maximum decomposition depth in horizontal and vertical direction. For randomized anisotropic wavelet packets, a joint maximum decomposition depth for both directions is preferable. The method by Xu and Do can be adapted for this case. We determine $A(l)$, the number of bases of joint horizontal and vertical decomposition level up to l , recursively. The root node may not be decomposed, or it may be decomposed either horizontally or vertically, forming two subtrees of $A(j-1)$ possible decompositions in each case, leading to $1 + 2A^2(l-1)$ possible bases.

There are $2A^4(l-2)$ isotropic decompositions that are counted twice in this formula. Therefore, half their number, $A^4(l-2)$, is subtracted, leading to the formula:

$$A(l) = 1 + 2 \cdot A^2(l-1) - A^4(l-2) \quad (1)$$

where $A(0) = 1$, $A(1) = 3$, and $A(l) = 0$ for $l < 0$.

In preparation for defining a uniform distribution later, we use a different approach to determine the number of bases: The decomposition decisions are split into mutually exclusive cases (Figure 1), then the number of anisotropic bases is determined recursively by the number of bases contained in the subtree for each case. To address the fact that there is more than one way to construct an isotropic decomposition, two modes of decomposition are distinguished: *unrestricted* and *restricted*. Without loss of generality, we define the admissible decompositions for a subband without restriction as horizontal decomposition with further unrestricted processing (U), or further processing with restriction for horizontal decomposition (R). Case (R) leads to the restricted case in which horizontal decomposition is forbidden for at least one of the resulting subbands, leading to four possible decompositions (R_A through R_D). Let $U(l)$ be the number of bases up to level l in the unrestricted case and $R(l)$ be the number of bases up to level l in the restricted case, with the number of possible subcases $R_A(l)$, $R_B(l)$, $R_C(l)$, and $R_D(l)$. The possible decompositions in unrestricted and restricted mode are illustrated by Figure 1. We recursively determine the number of bases for each possible decomposition in the following way:

$$U(l) = \begin{cases} R(0) & \text{for } l = 0 \\ U^2(l-1) + R(l) & \text{else} \end{cases} \quad (2)$$

$$R(l) = R_A(l) + R_B(l) + R_C(l) + R_D(l) \quad (3)$$

$$R_A(l) = 1 \quad (4)$$

$$R_B(l) = \begin{cases} 0 & \text{for } l = 0 \\ R^2(l-1) & \text{else} \end{cases} \quad (5)$$

$$R_C(l) = R_D(l) = \begin{cases} 0 & \text{for } l = 0 \vee l = 1 \\ U^2(l-2) \cdot R(l-1) & \text{else.} \end{cases} \quad (6)$$

This formula is more flexible than Equation (1) and it forms the foundation for the uniform distribution as well as for determining the number of bases in the compression-oriented approach later on.

2.2. Uniform Distribution

Using a uniform random distribution over the whole set of bases is one way to choose a basis. In this procedure, every decomposition structure up to a maximum decomposition level is equally likely to be chosen for the transformation step. From the perspective of security, a uniform distribution is a good choice, as it does not give a potential attacker any prior knowledge.

We use the case distinction which we introduced above to construct a uniform distribution for the selection of a random basis: the probability for any case to be chosen is the ratio of the number of bases contained in the case to the total number of bases ($U(l)$ for unrestricted processing, $R(l)$ for restricted processing).

2.3. Compression-oriented Distribution

We have presented the procedure for generating randomized AWP bases that are suited for compression in previous work.¹⁷ Here we give a short summary of this procedure that aims at discarding bases that are too specific in their frequency properties to be useful for the compression of natural images. The most important parameters are:

n	Minimum decomposition depth of the approximation subband
m	Maximum decomposition depth of the approximation subband
e	Minimum decomposition depth of the detail subbands
d	Maximum decomposition depth of the detail subbands
q	Squareness factor for approximation subband
r	Squareness factor for detail subbands
bv	Base value of decomposition probability
cf	Change factor of decomposition probability
s	Seed for pseudo-random number generator

Table 1. Parameters for generating randomized AWP Bases

The first four parameters, n, m, e, d , determine the maximum and minimum decomposition depths for the approximation and the detail subbands. They influence both, compression performance and key-space size. Note that the number of decompositions is given here as single decompositions in any direction, whereas for the isotropic wavelet packet transform the number of decompositions usually denotes pairs of horizontal and vertical decompositions. Therefore, a decomposition depth of $2k$ in the anisotropic case is comparable to a decomposition depth of k in the isotropic case.

The squareness factors q and r are necessary to prevent subbands from being decomposed excessively in a single direction, as, especially in the case of the approximation subband, this would lead to inferior energy compaction in the frequency domain for the other direction. For the squareness of a subband we use the following definition:

$$Q(h, v) = \frac{1}{2^{|h-v|}} \quad (7)$$

where h and v are the decomposition depths in horizontal and vertical direction, respectively.

If a decomposition in the randomly chosen direction would result in this ratio dropping below the squareness factor, the direction is changed. A squareness factor of 0 means that no restriction in squareness is applied. The squareness factors influence both compression performance and key-space size. This is not the case for the following three parameters, which only determine the probability distribution of the randomly generated bases.

The seed s initializes the pseudo-random number generator. The base value bv determines the basic probability with which a subband is decomposed. The change factor cf alters this probability based on the decomposition depth of the subband: for each level of decomposition, the change factor is added to the base value. If cf is positive, then the higher the level of the subband, the higher is the chance for it to be decomposed. If cf is negative, the chance for decomposition decreases with higher decomposition levels. In this way, the generation process can be tuned to favor deeper or more shallow decompositions. Generally, it is advisable to tune these parameters to produce a balanced distribution.

3. COMPRESSION PERFORMANCE

Settings for the parameters pertaining to compression performance, which are successful at eliminating bases that are unsuited for compression have been determined empirically for test images of the size 512×512 pixels:¹⁷ $n = 6, m = 12, e = 0, d = 8, r = 0$, and $q = 0.5$. These settings restrain the keyspace compared to a random selection from the whole set of bases. The important question is if this restriction in keyspace size compromises security in a degree that is unacceptable for real life application. Section 4 deals with this issue.

In order to be useful in applications the secret transform domain must be successful in compacting the energy of natural images. For the uniform distribution, we tested a number of test images to get an idea of the compression performance. We compare these compression results with the compression-oriented distribution with the parameter setting suggested above. Table 2 lists the minimum, average, and maximum PSNR for a number of 8 bpp grayscale images of 512×512 pixels, each tested with 10.000 random bases at a compression rate of 0.25 bpp. For the uniform distribution a maximum decomposition depth of 10 was used.

Image	Pyramidal	Comp.-oriented			Uniform		
		Avg.	Min.	Max.	Avg.	Min.	Max.
Lena	32.26	31.97	30.61	32.42	28.75	25.65	29.62
Barbara	28.35	27.94	26.83	29.14	25.37	23.31	26.12
Peppers	33.45	32.93	31.17	33.53	29.41	26.75	30.36
Houses	23.47	23.15	22.40	24.03	21.37	19.68	22.06
Graves	28.30	27.94	26.22	28.61	24.68	23.12	25.31

Table 2. Compression performance of randomized AWP at 0.25bpp

The bases that give the maximum compression quality could of course also be selected by the uniform distribution. However, it can be seen, that in the total set of bases there are many that yield inferior compression results. In terms of compression performance it is therefore important to limit the number of admissible bases. Only the compression-oriented approach ensures good compression results. In the following we will evaluate if it also yields acceptable security.

4. KEYSPACE SIZE

The keyspace with a uniform distribution comprises all possible bases and its magnitude can easily be determined by using Equation (2). For the compression-oriented approach, more work is required to substantiate what the impact of the different parameters is on keyspace size. We use the definition of squareness given above. Furthermore, to reflect the minimum decomposition depth n we define N as

$$N(l, a) = \begin{cases} 1 & \text{for } a = 0 \\ 0 & \text{for } a = 1 \wedge l < n \\ 1 & \text{for } a = 1 \wedge l \geq n \end{cases} \quad (8)$$

where a defines if the current subband is in the approximation tree ($a = 1$) or in the detail tree ($a = 0$). We adapt Equations (2) – (6) to take the compression parameters into account as follows:

$$U(l, h, v, a) = \begin{cases} 1 & \text{for } h + v + 1 > d \\ R(0, h, v, a) & \text{for } l = 0 \\ R(l, h, v, a) & \text{for } a = 1 \wedge \\ & Q(h + 1, v) < q \\ U(l - 1, h + 1, v, a) \cdot U(l - 1, h + 1, v, 0) & \\ + R(l, h, v, a) & \text{else} \end{cases} \quad (9)$$

$$R(l, h, v, a) = R_A(l) + R_B(l) + R_C(l) + R_D(l) \quad (10)$$

$$R_A(l, h, v, a) = N(l, a) \quad (11)$$

$$R_B(l, h, v, a) = \begin{cases} 0 & \text{for } l = 0 \\ 0 & \text{for } h + v + 1 > d \\ 0 & \text{for } a = 1 \wedge \\ & Q(h, v + 1) < q \\ R(l - 1, h, v + 1, a) \cdot R(l - 1, h, v + 1, 0) & \text{else} \end{cases} \quad (12)$$

$$R_C(l, h, v, a) = \begin{cases} 0 & \text{for } l = 0 \vee l = 1 \\ 0 & \text{for } h + v + 1 > d \\ 0 & \text{for } a = 1 \wedge \\ & Q(h, v + 1) < q \\ U(l - 2, h + 1, v + 1, a) \cdot U(l - 2, h + 1, v + 1, 0) \cdot \\ R(l - 1, h, v + 1, 0) & \text{else} \end{cases} \quad (13)$$

$$R_D(l, h, v, a) = \begin{cases} 0 & \text{for } l = 0 \vee l = 1 \\ 0 & \text{for } h + v + 2 > d \\ 0 & \text{for } a = 1 \wedge \\ & Q(h, v + 1) < q \\ R(l - 1, h, v + 1, a) \cdot U^2(l - 2, h + 1, v + 1, a) & \text{else.} \end{cases} \quad (14)$$

This formula reflects the minimum decomposition depth n of the approximation subband by checking against N . The squareness of the approximation subband is handled by comparing Q to the minimum squareness factor q . Maximum squareness of the detail subbands is not set for the suggested parameters, but it could be easily included in a similar way. The maximum decomposition of the detail subbands is handled by checking against e .

We can now give the exact impact on keyspace size of the compression-oriented approach and compare it to the uniform distribution which uses all possible bases. Table 3 lists the numbers for some parameter settings. The last line represents the suggested settings for the compression-oriented approach, which only include bases that perform well for image compression.

It can be seen that the minimum decomposition depth of the approximation subband has little influence on keyspace size. Specifying a minimum squareness factor for the approximation subband does have a considerable

m	Compression-Oriented				Uniform
	n	d	q	#Bases	#Bases
6	0	6	0.5	2^{75}	2^{78}
12	6	12	0.5	2^{5048}	2^{5055}
12	0	8	0.5	2^{364}	2^{5055}
12	6	8	0	2^{371}	2^{5055}
12	6	8	0.5	2^{364}	2^{5055}

Table 3. Number of Bases for Compression-Oriented vs. Uniform Distribution

impact on keyspace size. The biggest cut in the keyspace is due to setting the maximum decomposition depth d for the detail subbands. The bases with a complex decomposition of the higher frequency subbands are not suitable for the compression of natural images. As the subbands of high decomposition depth in the high frequency subbands are very numerous, a majority of bases is discarded by setting d to a low value.

From a practical point of view the keyspace size provided by the compression-oriented approach with the suggested parameter settings is sufficient. A brute-force attack that tries to find the correct basis has a complexity of 2^{363} . Compared to complexity of a brute-force attack on 256-bit AES, which is 2^{255} , this keyspace size is more than sufficient for applications.

5. KEYSPACE QUALITY

It should be noted that a full brute-force attack may not be necessary. If a basis can be found that is close enough to the basis that was used for transformation, there will only be very little distortion in the reconstructed image. Theoretically, such a basis can be searched for incrementally, by trying to decode as much of the data as possible in resolution-progressive mode and then vary the higher frequency subbands and rerun the decoding. It has, however, been pointed out that the coding of the transform coefficients in JPEG2000 strongly relies on the associated subband structure.¹⁶ Especially for the high frequency subbands, where there is a large number of possibilities it will be hard, if not impossible, to perform this incremental search.

Without the correct subband structure, the coefficients of the JPEG2000 bitstream cannot be interpreted correctly. Nevertheless, the question should be addressed, how dissimilar two decomposition structures need to be to ensure sufficient distortion for the wrong key. We perform a replacement attack: Two randomly generated subband structures are produced, one of which is used for encoding. Then the coefficients of the encoding process are decoded with the second subband structure. The distance of the trees is recorded along with a quality comparison of two reconstructed images: one reconstructed with the decomposition structure used for encoding and the other reconstructed with the second subband structure.


We measure the distance of two anisotropic subband structures s_1 and s_2 by comparing leaves common to both decomposition trees. First, all leaves of s_1 and s_2 are assigned a score. We define the score $L(b)$ as

$$L(b) = 2^{m-r(b)} \quad (15)$$


where b is a subband, $r(b)$ is the resolution level of b and m is the maximum resolution level in any of the two trees. The rationale for the weighting by resolution level is that differences that occur on the lower resolution levels have a greater impact on visual quality and therefore should be given a higher rating. The total score $L_t(s_1, s_2)$ is computed as the sum of the scores of all individual subbands contained in s_1 and s_2 . The score L_u for unique subbands in s_1 and s_2 is defined as the sum of the scores of all subbands that are not contained in s_1 and s_2 . The distance between s_1 and s_2 is defined as L_u/L_t . This measure is relatively crude but should suffice to give an idea of the distances the two distributions achieve on average.

Figure 2 shows the empirical results for the uniform and the compression-oriented distribution for the same five test images used above and 10000 tests (i.e. 20000 randomly selected bases) for each image. Figures 2(a) and 2(b) compare the tree distances of the two distributions by the number of samples. It can be seen that for both distributions, two randomly selected basis generally have a large distance. For the uniform distributions the minimum distance is approximately 0.5, the trees in the compression-oriented distribution show higher similarities. As the uniform distribution generally produces bases with a larger number of leaves than the compression-oriented distribution, the chance is higher for some subbands in the two bases to be the same. This is the reason why there is a smaller number of pairs of trees with a distance of 1. Figures 2(c) and 2(d) compare the average PSNR by tree distance. The distance is given on the abscissa, the ordinate gives the ratio of the PSNR of the image reconstructed with the wrong subband structure and the PSNR of the correctly reconstructed image. It can be seen that, as expected, the uniform distribution avoids smaller distances and therefore achieves greater reduction in visual quality.

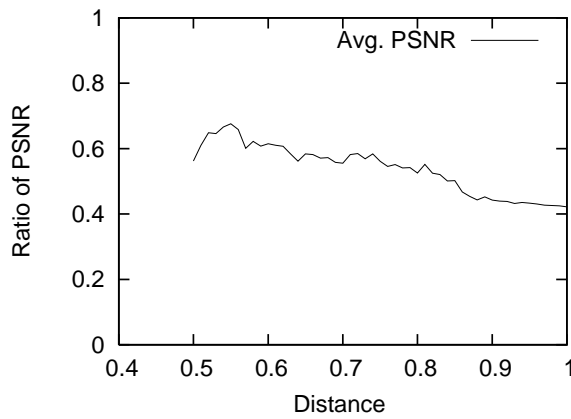
Figures 2(e) and 2(f) compare the number of bases by achieved reduction in visual quality. In this case the abscissa gives the ratio of the PSNR of the wrong subband structure and the correct subband structure. The ordinate reports the ratio of samples that fall into the corresponding PSNR slot and the total number of samples.




(a) Samples by Distance – Uniform Distribution



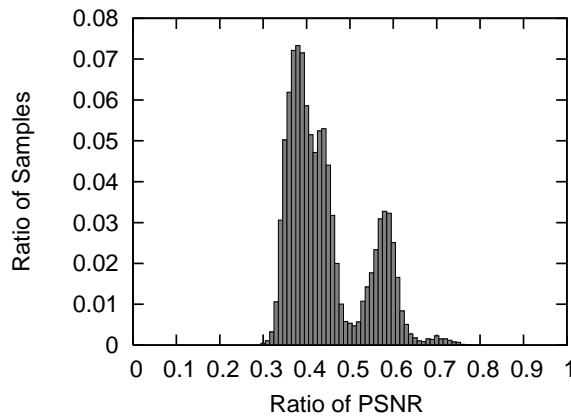
(b) Samples by Distance – Compression-oriented Distribution




(c) PSNR by Distance – Uniform Distribution



(d) PSNR by Distance – Compression-oriented Distribution



(e) Samples by PSNR – Uniform Distribution



(f) Samples by PSNR – Compression-oriented Distribution

Figure 2. Visual Quality and Tree Distance

It can be seen that both distributions on average produce a sufficient reduction in visual quality. The problem of the compression-oriented approach is that there are some cases in which the visual distortion is less than 20%. The chances for such a case to occur are very low, as can be seen in Figure 2(b): the number of pairs of basis that are not sufficiently distinct, i.e. the tree distance is lower than approximately 0.4, is very small. An attacker would have to perform a brute-force search on the most part of the keyspace to find a basis that produces less than 20% distortion. Furthermore, as stated earlier, the coefficients of the higher frequency subbands cannot be sensibly decoded from the bitstream without the subband structure, even for largely similar structures, so the replacement attack is hardly feasible.

The fact remains that even without the correct key, the lower resolutions can be reconstructed easily. This is the reason why the suggested approach is not suitable for providing full confidentiality but can only be used to provide sufficient or transparent encryption. For sufficient encryption, the availability of a preview image is tolerated, for transparent encryption it is even desired. Full protection is only needed for the high resolutions. The compression-oriented approach yields a keyspace that can provide this protection.

6. CONCLUSION

The compression-oriented approach for lightweight encryption with AWP discards a number of possible bases to produce good compression results with randomly selected AWP. By comparing compression performance to a method that randomly selects from all possible bases, we have shown that pruning the set of all bases is necessary, as compression results vary immensely. If all bases were used, the approach would not be suitable for application. In order to assess the loss in keyspace size we formally determined the total number of bases and the number of bases for the compression-oriented approach. A comparison shows that the number of bases that are suitable for image compression is small compared to the total number of possible bases but still above the complexity of a brute-force attack against 256-bit AES. A quality assessment of the set of available bases showed that although for very similar bases the reduction in visual quality is low, the effort for finding a basis for which the distortion of the reconstructed image is less than 20 percent is close to a full search.

As a side product of our evaluation, we have given an approach to generate uniformly distributed AWP decompositions. This could be useful for the investigation of their compression performance for different classes of images. Furthermore, we have discussed in detail the number of possible AWP decompositions.

The evaluation of the AWP lightweight encryption approach regarding compression performance and keyspace size shows that (a) a cut in the available bases is necessary to facilitate good compression results and (b) the loss in keyspace size is substantial but keyspace size remains large enough for application.

ACKNOWLEDGMENTS

We are very grateful to Rade Kutil for his help with the uniform distribution.

This work has been partially supported by the European Commission through the IST Programme under contract IST-2002-507932 ECRYPT and by the Austrian Science Fund (FWF) under project number P15170.

REFERENCES

1. R. Grosbois, P. Gerbelot, and T. Ebrahimi, "Authentication and access control in the JPEG 2000 compressed domain," in *Applications of Digital Image Processing XXIV*, A. Tescher, ed., *Proceedings of SPIE* **4472**, pp. 95–104, (San Diego, CA, USA), July 2001.
2. R. Norcen and A. Uhl, "Encryption of wavelet-coded imagery using random permutations," in *Proceedings of the IEEE International Conference on Image Processing (ICIP'04)*, IEEE Signal Processing Society, (Singapore), Oct. 2004.
3. S. Wee and J. Apostolopoulos, "Secure scalable streaming and secure transcoding with JPEG2000," in *Proceedings of the IEEE International Conference on Image Processing (ICIP'03)*, **I**, pp. 547–551, (Barcelona, Spain), Sept. 2003.

4. H. Kiya, D. Imaizumi, and O. Watanabe, "Partial-scrambling of image encoded using JPEG2000 without generating marker codes," in *Proceedings of the IEEE International Conference on Image Processing (ICIP'03)*, **III**, pp. 205–208, (Barcelona, Spain), Sept. 2003.
5. Y. Wu and R. H. Deng, "Compliant encryption of JPEG2000 codestreams," in *Proceedings of the IEEE International Conference on Image Processing (ICIP'04)*, IEEE Signal Processing Society, (Singapore), Oct. 2004.
6. L. Vorwerk, T. Engel, and C. Meinel, "A proposal for a combination of compression and encryption," in *Visual Communications and Image Processing 2000, Proceedings of SPIE 4067*, pp. 694–702, (Perth, Australia), June 2000.
7. J. Fridrich, A. C. Baldoza, and R. J. Simard, "Robust digital watermarking based on key-dependent basis functions," in *Information hiding: second international workshop*, D. Aucsmith, ed., *Lecture notes in computer science 1525*, pp. 143–157, Springer Verlag, Berlin, Germany, (Portland, OR, USA), Apr. 1998.
8. J. Fridrich, "Key-dependent random image transforms and their applications in image watermarking," in *Proceedings of the 1999 International Conference on Imaging Science, Systems, and Technology, CISST '99*, pp. 237–243, (Las Vegas, NV, USA), June 1999.
9. I. Djurovic, S. Stankovic, and I. Pitas, "Digital watermarking in the fractional fourier transformation domain," *Journal of Network and Computer Applications 24*, pp. 167–173, 2001.
10. G. Unnikrishnan and K. Singh, "Double random fractional fourier-domain encoding for optical security," *Optical Engineering 39*, pp. 2853–2859, Nov. 2000.
11. A. Pommer and A. Uhl, "Wavelet packet methods for multimedia compression and encryption," in *Proceedings of the 2001 IEEE Pacific Rim Conference on Communications, Computers and Signal Processing*, pp. 1–4, IEEE Signal Processing Society, (Victoria, Canada), Aug. 2001.
12. T. Köckerbauer, M. Kumar, and A. Uhl, "Lightweight JPEG 2000 confidentiality for mobile environments," in *Proceedings of the IEEE International Conference on Multimedia and Expo, ICME '04*, (Taipei, Taiwan), June 2004.
13. D. Engel and A. Uhl, "Parameterized biorthogonal wavelet lifting for lightweight JPEG 2000 transparent encryption," in *Proceedings of ACM Multimedia and Security Workshop, MM-SEC '05*, pp. 63–70, (New York, NY, USA), Aug. 2005.
14. A. Pommer and A. Uhl, "Selective encryption of wavelet packet subband structures for secure transmission of visual data," in *Multimedia and Security Workshop, ACM Multimedia*, J. Dittmann, J. Fridrich, and P. Wohlmacher, eds., pp. 67–70, (Juan-les-Pins, France), Dec. 2002.
15. 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)*, pp. 279–287, 2003.
16. D. Engel and A. Uhl, "Secret wavelet packet decompositions for JPEG 2000 lightweight encryption," in *Proceedings of 31st International Conference on Acoustics, Speech, and Signal Processing, ICASSP '06*, **V**, pp. 465–468, (Toulouse, France), May 2006.
17. D. Engel and A. Uhl, "Lightweight JPEG2000 encryption with anisotropic wavelet packets," in *Proceedings of International Conference on Multimedia & Expo, ICME '06*, pp. 2177–2180, (Toronto, Canada), July 2006.
18. B. M. Macq and J.-J. Quisquater, "Cryptology for digital TV broadcasting," *Proceedings of the IEEE 83*, pp. 944–957, June 1995.
19. R. Kutil, "Zerotree image compression using anisotropic wavelet packet transform," in *Visual Communications and Image Processing 2003 (VCIP'03)*, T. Ebrahimi and T. Sikora, eds., *SPIE Proceedings 5150*, pp. 1417–1427, SPIE, (Lugano, Switzerland), July 2003.
20. D. Xu and M. N. Do, "Anisotropic 2-D wavelet packets and rectangular tiling: theory and algorithms," in *Proceedings of SPIE Conference on Wavelet Applications in Signal and Image Processing X*, M. A. Unser, A. Aldroubi, and A. F. Laine, eds., *SPIE Proceedings 5207*, pp. 619–630, SPIE, (San Diego, CA, USA), Aug. 2003.
21. R. Kutil, "Anisotropic 3-D wavelet packet bases for video coding," in *Proceedings of the IEEE International Conference on Image Processing (ICIP'03)*, (Barcelona, Spain), Sept. 2003.