Lightweight Protection of Visual Data Using High-Dimensional Wavelet Parametrization*

Andreas Pommer and Andreas Uhl

Salzburg University, Department of Scientific Computing,
Jakob Haringer-Str. 2, A-5020 Salzburg, Austria
{apommer, uhl}@cosy.sbg.ac.at
http://www.scicomp.sbg.ac.at/

Abstract. A lightweight encryption scheme for visual data based on wavelet filter parametrization is discussed. Being a special variant of header encryption, the technique has an extremely low computational demand. Security assessement of low-dimensional parametrizations schemes show severe weaknesses. We show that using high-dimensional parametrizations the scheme may be employed in applications requiring a medium security level.

1 Introduction

Encryption schemes for multimedia data need to be specifically designed to protect multimedia content and fulfil the security requirements for a particular multimedia application. For example, real-time encryption of an entire video stream using classical ciphers requires heavy computation due to the large amounts of data involved, but many multimedia applications require security on a much lower level (e.g. TV news broadcasting [5]). In this context, several selective encryption schemes have been proposed recently which do not strive for maximum security, but trade off security for computational complexity (see [12] for a comprehensive overview). For example, we mention selective encryption of MPEG streams [1] and of JPEG 2000 data [2, 7]. However, in case that one or more parties which are involved in an application have strong limits on their processing capacities (e.g., a mobile device with a small battery and a slow processor), even encrypting a small fraction of the image data may still be out of reach. In such environments, confidentiality may be provided by an extreme case of selective encryption as will be described in the following.

In recent work, we have proposed lightweight encryption schemes based on orthogonal [3] and biorthogonal [11] wavelet filter parametrizations for applications requiring a low to medium security level. In this work we investigate higher-dimensional wavelet parametrizations in order to increase the key-space and attack resistance of the scheme. In Section 2, we shortly review wavelet

* This work has been partially supported by the Austrian Science Fund FWF, project no. P15170.

F. Roli and S. Vitulano (Eds.): ICIAP 2005, LNCS 3617, pp. 645–652, 2005.

compression and the wavelet filter parametrization scheme in use. Section 3 introduces the encryption scheme, evaluates the resulting compression quality and security, and discusses the use of higher-dimensional parametrizations.

2 Wavelet Compression and Filter Parametrization

Wavelet-based still image compression has to be considered state of the art nowadays, especially in applications requiring low bit rates and bitstream scalability. In the area of standardization the two most prominent techniques are JPEG 2000 and MPEG-4 VTC. The SMAWZ codec [4] used in our experiments is a variant of the well known SPIHT algorithm which has been optimized for efficient implementation using bitplanes instead of lists. In all these compression schemes, filters especially tuned for that specific purpose are employed. However, there exists an almost infinite richness of different wavelet filters to choose from.

For the construction of compactly supported orthonormal wavelets, solutions for the dilation equation have to be derived, satisfying two conditions on the coefficients c_k ($\phi(t) = \sum_{k\in\mathbb{Z}} c_k \phi(2t-k)$, with $c_k \in \mathbb{R}$). In our work we use a family of parameterized filters generated according to an algorithm proposed by Scheid and Pittner [10]:

Given N parameter values $-\pi \leq \alpha_i < \pi$, $0 \leq i < N$, the following recursion

$$c_0^0 = \tfrac{1}{\sqrt{2}} \text{ and } c_1^0 = \tfrac{1}{\sqrt{2}}$$

$$c_k^n = \tfrac{1}{2}\Big((c_{k-2}^{n-1} + c_k^{n-1}) \cdot (1+\cos\alpha_{n-1}) + (c_{2(n+1)-k-1}^{n-1} - c_{2(n+1)-k-3}^{n-1})(-1)^k \sin\alpha_{n-1}\Big)$$

can be used to determine the filter coefficients c_k^N, $0 \leq k < 2N+2$. We set $c_k = 0$ for $k < 0$ and $k \geq 2N+2$. Example filters which can be generated using this formula are the Daubechies-6 filter, which can be constructed using the parameters (0.6830127, −0.1830127), or the Haar filter which is generated with the parameter 0.

Note that the number N of parameter values α_i is denoted as the dimensionality of the parametrization scheme. Larger N lead to longer wavelet filters.

3 A Lightweight Encryption Scheme

As we have seen, wavelet-based compression can be performed using a wide variety of different wavelet transforms. This degree of freedom may be exploited to add security to wavelet-based applications by only encrypting the header information defining the wavelet transform in use and keeping the rest in plaintext. Following this general idea, selective encryption schemes based on encrypting the secret wavelet packet subband structure [9] or NSMRA decomposition scheme [8] have been proposed recently.

Besides their use in encryption, secret wavelet filters generated by filter parametrizations like those reviewed in section 2 have been proven to increase security in wavelet-based watermarking schemes [6]. In this paper we investigate the properties of a header encryption variant where we keep the parameter to generate the filters for the wavelet transform secret. For example, this can be easily achieved in the context of JPEG 2000 Part II by simply encrypting the corresponding field containing the user-defined custom filters in the header using a cryptographically strong cipher. As a consequence, the amount of data subject to encryption is minimal, since no actual image data but only filter coefficients are encrypted.

In the following subsections, we investigate the compression quality and the security of the resulting scheme.

3.1 Compression Quality

Whereas the traditional filters used for wavelet compression are tuned for optimal concentration of the energy contained in the image and the separation of high- and low-frequency parts, parameterized filters provide a wide quality range. The advantage as well as the disadvantage of parameterized filters is their variety, not all filters within such a family are equally suited for a specific purpose like image compression. Fig. 1(a) shows the resulting compression ratios when compressing the 8 bpp 512×512 pixels Lena image using different parameter values to a set of fixed quality levels in the range between 38 and 46 dB PSNR.

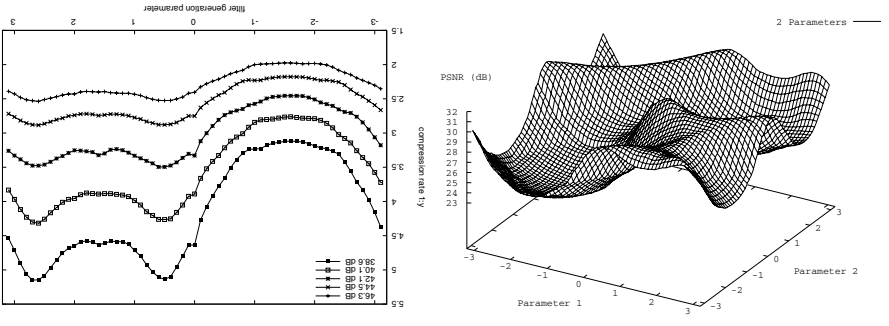

(a) Compression rate for various quality levels, 1-d parameterization

(b) Compression quality at fixed bitrate, 2-d parameter space

Fig. 1. Compression results employing parametrized filters within SMAWZ.

It is clearly displayed that the file sizes obtained by the filters resulting from the parametrization algorithm described in Section 2 vary to a large extent. Obviously, the differences increase when decreasing the bitrate. Among other (smaller) variations, the left half of the parameter range leads to poor filter quality. In Fig. 1(b) we display the PSNR quality when compressing the Lena

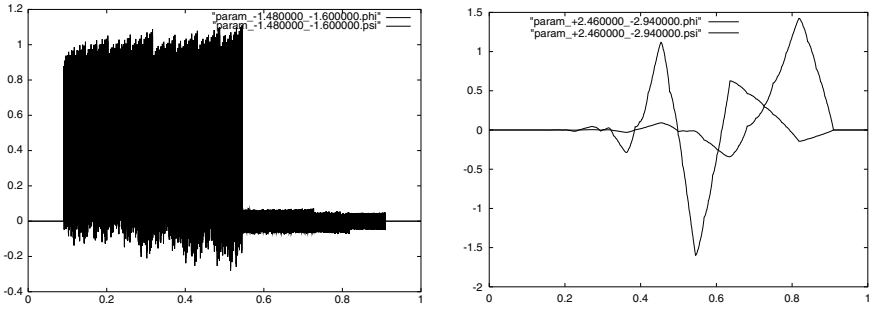

(a) 2-parameter filter giving poor compression results

(b) 2-parameter filter giving good compression results

Fig. 2. Filters generated by 2 parameters

image to 40000 bits (the same bitrate is used in all subsequent experiments) with different filters generated by the parametrization scheme using 2 parameters.

Similar behaviour with respect to quality variations can be observed in the two-dimensional case. Again the quality varies in a wide range, here we observe a minimum of 23.4 dB (generated by the filter shown in Fig. 2(a)) and a maximum of 32 dB (for the corresponding filter see Fig. 2(b)), with an average of 27.7 dB.

As a consequence of these findings, a strategy is required to limit the possible loss in compression quality introduced by randomly chosen parameters. The most desirable approach would be a heuristic which – given either the parameters to generate the filters or the actual filter coefficients themselves – could determine an approximation of the compression quality to be expected in advance (i.e. without performing the actual compression). Unfortunately, besides restricting the parameter to positive values in the one-dimensional case, no such heuristic could be found.

Generating the parameters and performing the actual compression stage to determine the corresponding quality is too time consuming (since only one failure in parameter choice (i.e. one bad quality filter) makes the scheme significantly more expensive than a full AES encryption of a classically encoded bitstream). Therefore, we determine parameter values of good quality in advance and restrict the admissible parameters to regions close to that values. Fortunately, the quality of parameters is very much image independent, which makes this approach a feasible and efficient one. However, the decrease of the amount of admissible parameter values is known in advance (also to a potential attacker) and reduces the overall security of the system since it corresponds to a smaller keyspace.

3.2 Security

The data type of the parameters is $\mathbb{R}$ (in theory), in practice it is $\mathbb{Q}$ which means we need to discretize the parameter space applying a fixed size grid onto it. Close parameters lead to similar filters which in turn lead to similar wavelet transform coefficients. Of course, this might be a threat to the security of the system since

an attacker does not need to know the compression parameter exactly to get a "decrypted" image with sufficient quality. Therefore, the discretization (i.e. the grid size) needs to be defined in a way that different parameters lead to sufficiently different filters.

In Fig. 3 we illustrate this problem. The Lena image is compressed with filters generated by six parameters, and subsequently decompressed with a large number of different filters derived from parameters covering the range $\pm\pi$ centered at the "correct" parameter. We plot the PSNR of the resulting images against the parameter used for decompression. The desired result would be an isolated single PSNR peak at the position of the "correct" parameter (that one used for compression) and low values everywhere else.

The result of this experiment is not an isolated PSNR peak but an entire region centered around the correct parameter where the PSNR values are decreasing with increasing distance from the correct value. In Fig. 4 we visualize images where the Lena image was compressed using the parameter 1 and decompressed with parameters displaced from the correct one by a certain amount.

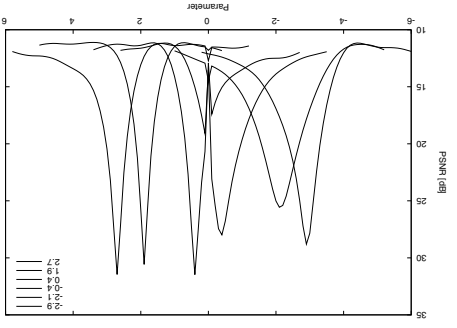

Fig. 3. Attack against a 1-D parameter scheme

(a) parameter distance of 0 (b) parameter distance of 0.2 (c) parameter distance of 1.0 (d) parameter distance of 3.1

Fig. 4. Images resulting from compression and decompression with similar parameterized filters

Obviously, the quality of the image in Fig. 4(b) is too high to provide any kind of confidentiality (compression and decompression paramters are too close), whereas the quality of Figs. 4(c) and 4(b) is low enough for applications requiring a low to medium confidentiality level.

As a consequence, the number of admissible parameter values needs to be restricted to a rather sparse discretization grid. Taken this fact together with the beforementioned restrictions due to low quality filters (subsection 3.1), of course the keyspace is too small for a reasonable application in case of the 1-D

parameter scheme. However, higher dimensional parametrizations can provide a sufficiently large amount of different parameters (see next subsection).

An attacker wanting to recreate the image without knowledge of the correct parameters (and without the plaintext image) has to test a grid spanning the search space of possible parameters and tries to deduce information about the correct parameters. In order to do this, the (possibly only partly) reconstructed images need to be automatically evaluated with respect to their "perceptual quality". One possibility to achieve this is to exploit the fact that incorrect filters generate more high-frequency noise. A simple technique in this context is to compute the differences between neighbouring pixels of image reconstructions.

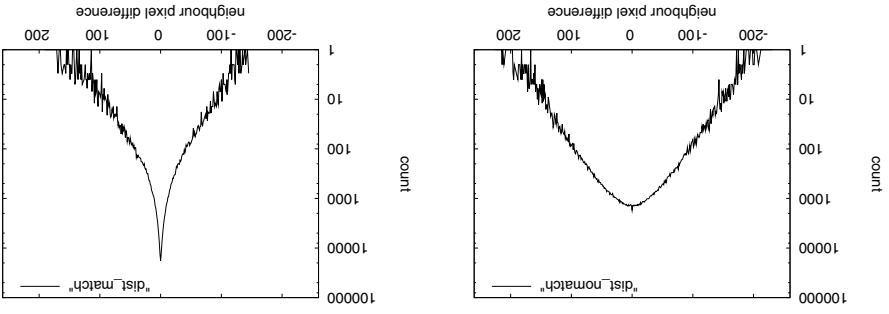

(a) matching filter (b) non-matching filter

Fig. 5. Statistical distribution of pixel differences for matching and non-matching filters in the 1-dimensional case

Fig. 5 shows histograms of the magnitude of the pixel differences, Fig. 5.a in case the correct filter is used for reconstruction, Fig. 5.b in case the filter used for reconstruction is far away from the correct one. In the matching case the amount of low differences is higher, and high differences do not exist, in the non-matching case the opposite holds. For automated search the distance to an histogram like that in Fig. 5.b can be minimized using a gradient technique, for example.

3.3 Higher-Dimensional Parametrizations

A possible strategy to increase the available keyspace significantly is to move to schemes with more parameters (leading to longer filters). Of course, the problem of varying filter qualities also exists in higher dimensions. To quantify the corresponding properties, we randomly select parameters and perform the compression with the same fixed bit limit as before and record for each number of parameters N the minimal, average, and maximal PSNR. Table 1 shows evidence that the maximum PSNR slightly increases, whereas the minimum and average values significanlty decrease for an increasing number of dimensions N. This result shows that for larger N a strategy to avoid low quality filters is even more important.

Table 1. Quality (PSNR in dB) tests with random parameters

#param	#samples	min	avg	max
1	25	24.94	28.36	31.50
2	69	23.40	27.49	31.44
3	193	22.22	26.84	31.99
4	603	21.67	26.37	32.14
5	1875	20.59	25.87	32.10
6	5470	20.14	25.57	32.07
7	8766	20.16	25.26	32.16
8	13185	19.88	24.98	32.09

On the other hand, the results available so far do not guarantee the existence of an increasing number of high quality filters for increasing dimension N. In Fig. 6 we plot the empirical cumulative distribution function of the filters with parameter dimensions $N = 1, 2, \ldots, 8$ with respect to their compression quality. The relative number of high quality filters obviously decreases with increasing N. For example, we see that for $N = 1$, about 45% of all filters show quality ≥ 30 dB, for $N = 2$ about 20%, and for $N = 8$ approximately 2% of all filters exhibit quality ≥ 30 dB.

However, considering the exponential growth of the overall number of available filters with respect to the dimension of the parametrization scheme (assuming a constant, dimension-independent discretization of the parameter range), a significant increase of the absolute number of high quality filters is guaranteed when increasing N.

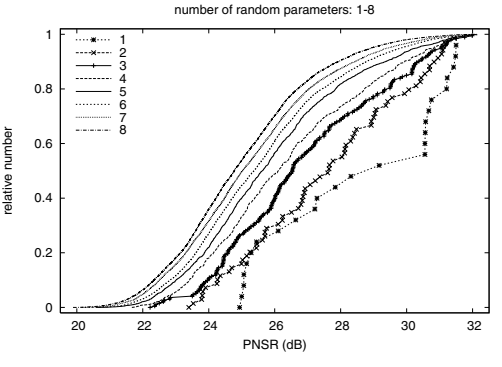

Fig. 6. Relative number of good and poor quality filters

But does the larger keyspace help to avoid an attack as described in the previous subsection? An exhaustive search through the parameter space is simply too costly, and varying all parameters simultanously during a random search does not enable a steepest descent search technique due to the high dimensionality. Once a single parameter combination with promising smoothness properties is found, it is not clear how to further improve the result. On the other hand, fixing all parameters but that in one dimension does not lead to clearly improved smoothness behaviour once that single parameter is near to its optimum. Therefore, a separable search (i.e. optimizing each parameter separately) is not successful as well. It turns out that all considered automated techniques fail to reliably identify an approximation to the correct parameter set for larger N due to the size of the key-space and the increasing non-linearity of the search space with increasing dimension.

4 Conclusions and Future Work

We have shown that higher dimensional wavelet filter parametrizations may solve the security problems of recently proposed lightweight encryption schemes for visual data. As the compression performance of the orthogonal filters in use is inferior to the standard biorthogonal filters employed within codecs, we will focus in future work on parametrization techniques directly related to the lifting scheme to integrate the approach into JPEG 2000.

References

[1] B. Bhargava, C. Shi, and Y. Wang. MPEG video encryption algorithms. *Multimedia Tools and Applications*, 24(1):57–79, 2004.

[2] Raphaël Grosbois, Pierre Gerbelot, and Touradj Ebrahimi. Authentication and access control in the JPEG 2000 compressed domain. In A.G. Tescher, editor, *Applications of Digital Image Processing XXIV*, volume 4472 of *Proceedings of SPIE*, pages 95–104, San Diego, CA, USA, July 2001.

[3] 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.

[4] R. Kutil. A significance map based adaptive wavelet zerotree codec (SMAWZ). In S. Panchanathan, V. Bove, and S.I. Sudharsanan, editors, *Media Processors 2002*, volume 4674 of *SPIE Proceedings*, pages 61–71, January 2002.

[5] Benoit M. Macq and Jean-Jacques Quisquater. Cryptology for digital TV broadcasting. *Proceedings of the IEEE*, 83(6):944–957, June 1995.

[6] P. Meerwald and A. Uhl. Watermark security via wavelet filter parametrization. In *Proceedings of the IEEE International Conference on Image Processing (ICIP'01)*, volume 3, pages 1027–1030, Thessaloniki, Greece, October 2001. IEEE Signal Processing Society.

[7] Roland Norcen and Andreas Uhl. Selective encryption of the JPEG2000 bitstream. In A. Lioy and D. Mazzocchi, editors, *Communications and Multimedia Security. Proceedings of the IFIP TC6/TC11 Sixth Joint Working Conference on Communications and Multimedia Security, CMS '03*, volume 2828 of *Lecture Notes on Computer Science*, pages 194 – 204, Turin, Italy, October 2003. Springer-Verlag.

[8] 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*, pages 1–4, Victoria, Canada, August 2001. IEEE Signal Processing Society.

[9] 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.

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

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

[12] A. Uhl and A. Pommer. *Image and Video Encryption. From Digital Rights Management to Secured Personal Communication*, volume 15 of *Advances in Information Security*. Springer-Verlag, 2005.