Are Parameterised Biorthogonal Wavelet Filters Suited (Better) for Selective Encryption?

Andreas Uhl
Department of Scientific Computing
University of Salzburg
J.Haringer Str.2
5020 Salzburg, Austria
uhl@cosy.sbg.ac.at

Andreas Pommer
Department of Scientific Computing
University of Salzburg
J.Haringer Str.2
5020 Salzburg, Austria
apommer@cosy.sbg.ac.at

ABSTRACT

Selective encryption is used to encrypt parts of a bitstream, in our case images which are compressed by a wavelet based method. One approach is to keep the filter secret which is used for the transformation. Parameterised wavelet filters can be used to generate a large keyspace, however, in the case of orthogonal filters obtained by a variant of Pollen's factorisation it turns out that different parameters yield filters with very different quality and in particular worse quality as compared to the standard biorthogonal filters usually used for compression. To eventually overcome these limitations, we consider parametrisations of biorthogonal filters in this work. We discuss methods to create such filters, and show their properties regarding compression and encryption.

Categories and Subject Descriptors

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

General Terms

Security, Performance, Experimentation

Keywords

selective encryption, image encryption, filter parameterisation, biorthogonal wavelet filters, image compression

1. INTRODUCTION

With the rise of wavelet based compression standards like JPEG2000 and MPEG-4 VTC wavelets have become very popular in the field of image compression. There are also many requests for means of confidential transmission of image data. The straight-forward approach is to compress the image and encrypt the complete compressed bitstream afterwards. In some application scenarios [15] this approach cannot be used, sometimes "selective encryption" can be the method of choice.

Selective encryption is a method to encrypt just some parts of the bitstream. Sometimes this approach is used because one tries to save CPU time, sometimes one sends an embedded bitstream with two versions of the data, an unencrypted low-quality version and encrypted better version for subscribers ("transparent encryption"). Common to all these selective encryption methods is the requirement to identify the parts of the bitstream which are crucial and therefore must be encrypted, and parts which are less important (or depend on encrypted crucial data) and therefore they can be left unencrypted. Usually this division is accompanied by an assessment of the security of the remaining bitstream: is it possible to reconstruct the encrypted data just with the help of the unencrypted data? How difficult is it to get an unencrypted estimation of the encrypted data? The provider of the data tries to exploit the structure of the data to maximise the efficiency of the compression algorithm (the usual thing when multimedia data compression is involved), but also tries to exploit the structure to maximise the efficiency of the encryption algorithm. An attacker tries to exploit this structure as well, but from a different point of view and usually with more restrictions concerning the knowledge of this structure.

When one wants to perform selective encryption on images compressed with wavelet methods there are two general approaches to accomplish this. The first approach is to process an already existing bitstream by simply encrypting parts of it. This has been discussed in the context of JPEG 2000 [1, 9]. The second approach is a special variant of header encryption. Wavelet-based compression can be potentially 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 [16] or NSMRA decomposition scheme [14] have been proposed.

Another possible approach in this direction is to keep the choice of filters secret which are used for the low- and high-pass filtering. There exist established filters for compression, it has been shown that they are good for the compression of a wide variety of images, but there are many more filters available to choose from. One example are entire families of

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
MM-SEC'04, September 20-21, 2004, Magdeburg, Germany.
Copyright 2004 ACM 1-58113-854-7/04/0009 ...$5.00.

filters, where individual filters are generated by an algorithm depending on one or more parameters. In recent work [5] we have used secret orthogonal filters derived from a variant of Pollen's factorisation in a lightweight JPEG 2000 encryption scheme. It turned out that the compression quality of these filters varies by a large extent and is particularly worse as compared to the standard biorthogonal filters usually used for compression. This is shortly reviewed in Section 2. As a possibility to overcome these problems, we investigate parametrisations of biorthogonal filters in this work. We describe their construction and discuss some problems we encountered in section 3. In section 4 we show some experimental results obtained by tests with these biorthogonal filters, in particular we show how good they perform when they must satisfy two requirements at the same time: being secure enough for selective encryption and performing a good at the compression. Finally we draw conclusions and give outlook to further work in this direction.

2. VARIATION 1: ORTHOGONAL FILTERS

The SMAWZ codec [6] used in our experiments is a variant of the well known SPIHT algorithm which has been optimised for efficient implementation using bitplanes instead of lists. In all wavelet based compression schemes (JPEG 2000, MPEG-4 VTC, SPIHT, SMAWZ), 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 parameterised filters generated according to an algorithm proposed by Schneid and Pittner [18]:

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

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

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

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.

When we use this selective encryption scheme to transmit an image we have to encrypt the key which is used to encrypt the remaining data. In this case the key is the set of parameters to generate the parameterised filters. Everything else, like the actual coefficients, can be transmitted as it comes out of the compression algorithm. This is the benefit of such a selective encryption scheme: there is no need to spend computing time to encrypt the complete bitstream, it is sufficient to encrypt the filter generation parameters, here 32 bits per parameter are enough. On the other hand we loose some of the security when compared to full encryption. The question is just how easy it is for an attacker to reconstruct the original image. Within the right threat model and other requirements this selective encryption scheme for images is can be a solution.

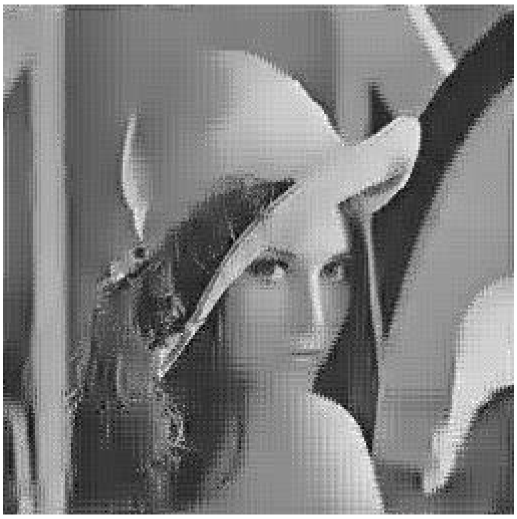

Figure 1: image obtained using a 1-parameter filter set with similar characteristics as in figure 2(a)

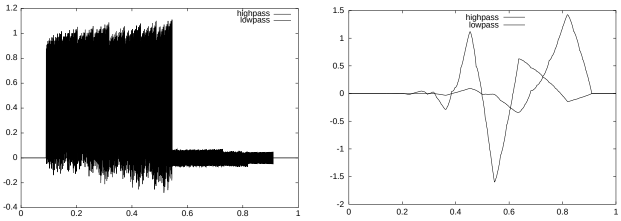

(a) filter giving poor compression results

(b) filter giving good compression results

Figure 2: Filters generated by 2 parameters

2.1 Some results

Our experiments with these parameterised orthogonal filters showed that the quality varies by a large amount, it depends heavily on the filter. For a filter giving good results the uncompressed image looks as expected: no visible degradation, an example for a good filter and its frequency response is shown in figures 2(b) and 3(b). An example compression result for a poor filter is shown in figure 1, a poor quality filter itself and its frequency response in shown in figures 2(a) (oscillating filter) and 3(a). This figures of frequency response give an indication about the expected compression quality performance: For a good compression with wavelets it is necessary that the original signal is partitioned in a recursive manner into the high and low frequency parts. Filters which perform good show a good frequency separation: the lowpass filter allows all low frequencies to pass (ideally all at the same level) and does not let pass any of the high frequencies, and the reverse is true for the highpass filter.

For more detailed results (especially also focusing on the security of such a system) please refer to [5] and [13]. The

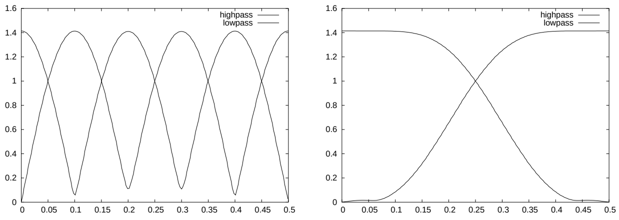

(a) frequency response for the poor filter

(b) frequency response for the good filter

Figure 3: Frequency response for the good and the bad filter, as shown in figure 2.

bottom line is that the quality varies by a large amount and that it is not easy to predict the parameters which give reasonably good results. Additionally, even the best quality filters are significantly inferior to the biorthogonal filters usually used in wavelet codecs. This of course reduces the applicability for selective encryption: for a legitimate user the received image should have good quality!

3. VARIATION 2: BIORTHOGONAL FILTERS

In the last section we showed that the orthogonal filters do not perform sufficiently well at the compression quality. So we should look for some replacement and biorthogonal filters come to our mind. In the following we try the same selective encryption approach as before, but with a different filter class.

Theoretical and practical work in the field of image and video compression usually prefer biorthogonal filters. In most compression applications the well-known “Biorthogonal 7,9” filter is used. Therefore our hope was that this filter is not an exception but an indication about the superiority of the overall class of biorthogonal filters in this context [19].

For the creation of biorthogonal filters we relied mainly on a paper by Hartenstein [2] because a method was presented which allowed an “easy” implementation: no symbolic computations with programs like *Mathematica* or *Mathlab* were required. Such a prerequisite would have negated to requirement that many tests with different parameters should be performed, and that the program should be able to perform a very quick filter exchange. Additionally, this application should fit into the context of the C++ based compression library and framework developed at our department (called `libganesh++`) which includes the aforementioned SMAWZ-codec. Besides Hartenstein some other authors have proposed additional methods for parameterising biorthogonal wavelet filters [17, 11, 12, 8, 7, 4, 3, 10].

3.1 Generation of Even-Length Filters

Even length filters require that the difference between high- and low-pass filter is a multiple of 4, i.e. $4K$. The general formula to generate these filters is

$$\begin{bmatrix} H(z) \\ G(z) \end{bmatrix} = H_p(z^2) \begin{bmatrix} 1 \\ z^{-1} \end{bmatrix} \tag{1}$$

$$\text{with } H_p(z) = A\Lambda(z)S_{L-1}\Lambda(z)\dots\Lambda(z)S_0 \tag{2}$$

$$\text{and } \Lambda(z) = \begin{bmatrix} 1 & 0 \\ 0 & z^{-1} \end{bmatrix} \tag{3}$$

$$\text{and } S_i = \frac{1}{\cos^2\theta_i - \sin^2\theta_i} \begin{bmatrix} \cos\theta_i & \sin\theta_i \\ \sin\theta_i & \cos\theta_i \end{bmatrix}. \tag{4}$$

It is obvious that the denominator above must not be 0, therefore the θ_i are limited to $\theta_i \neq (2k+1)\frac{\pi}{4}, k \in \mathbb{Z}$. Hartenstein had two errors in his paper, one in his equation (2) where he had an excess matrix S_L, and the other one in his equation (3) where the restriction for θ_i was too strict. Additionally, he didn’t care about the energy-preserving property of his matrices: the value of the determinant must be 1. So Hartenstein didn’t give the fraction part for the matrices S_i. The above limitation for θ_i has to be extended in practise, so that it can be formulated like “θ_i should not lie within a neighbourhood of ϵ, centred at odd multiples of π”. The result is undefined right at these multiples, but within the neighbourhoods numerical instabilities occur which make it difficult to calculate reasonable results. We discovered that ϵ must be increased for increasing absolute values of K.

We can distinguish between three cases, for each one a different Matrix A must be constructed:

$\mathbf{K = 0}$: this is the simplest case

$$A = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix} \tag{5}$$

$\mathbf{K > 0}$: the high-pass filter is longer than the low-pass filter

$$A = \begin{bmatrix} 1 & 1 \\ P_{0K}(z) & P_{1K}(z) \end{bmatrix} \tag{6}$$

(this was equation (5) in the Hartenstein paper)

$$P_{0K} = \begin{cases} 1 & K = 0 \\ 1 + \tan\beta_1 z^{-1} - z^{-2} & K = 1 \\ P_{01}(z^k) + \sum_{j=2}^{K} z^{j-K-1} q_j(z) & K > 1 \end{cases}$$

$$P_{1K} = \begin{cases} -1 & K = 0 \\ 1 - \tan\beta_1 z^{-1} - z^{-2} & K = 1 \\ P_{11}(z^k) + \sum_{j=2}^{K} z^{j-K-1} q_j(z) & K > 1 \end{cases}$$

with $q_j(z) = \tan\beta_j - \tan\beta_j z^{2(1-j)}$.

$\mathbf{K < 0}$: the low-pass filter is longer, the formula is almost identical (and we set $K = -K$ to be positive):

$$A = \begin{bmatrix} Q_{0K}(z) & Q_{1K}(z) \\ 1 & -1 \end{bmatrix} \tag{7}$$

with $Q_{0K}(z) = P_{0K}(z)$ and $Q_{1K}(z) = -P_{1K}(z)$

Of course the matrix A must be normalised again, otherwise the subsequently generated filter will not preserve the energy of the signals.

3.2 Generation of Odd-Length Filters

In [2] Hartenstein shows a construction method for parameterised biorthogonal wavelet filters with an odd number of coefficients, too. The length difference between decomposition and reconstruction filter is a multiple of 2, but not a multiple of 4: diff $= 2 * (2N + 1)$. However in this paper the step from equation (10) to (11) is not correct, therefore his method cannot work. Hartensteins equation (10) is as follows:

$$A = \left[\begin{array}{cc} 1 + z^{-1} & \tan\alpha \\ \sum_{i=0}^{K+1} b_i z^{-i} + b_i z^{i-2(K+1)} & \sum_{i=0}^{K} c_i z^{-i} + c_i z^{i-2K-1} \end{array} \right] \tag{8}$$

Hartenstein requires that the determinant of this matrix should be monomial and therefore comes to his equation (11), which looks as follows:

$$A = \left[\begin{array}{cc} 1 + z^{-1} & \tan\alpha \\ P_{0K}(z) & P_{1K}(z) \end{array} \right], \tag{9}$$

$$P_{0K}(z) = \sum_{i=0}^{K+1} \tan\beta_i z^{-i} + \tan\beta_i z^{i-2K-2}, \tag{10}$$

$$P_{1K}(z) = \sum_{i=0}^{K} a_i z^{-i} + a_i z^{i-2K-1}, \tag{11}$$

$$a_i = \tan\alpha \sum_{j=0}^{i} (-1)^{i+j} \tan\beta_j \qquad i < K \text{ and} \tag{12}$$

$$a_K - \tan\alpha \tan\beta_{K+1} \neq 0. \tag{13}$$

The user has to provide $K + 4$ parameters: $\alpha, a_K, \beta_i, i = 0, \ldots, K + 1$.

However, when we recalculated his step from equation (8) to (9) we discovered that the determinant of the resulting matrix A is not monomial: Even for the most simple version of $K = 0$ the determinant is no monomial in general, the same applies for $K = 1$:

$$\begin{aligned} \det(A) \;\; = \;\; & (a_1 + \tan\alpha \tan\beta_0 - \tan\alpha \tan\beta_1) z^{-1} + \\ & 2(a_1 - \tan\alpha \tan\beta_2) z^{-2} + \\ & (a_1 + \tan\alpha \tan\beta_0 - \tan\alpha \tan\beta_1) z^{-3} \end{aligned} \tag{14}$$

with the user-supplied parameters $\alpha, a_1, \beta_0, \beta_1, \beta_2$.

So we see that the statements are not generally true. Therefore we created a new version of A with the following coefficients for the polynoms in equation (8), this version of A has a monomial determinant:

$$c_i = \tan\alpha \tan\beta_i \text{ for } 0 \leq i \leq K, \tag{15}$$

$$b_i = \tan\beta_{i-1} + \tan\beta_i \text{ for } 0 < i \leq K, \tag{16}$$

$$\text{and } b_0 = \tan\beta_0 \tag{17}$$

Here we need $K+3$ parameters: α, b_{K+1}, β_i for $0 \leq i \leq K$ with the limitations: $\alpha \neq 0, b_{K+1} \neq \tan\beta_K$. For a more homogeneous style we can rephrase the last sentence: we need $K + 3$ parameters: α, β_i for $0 \leq i \leq K + 1$ with the limitations: $\alpha \neq 0, b_{K+1} = \tan\beta_{K+1} \neq \tan\beta_K$. And now all parameters lie in the same open interval $(-\frac{\pi}{2}, \frac{\pi}{2})$.

Despite these attempts we could not create filters which meet the usual specifications. The first necessary condition is that the filtering and the subsequent inverse filtering recreate the original (some minor differences are allowed because of the involved numerics, of course). Hartenstein gives some conditions to create only low-/high-pass pairs (following his equation (11)), but gives no rationale for those equations. Additionally it seems that the equations lack some symmetry. Besides that it is our observation that the remarks of Hartenstein are quite unspecific at this part, it looks like Hartenstein himself was not very confident about that part.

4. EXPERIMENTAL RESULTS

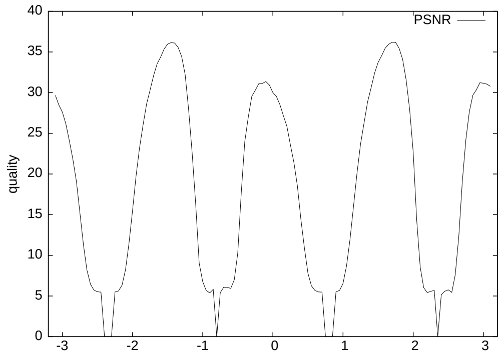

Figure 4: quality values for $L = 1, K = 0$ (one θ value).

In this section we look at the quality that we can achieve when we use parameterised biorthogonal filters as described earlier. Our initial hope was that these filters turn out to be better than the parameterised orthogonal filters which we have researched earlier.

The quality obtained by compression using even-length parameterised biorthogonal filters according to the construction method of Hartenstein varies by a very large amount, even more as compared to the orthogonal case. As can be seen on figures 4 and 5 the maximum value is near 35dB, but values go down to 5 dB as well. Note that at the instances where the filter could not be generated because of numerical instabilities a quality value of 0 was assumed. Some of the figures shown in this paper were generated from the results with tests with the Lena image, some with the baboon image — there are no significant differences between these two result sets. In all experiments the images were compressed with a target bitrate of 80000 bits, this leads to a compression rate of about 6.5 for 8-bit gray-level images with 256*256 pixels.

First we look at the most simple case where both filters have the same length ($K = 0$). We examine the results with $L = 1$ and $L = 2$. Figure 4 shows the first case, we observe a very high variance of the PSNR values. Figure 5 shows the results obtained by setting $L = 2$, the results look very similar to the previous figure. One can also observe some regular pattern with high-quality areas which could be used for later encryption tests.

When we compare the figures 6 and 7 we see a difference

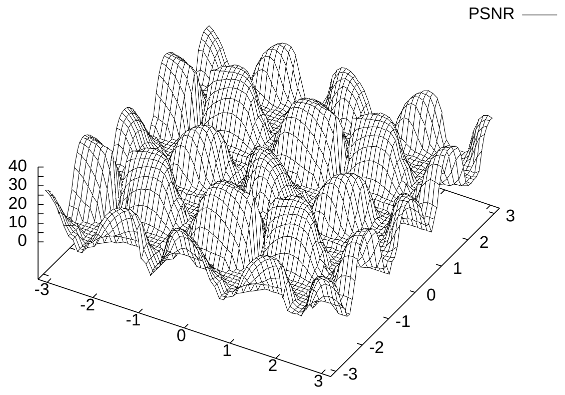

Figure 5: quality values for $L = 2, K = 0$ (two θ_i values).

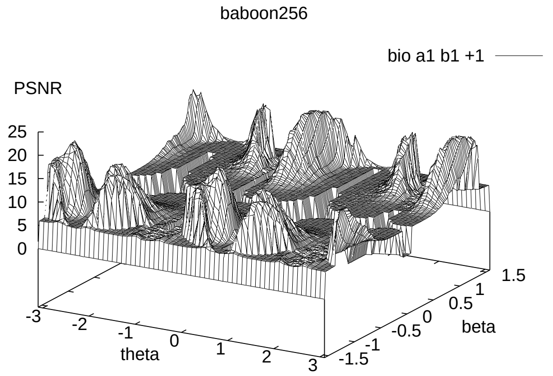

Figure 6: parameterised biorthogonal 4/8 filters: $K = +1, L = 1$.

in the maximum PSNR of about 10dB: 24.7 versus 34.6 on the other hand. This shows that it is important to make the right choice between a long high-pass filter together with a short low-pass filter or the short high-pass with the long low-pass filter: the variant with the shorter high-pass filter for decomposition is the better choice. Figure 8 shows the frequency response for both filters. We see that the frequency separation fails almost completely in figure 8(a): the low-pass filter allows certain higher frequencies to pass, but the high-pass filter blocks them. A second point is the high variation in the amplitude, the optimum (like shown in figure 3(b)) should be a level of $\sqrt{2}$ for the pass-band and it should close to 0 for the stop-band. The quality of a wavelet compression scheme is connected to the quality of the filter it uses, as long as the filter do not handle the frequency separation well the compression quality will be low.

Another interesting point is the comparison of the aforementioned filters with the well-known 7/9 filter: the PSNR in the same experiment lies at 37.7dB. Figure 9 shows the frequency response of the 7/9 filter. In comparison to figure 8 this looks much better: higher degree of symmetry, and the frequency separation into two bands is much higher. On the other hand the symmetry is lower when we compare it with figure 3(b), so this attribute can be a hint towards high-quality filters but it is widely known that there are other parameters important for compression performance.

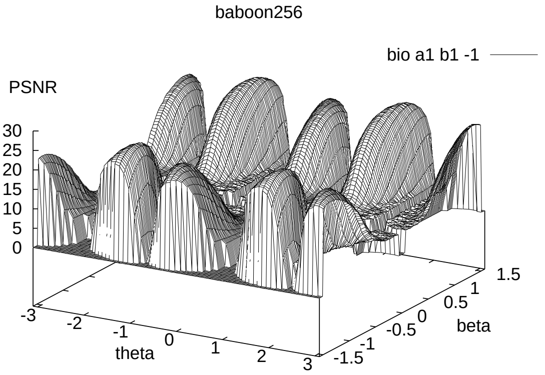

Figure 7: parameterised biorthogonal 4/8 filters: $K = -1, L = 1$.

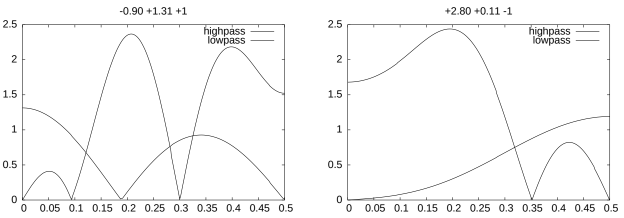

(a) frequency response of the filter giving the best result in figure 6.

(b) frequency response of the filter giving the best result in figure 7.

Figure 8: frequency response

So we see that when one wants to implement parameterised biorthogonal filters for selective encryption one will be faced with a decreased quality when using the same compression rate, at least when applying even-length filters derived from Hartensteins parametrisation.

4.1 Key Space

Since the input parameter is a real number in the range of $[-\pi, +\pi]$ it is theoretically possible to have an infinite key space. Filters generated by parameters which lie close to each other are very similar. This means that when the attacker hits a parameter very close to the original he will get almost the same filter and subsequently almost the same decompressed image as the original.

So an initial brute force search with a step size of say 0.1 will need 63 attempts to scan the interval. Then it would be possible to refine the search near the best matches. One problem is to determine the best match in a general manner. One heuristic is to measure the amount of smoothness in the generated image, images which close parameters tend to have intensity differences between neighbouring pix-

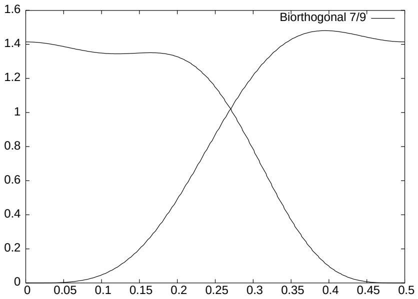

Figure 9: Frequency response of the Biorthogonal 7/9 filter

els which are smaller than at images generated without a good parameter. Each attempts includes a complete wavelet reconstruction (usually 5 levels), and some postprocessing of the complete image for the heuristics. This adds a huge multiplicative constant to the complexity formula, the attack with one filter parameter is linear: $O(s)$ where s is the number of steps.

The counter measure against this brute force search is to increase the search space. This can be achieved by generating filters using more than 1 parameter. With n parameters the search space is $[-\pi, +\pi]^n$. With such an increasing search space a brute force attack must also exponentially increase the number of attempts to $O(s^n)$. The attacker faces two problems: he does not know n, and the second problem is that the attack complexity is exponential. To make this attack feasible at higher dimensions the size of the steps must be decreased (since he cannot change n). This leads to less accurate results since the linearity of the search space decreases with the increasing number of dimensions, therefore the chances increase that the attacker will miss the correct choice. Since the attacker does not know n he must guess or try. The best advice for the attack attempts is to start with n=1 and increase it by 1 until the attacker can be reasonably sure to have recovered the correct parameter. The downside of choosing a bigger value for n is that with this increasing number of parameters the expected average quality of the filter decreases. So in the role of the provider one wants longer filter to be more secure, but also wants shorter filters to achieve a better compression quality. See [13] for some information on this topic.

When a user wants to apply selective encryption with such filters it is reasonable to exclude areas of bad quality. An attacker knows this, too, so he can reduce the search for the correct parameters to the areas of high quality. This reduces the key search space significantly, e.g. when we look at figure 4 and when we set a quality limit of 25 dB the set of admissible parameters is reduced to approximately 40%. This reduction increases with the number of parameters n.

5. CONCLUSIONS

For selective encryption we have researched the creation of parameterised wavelet filters, first of orthogonal filters, then biorthogonal filters. We tried to create several flavours of biorthogonal filters, first with even-length where we had some promising results, then with odd-length filters where we encountered some problems. Such parameterised filters were researched how arbitrary parameters affect the compression rate, the security and also important the compression performance.

We saw that a completely random choice of the filter generation parameter(s) is no good idea: Even-length biorthogonal wavelet filters derived from a parameterisation proposed by Hartenstein have turned out to give extremely varying (and also generally poor) compression results thus making them inappropriate for a selective encryption approach which only protects the filters in use during wavelet decomposition and compression. However the initial question is still open, there are indicators that at least some biorthogonal filters are better suited for selective encryption than orthogonal filters are.

In future work we will focus on the question whether there exist further parametrisations of biorthogonal wavelet filters in literature which are correct, which can be implemented in a programming language like C, and which provide a reasonable compression performance across a wide range of filters to make them useful in the context of a selective encryption application.

6. ACKNOWLEDGEMENTS, DISCLAIMER

This work was partially funded by the Austrian Science Fund FWF, project number P15170.

The work described in this paper has been supported in part by the European Commission through the IST Programme under Contract IST-2002-507932 ECRYPT. The information in this document reflects only the author's views, is provided as is and no guarantee or warranty is given that the information is fit for any particular purpose. The user thereof uses the information at its sole risk and liability.

7. REFERENCES

[1] 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 *Proceedings of SPIE*, pages 95–104, San Diego, CA, USA, July 2001.

[2] F. Hartenstein. Parametrization of discrete finite biorthogonal wavelets with linear phase. In *Proceedings of the 1997 International Conference on Acoustics, Speech and Signal Processing (ICASSP'97)*, Apr. 1997.

[3] B. Jawerth and W. Sweldens. Biorthogonal smooth local trigonometric bases. *J. Fourier Anal. Appl.*, 2(2):109–133, 1995.

[4] J. Kautsky and R. Turcajova. Pollen product factorization and construction of higher multiplicity wavelets. *Linear Algebra and its Applications*, 222:241–260, 1995.

[5] 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. To appear.

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

[7] S. Maslakovic, I. R. Linscott, M. Oslick, and J. D. Twicken. A library-based approach to design of

smooth orthonormal wavelets. In *Proceedings of the IEEE Digital Signal Processing Workshop, DSP '98*, Bryce Canyon, USA, Aug. 1998.

[8] S. Maslakovic, I. R. Linscott, M. Oslick, and J. D. Twicken. Smooth orthonormal wavelet libraries: design and application. In *Proceedings of the 1998 International Conference on Acoustics, Speech and Signal Processing (ICASSP'98)*, pages 1793–1796, Seattle, WA, USA, May 1998.

[9] R. Norcen and A. 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, Oct. 2003. Springer-Verlag.

[10] J. Odegard and C.S.Burrus. Smooth biorthogonal wavelets for applications in image compression. In *IEEE DSP Workshop, Loen, Norway*, Sept. 1996.

[11] M. Oslick, I. R. Linscott, S. Maslakovic, and J. D. Twicken. A general aproach to the generation of biorthogonal bases of compactly-supported wavelets. In *Proceedings of the 1998 International Conference on Acoustics, Speech and Signal Processing (ICASSP'98)*, Seattle, WA, USA, May 1998.

[12] S.-M. Phoong, C. Kim, P. Vaidyanathan, and R. Ansari. A new class of twochannel biorthogonal filter banks and wavelet bases. *IEEE Transactions on Signal Processing*, 43(3), Mar. 1995.

[13] A. Pommer. *Selective Encryption of Wavelet-compressed Visual Data*. PhD thesis, University of Salzburg, Austria, June 2003.

[14] 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, Aug. 2001. IEEE Signal Processing Society.

[15] A. Pommer and A. Uhl. Application scenarios for selective encryption of visual data. In J. Dittmann, J. Fridrich, and P. Wohlmacher, editors, *Multimedia and Security Workshop, ACM Multimedia*, pages 71–74, Juan-les-Pins, France, Dec. 2002.

[16] 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.

[17] H. L. Resnikoff, J. Tian, and R. O. Wells. Biorthogonal wavelet space: parametrization and factorization. *SIAM Journal on Mathematical Analysis*, Aug. 1999.

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

[19] J. D. Villasenor, B. Belzer, and J. Liao. Wavelet filter evaluation for image compression. *IEEE Transactions on Image Processing*, 4(8):1053–1060, Aug. 1995.