LIGHTWEIGHT JPEG2000 ENCRYPTION WITH ANISOTROPIC WAVELET PACKETS

Dominik Engel and Andreas Uhl

Department of Computer Sciences
University of Salzburg, Austria
Email: {dengel, uhl}@cosy.sbg.ac.at

ABSTRACT

A lightweight encryption technique for JPEG2000 with optional support for transparent encryption is proposed. Randomized anisotropic wavelet packet bases are used to construct a secret frequency domain, leading to a situation in which only a minimal amount of data needs to be encrypted. Results and calculations are presented to evaluate the suggested approach in terms of compression performance, security, and applicability.

1. INTRODUCTION

With the spreading use of multimedia applications, security techniques specifically tailored for multimedia data have become an important research area. This is especially true for scalable formats that become crucial as a multitude of viewing devices of varying resources in terms of display and computational processing power have to be supported by a single bitstream.

While for multimedia data like any other type of data, full encryption with a traditional cipher remains the most secure option, 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. Lightweight encryption aims at striking a balance between security and these other requirements. Traditional approaches assume that confidentiality is required in the sense that without key-data no discernible version of the visual data can be obtained. In some contexts, this strict requirement is not necessary; it is sufficient that the full quality of the visual data is protected, but the possibility to decode versions of (substantially) degraded quality or low resolution from the encrypted bitstream without key-data does not pose a problem. In the context of wavelet coded image data, various methods have been proposed for lightweight encryption: Selective encryption of crucial parts of the JPEG2000 packet data is proposed by [1]. [2] propose introducing pseudo-random noise in the high-frequency subbands by inverting coefficient signs in the high resolutions. [3] investigate bitstream encryption in the context of motion JPEG2000 coding, integrated into their scalable streaming concept. [4] and [5] discuss the problem of marker emulation in the context of JPEG2000 scrambling and encryption and propose solutions that allow to retain standard bitstream compliance.

For some applications the possibility of decoding a version of degraded quality is not only tolerable but it is desired to provide a preview image as an incentive for potential customers to purchase the full quality version. In this case a minimum quality has to be guaranteed for images decoded without key or with the wrong key. Encryption schemes of this kind have been termed "transparent encryption" schemes [6]. [7] propose to encrypt about 85% of the packet data in resolution progressive mode for JPEG2000 transparent encryption.

In the approach presented here, we aim at providing lightweight encryption with optional support for transparent encryption by the use of the anisotropic wavelet packet transform. The suggested approach is compliant to JPEG2000 part II. Security for lower resolutions is weaker, full security is provided for the higher resolutions. A main advantage of this scheme is that only an extremely small amount of data needs to be encrypted: the parameters of the randomized wavelet packet bases generation and the seed of the pseudo-random number generation. Because the wavelet packet transform is of higher complexity than the pyramidal wavelet transform, computational demands are transferred from the encryption to the compression stage. In Sec. 2 we outline the algorithm for randomized generation of anisotropic wavelet packet bases. We then discuss its parameter settings in the light of compression results in Sec. 3. Sec. 4 addresses security issues. Sec. 5 discusses the applicability of the proposed approach and concludes.

2. ANISOTROPIC WAVELET PACKETS (AWP)

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 idea of using randomized wavelet packet bases for lightweight encryption by creating a secret frequency domain is proposed by [8] in the context of a significance-map-based compression algorithm. We transfer this approach to the domain of JPEG2000 and show that it can be improved significantly by the use of anisotropic wavelet packets.

The anisotropic wavelet packet 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 anisotropic wavelet packets. Anisotropic wavelet packets have been proposed for the compression of image [9, 10] and video [11] data. The main motivation to introduce anisotropic wavelet packets for lightweight encryption is a substantial increase in keyspace size: the space of possible bases is not only spanned by the decision of decomposing or not (as is the case for the isotropic transform), but also by the direction of each decomposition. During compression 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 encrypted with a traditional cipher and inserted into the bitstream. Only a minimal amount of data needs to be encrypted.

As not all anisotropic wavelet packet decompositions produce

This work has partially been funded by the Austrian Science Fund (FWF), project no. 15170. Dominik Engel gratefully acknowledges funding by the Austrian Academy of Sciences (ÖAW).

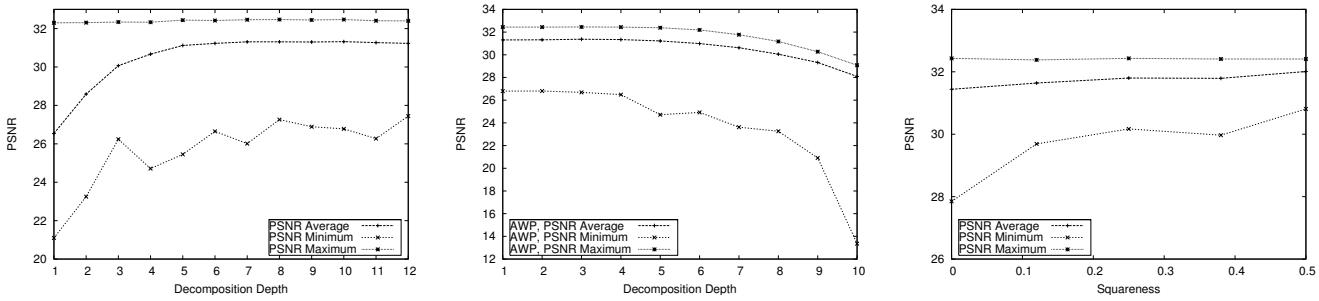

(a) Minimum decomposition depth for approximation subband

(b) Minimum decomposition depth for detail subbands

(c) Squareness factor for approximation subband

Fig. 1. Parameter settings and compression performance

good compression results, we introduce parameters that can be used to constrain the possible anisotropic decompositions. This reduces the size of the available keyspace, so the goal here is to strike a good balance between satisfactory compression performance and keyspace size. In the next section we propose settings that achieve such a balance. The following parameters are of relevance:

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 keyspace size. Note that the number of decomposition is given here as single decompositions in any direction, whereas for the isotropic case 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 into a single direction excessively, as, especially in the case of the approximation subband, this would lead to inferior energy compaction in the frequency domain for the other direction. The squareness parameters reflect a threshold for the ratio of the longer side of the subband to the shorter side. 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 checking is done. The squareness factors influence both compression performance and keyspace 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. In this way, the generation process can be tuned to favor deeper or more shallow decompositions.

Transparent encryption can be accommodated in the proposed scheme by introducing a parameter p that reflects the number of resolutions that can be decoded without knowledge of the anisotropic decomposition structure. For this purpose, $2r$ decompositions, alternating between horizontal and vertical direction, are applied recursively to the LL-Subband, where r is the total number of resolutions. Of the $2r$ detail subbands generated in this way, only the first $2r-2p$ are subject to further decomposition. The resulting LL-subband, and the corresponding detail subbands for the resolutions R_0 to R_{p-1} are the same as that produced by the pyramidal wavelet transform. Any decoder compliant to JPEG2000 part I can be used to decode the first p resolutions.

3. COMPRESSION PERFORMANCE

In order to determine parameter settings that favor good compression results we use a set of grayscale test images and vary the parameters in their respective ranges and progressively eliminate settings that produced inferior compression results. The implementation is based on the JJ2000 reference implementation. The plots shown for illustration are for the image Lena, 512×512 pixels at a compression ratio of 0.25 bpp, and present the minimum, average, and maximum PSNR of the randomly generated decompositions. At the end of this section we present compression results for other images.

Fig. 1(a) shows the compression performance by minimum decomposition depth of the approximation subband. As can be seen, setting n is important, as only a sufficient number of decompositions ensures competitive compression results. For the image size used in our tests, setting the maximum decomposition depth m to 12 produced favorable compression results.

For most natural images, the situation for the detail subbands is different, as illustrated by Fig. 1(b). A minimum decomposition depth does not improve compression results in this case. If the decomposition depth of the detail subbands is set too high, however, significant overhead is introduced by the large number of subbands that leads to a deterioration of compression performance, so the maximum decomposition depth of the detail subbands, d, should be confined. On the other hand, d must not be set much lower than m, as this affects security for the lower resolutions. In our tests, $d = 8$ produced acceptable results. In order to increase security for the lower resolutions (e.g. when only one level of transparency is desired, i.e. p is set to 1), an alternative definition of d can be used, in which d defines the maximum decomposition depth in *addition* to the depth induced by the approximation subband. In this way, a number of different decompositions in the lower resolutions (which are further decomposed than d) become possible, increasing security at the cost of a slight loss in compression performance. An additional parameter needs to be introduced in this case to regulate the maximum global decomposition depth for any subband.

Fig. 1(c) shows the compression performance for $n = 6, m = 12, e = 0, d = 8$ and varying settings for the squareness factor of the approximation subband (q). It can be seen, that a high degree of

squareness, i.e. a high similarity to the isotropic decomposition, increases minimum and average compression performance. We therefore propose to set q to the maximum of 0.5. For the detail subbands, our results show that the situation is different: the squareness factor r cannot improve compression results – the curves for average, minimum and maximum are nearly parallel to the abscissa. While it is surprising that good compression results can be achieved with most combinations of horizontal and vertical decompositions in the detail subbands, the situation obviously yields an advantage in terms of keyspace size. Fig. 2 compares the parameter settings by plotting on the ordinate the percentage of samples for which the compression quality lies below the PSNR-value of the abscissa. It can be seen that setting the squareness factor for the approximation subband puts the finishing touch on compression performance.

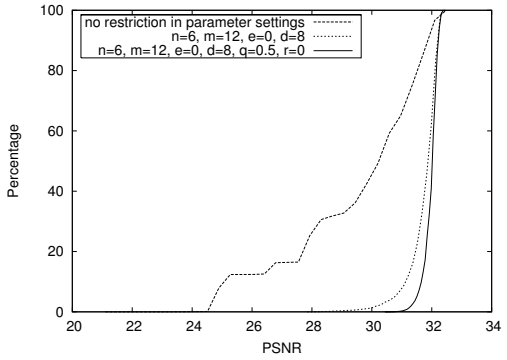

Fig. 2. Comparison of compression performance

Table 2 shows the compression performance for some of the tested images. It can be seen that on average, the compression quality is competitive with the pyramidal wavelet transform. Furthermore, the performance is comparable to randomized isotropic wavelet packets (IWP).

Image	Avg	Min	Max	Pyr (lev. 6)
Lena (AWP)	31.97	30.61	32.42	32.26
Barbara (AWP)	27.94	26.83	29.14	28.35
Peppers (AWP)	32.93	31.17	33.53	33.45
Houses (AWP)	23.15	22.40	24.03	23.47
Graves (AWP)	27.94	26.22	28.61	28.30
Lena (IWP)	31.74	29.91	32.45	32.26
Barbara (IWP)	28.21	26.65	29.19	28.35

Table 2. Compression performance of randomized AWP

4. SECURITY EVALUATION

Possible attacks on the proposed scheme are to (a) break the cipher with which the key was encrypted, (b) infer the wavelet packet structure from statistical properties of the wavelet coefficients, (c) infer the wavelet packet structure from the codestream or to (d) perform a full search. The feasibility of attack (a) is equivalent to the feasibility of breaking the used cipher.

For JPEG2000 attacks (b) and (c) are also not feasible. Apart from the encrypted parameters, no information on the anisotropic wavelet packet structure is contained in the header data. Inferring the decomposition structure from the codestream is impossible, because JPEG2000 employs so-called tag-trees [12] to signal inclusion information: In a highly contextualized coding scheme, the contributions of each code-block contained in a packet are linked to the subband structure. Thereby the subband structure is used as context to interpret the output of the tagtrees. If the subband decomposition structure is unknown, the attacker has no way of correctly interpreting this output: the attacker can only see the answer given to an inclusion question, but, lacking the decomposition structure, does not know the right question. Furthermore, the fact that the inclusion information cannot be decoded eliminates access to the raw coefficient data, as an attacker cannot correctly associate the contributions of a code-block to the correct coefficients. Attack (b), inferring the wavelet packet structure from the wavelet coefficients, therefore is not feasible for the scheme proposed here. In this respect the security of the secret frequency domain is strongly dependent on the used codec: for JPEG2000 attacks (b) and (c) do not pose a threat, whereas for other codecs this may well be the case. For the scheme proposed by [8], for example, option (b) presents a successful attack. One important exemption to what has been detailed above has to be made for the single subband of the lowest resolution (the LL-subband). If this subband is of quadratic shape, then its packets will be the same as when the pyramidal decomposition is used. Therefore, the proposed scheme cannot be used to provide strict confidentiality without encrypting the LL-subband.

The feasibility of attack (d) depends on the size of the keyspace, which is the number of anisotropic bases for the used parameters. Following [10] we determine G_j, the number of bases of decomposition level up to j, recursively. The root node may not be decomposed, or it may be decomposed either horizontally or vertically, forming two subtrees of G_{j-1} possible decompositions in each case, leading to $1+2\cdot(G_{j-1})^2$ possible bases. There exist, however, some decompositions that result in the same basis: a horizontal decomposition (**r**) followed by two vertical decompositions (**c**) on the resulting subtrees is equivalent to the case in which the vertical decomposition is done first followed by two horizontal decompositions, as illustrated in Fig. 3. As these bases should be counted only once, half their number, $(G_{j-2})^4$, is subtracted, leading to the formula:

$$G_j = 1 + 2 \cdot (G_{j-1})^2 - (G_{j-2})^4 \qquad (1)$$

where $G_0 = 1$ and $G_1 = 3$.

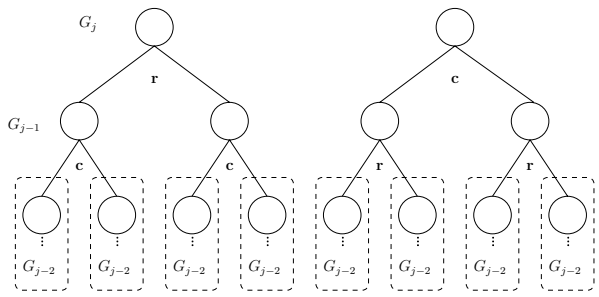

Fig. 3. Number of equivalent AWP bases

Anisotropic			Isotropic		
G_6	$\approx 10^{23}$	$\approx 2^{78}$	K_3	$\approx 10^{5}$	$\approx 2^{16}$
G_8	$\approx 10^{95}$	$\approx 2^{315}$	K_4	$\approx 10^{19}$	$\approx 2^{65}$
G_{10}	$\approx 10^{380}$	$\approx 2^{1263}$	K_5	$\approx 10^{78}$	$\approx 2^{261}$
G_{14}	$\approx 10^{6088}$	$\approx 2^{20225}$	K_7	$\approx 10^{1260}$	$\approx 2^{4185}$

Table 3. Number of anisotropic and isotropic wavelet packet bases

Table 3 compares G_j, the number of anisotropic wavelet packet bases with j horizontal or vertical decompositions, to K_i, the number of isotropic wavelet bases with i pairs of horizontal and vertical decomposition. It can be seen that the increase in keyspace size introduced by the use of anisotropic wavelet packets is substantial.

It has to be noted that the numbers given here do not reflect the reduction by the parameter settings that control compression performance. However, with the suggested maximum decomposition depth of 12 for the approximation subband and 8 for the detail subbands, the full search option is still in a higher order of complexity than a full search attack on AES with a 256-bit key (2^{255}). This can be shown as follows: To obtain a lower bound, we only regard the number of bases induced by the decomposition of the detail subbands, as the restrictions that have most impact on the keyspace size pertain to the approximation subband (see the proposed settings for minimum decomposition depth n and squareness factor q). In correspondence to the suggested parameter settings, we assume that no minimum decomposition depth or squareness factor is set for the detail subbands, and that the maximum decomposition depth for the approximation subband (m) is greater or or equal the maximum decomposition depth for the details subbands (d). We then assume the root to be decomposed either horizontally or vertically (reflected by the factor 2 in the formula below). The approximation subband generated by this initial decomposition is decomposed up to the depth of d in alternating directions, starting with the inverse direction of the root subband. The resulting approximation subband does not violate the squareness requirement q. The decomposition leads to d detail subbands. We assume that each of these detail subbands is either not decomposed or further decomposed an arbitrary amount of times up to level d, with the first decomposition being in the inverse direction of the neighboring approximation subband to avoid isotropic decompositions. For the detail subband at level i this corresponds to at least $1 + (G_i - 1)/2$ possibilities. The combination of the possibilities in the subtree of each of the d detail subbands gives a lower bound for P_d, the number of possible bases that can be obtained with $e = 0$, $r = 0$, $m \geq d$, and arbitrary settings for q and n:

$$2 * \prod_{i=0}^{d-1} (1 + \frac{G_i - 1}{2}) \leq P_d. \quad (2)$$

For $d = 8$, the keyspace size is greater than 2^{302}, and thus above the full search complexity of AES with a 256-bit key.

5. CONCLUSION

The proposed scheme successfully retains compression performance and increases keyspace size as compared to previously suggested approaches that only use isotropic wavelet bases. As the the anisotropic wavelet transform and the isotropic wavelet transform are of the same computational complexity, no processing overhead is introduced. In other words, compared to isotropic wavelet packets, the same keyspace size can be achieved with anisotropic wavelet packets at a significantly lower cost of computational complexity.

The proposed scheme is not suitable for applications that require full confidentiality in the sense that without the key, the encrypted bitstream may not be decoded in any way that yields an image in which the original data is discernible. Transparent encryption, on the other hand, in which a minimum quality for a preview image has to be guaranteed, can be easily incorporated into the proposed scheme. As the amount of data to be encrypted is minimal, the proposed scheme is well suited to be combined with public key cryptography and benefit from the superior key management of this approach.

In future work, we will investigate the advantages of combining the anisotropic wavelet packet transform with other techniques, such as parameterized filters [13], to increase security. Other areas of multimedia security can also benefit from the techniques explored in the present work. Specifically, we will investigate the use of wavelet packets to improve security for perceptual hashing and watermarking techniques.

6. REFERENCES

[1] R. Norcen and A. Uhl, "Selective encryption of the JPEG2000 bitstream," in *Communications and Multimedia Security. Proceedings of the IFIP TC6/TC11 Sixth Joint Working Conference on Communications and Multimedia Security, CMS '03* (A. Lioy and D. Mazzocchi, eds.), vol. 2828 of *Lecture Notes on Computer Science*, (Turin, Italy), pp. 194 – 204, Springer-Verlag, Oct. 2003.

[2] 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.), vol. 4472 of *Proceedings of SPIE*, (San Diego, CA, USA), pp. 95–104, July 2001.

[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)*, vol. I, (Barcelona, Spain), pp. 547–551, 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)*, vol. III, (Barcelona, Spain), pp. 205–208, 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)*, (Singapure), IEEE Signal Processing Society, Oct. 2004.

[6] B. M. Macq and J.-J. Quisquater, "Cryptology for digital TV broadcasting," *Proceedings of the IEEE*, vol. 83, pp. 944–957, June 1995.

[7] A. Uhl and C. Obermair, "Transparent encryption of JPEG2000 bitstreams," in *Proceedings EC-SIP-M 2005 (5th EURASIP Conference focused on Speech and Image Processing, Multimedia Communications and Services)* (P. Podhradsky *et al.*, eds.), (Smolenice, Slovak Republic), pp. 322–327, 2005.

[8] A. Pommer and A. Uhl, "Selective encryption of wavelet-packet encoded image data — efficiency and security," *ACM Multimedia Systems (Special issue on Multimedia Security)*, vol. 9, no. 3, pp. 279–287, 2003.

[9] 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.), vol. 5150 of *SPIE Proceedings*, (Lugano, Switzerland), pp. 1417–1427, SPIE, July 2003.

[10] 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*, (San Diego, USA), Aug. 2003.

[11] 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.

[12] D. Taubman, "High performance scalable image compression with EBCOT," *IEEE Transactions on Image Processing*, vol. 9, no. 7, pp. 1158 – 1170, 2000.

[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*, (New York, NY, USA), pp. 63–70, Aug. 2005.