Selective Encryption of the JPEG2000 Bitstream

Roland Norcen and Andreas Uhl

Department of Scientific Computing, Salzburg University,
Jakob-Haringer-Str. 2, A-5020 Salzburg, Austria

Abstract. In this paper, we propose partial encryption of JPEG2000 coded image data using an AES symmetric block cipher. We find that encrypting 20% of the visual data is sufficient to provide a high level of confidentiality. This percentage of encryption also provides security against replacement attacks, which is discussed at length.

Introduction

Images and videos (often denoted as visual data) are data types which require enormous storage capacity or transmission bandwidth due to the large amount of data involved. In order to provide reasonable execution performance for encrypting such large amounts of data, only symmetric encryption (as opposed to public key cryptography) can be used. The Advanced Encryption Standard AES [3] is a recent symmetric block cipher potentially used in such applications. Nevertheless, real-time encryption of an entire video stream using such symmetric ciphers requires much computation time due to the large amounts of data involved.

Since many multimedia applications require security on a low level (e.g. TV broadcasting [6]) or should protect their data just for a short period of time (e.g. news broadcast), faster encryption procedures specifically tailored to the target environment should be designed for such multimedia security applications.

A good step in this direction is the introduction of a "soft" encryption scheme. Such a scheme does not strive for maximum security and trades off security for computational complexity.

Selective or partial encryption (SE) of visual data is an example for such an approach. Here, application specific data structures are exploited to create more efficient encryption systems (see e.g. SE of MPEG video streams [1,5,10,11,13, 19], of wavelet-based encoded imagery [2,4,8,9,18,19], and of quadtree decomposed images [2]). Consequently, SE only protects (i.e. encrypts) the visually most important parts of an image or video representation relying on a secure but slow "classical" cipher.

In this work we discuss the selective encryption of JPEG2000 coded image data using an symmetric AES cipher. Section 1 introduces the JPEG2000 algorithm with a special focus on the bitstream assembling part, since this is the starting point for our selective encryption approach. Then, in section 2, our partial encryption scheme of the JPEG2000 bitstream is presented and the obtained results are discussed in detail with emphasis on the achieved security.

A. Lioy and D. Mazzocchi (Eds.): CMS 2003, LNCS 2828, pp. 194–204, 2003.

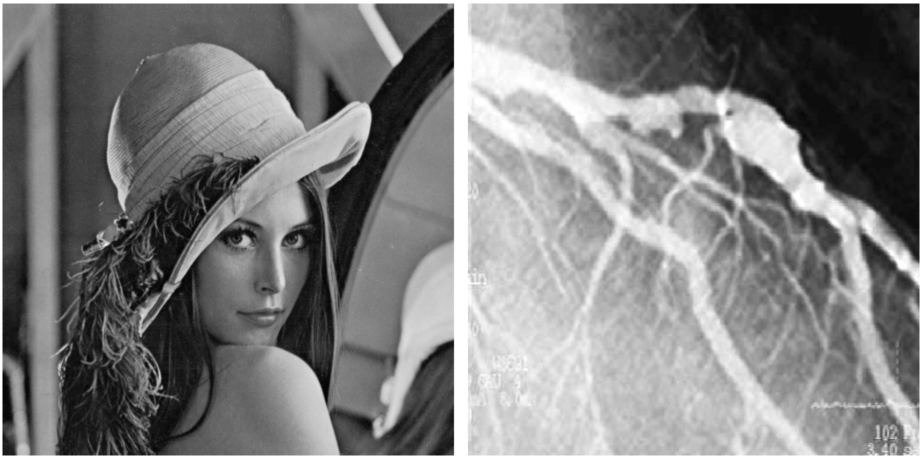

(a) Lena (256 × 256 and 512 × 512 pixels)

(b) Angiogram (512 × 512 pixels)

Fig. 1. Testimages used in the experiments

In Fig. 1 we display the testimages which we use in our experiments. We have decided to discuss the encryption efficiency of our proposed partial encryption scheme with respect to two different classes of visual image data. The first class of visual data discussed is typical still image data and the testimage representing this class is the lena image (see Fig. 1.a) at different resolutions including 256×256 and 512×512 pixels. Since this special type of visual data is usually encoded in lossy mode, the lena image is lossy coded in our experiments (at a fixed rate of 2 bpp). The second type of digital visual data should represent an application class where lossless coding is important. We have therefore decided to use an angiogram as testimage in this case (see figure 1.b), since angiograms represent an important class of medical image data where lossless coding is usually a mandatory requirement.

In order to make the explanations and experiments of the proposed techniques simpler, we assume the testimages to be given in 8bit/pixel (bpp) precision and in a squared format. Extensions to images of different acquisition types (e.g. [17]), higher bitdepth or non-squared format are straightforward.

1 JPEG2000

Image compression methods that use wavelet transforms [16] (which are based on multiresolution analysis – MRA) have been successful in providing high compression ratios while maintaining good image quality. Therefore, they have replaced

DCT based techniques in recent standards for still image coding: JPEG2000 [15] and VTC (visual texture coding in MPEG-4 [12]).

The JPEG2000 image coding standard is based on a scheme originally proposed by Taubman and known as EBCOT ("Embedded Block Coding with Optimized Truncation" [14]). The major difference between previously proposed wavelet-based image compression algorithms such as EZW or SPIHT (see [16]) is that, after performing a global wavelet transform, EBCOT as well as JPEG2000 operate on independent, non-overlapping blocks of transform coefficients which are coded in several bit layers to create an embedded, scalable bitstream. Instead of zerotrees, the JPEG2000 scheme depends on a per-block quad-tree structure since the strictly independent block coding strategy precludes structures across subbands or even code-blocks. These independent code-blocks are passed down the "coding pipeline" shown in Fig. 2 and generate separate bitstreams. The wavelet coefficients inside a code-block are processed from the most significant bitplane towards the least significant. Furthermore, in each bitplane the bits are scanned in a maximum of three passes called coding passes. Finally, during each coding pass, the scanned bits with their context value are sent to a context-based adaptive arithmetic encoder that generates the code-block's bitstream. This procedure is called Tier-1 encoding.

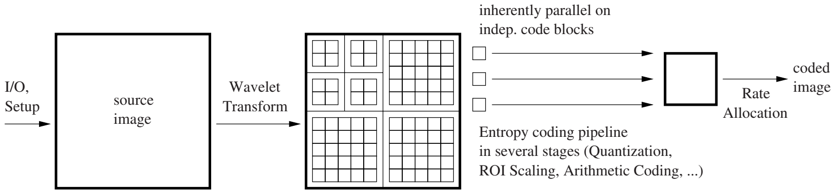

Fig. 2. JPEG2000 coding pipeline

The rate-distortion optimal merging of these bitstreams into the final one is based on a sophisticated optimization strategy and is called Tier-2 encoding. This last procedure carries out the creation of the so-called layers which roughly stand for successive qualities at which a compressed image can be optimally reconstructed. These layers are built in a rate-allocation process that collects, in each code block, a certain number of coding-passes codewords. Hence, in a code-block, the bitstream is distributed into a certain number of layers.

The final JPEG2000 bitstream is organized as follows: A set of different main headers (including a main header (SIZ), a coding style header (COD), a quantization header (QCD), a comments header (COM), a start of a tile parts header (SOT)) is followed by packets of data which are all preceded by a packet header. In each packet appear the codewords of the code-blocks that belong to the same image resolution and layer, the header identifies the data. Depending on the arrangement of the packets, different progression orders may be specified. Among others, resolution and layer progressive are most important for grayscale images. In layer progression order, the packets corresponding to the first layer are arranged first and cover data contained in all resolutions, followed by packets corresponding to the second layer and so on. Vice versa, in resolution progression order the packets corresponding to the first resolution level are arranged first (these contain data of all layers).

2 Selective Encryption of JPEG2000 Coded Data

For selectively encrypting the JPEG2000 bitstream we have two general options. First, we do not care about the structure of the bitstream and simply encrypt a part, e.g. the first 10% of the bitstream. In this case, the main header and a couple of packets including packet header and packet data are encrypted. Since basic information necessary for reconstruction usually located in the main header is not available at the decoder, encrypted data of this type can not be reconstructed using a JPEG2000 decoder. Although this seems to be desirable at first sight, an attacker could reconstruct the missing header data using the unencrypted parts, and, additionally, no control over the quality of the remaining unencrypted data is possible. Therefore, the second option is to design a JPEG2000 bitstream format compliant encryption scheme which does not encrypt main and packet header but only packet data. This is what we propose in the following.

In order to achieve format compliance, we need to access and encrypt data of single packets. Since the aim is to operate directly on the bitstream without any decoding we need to discriminate packet data from packet headers in the bitstream. This can be achieved by using two special JPEG2000 optional markers which were originally defined to achieve transcoding capability, i.e. manipulation of the bitstream to a certain extent without the need to decode data. Additionally, these markers of course increase error resilience of the bitstream. These markers are “start of packet marker” (SOP - 0xFF91) and “end of packet marker” (EPH - 0xFF92). The packet header is located between SOP and EPH, packet data finally may be found between EPH and the subsequent SOP. For example, using the official JAVA JPEG2000 reference implementation (JJ2000 - available at `http://jj2000.epfl.ch`) the usage of these markers may be easily invoked by the options `-Peph on -Psop on`.

Having identified the bitstream segments which should be subjected to encryption we note that packet data is of variable size and does not at all adhere to multiples of a block ciphers block-size. We have to employ AES in CFB mode for encryption, since in this mode, an arbitrary number of data bits can be encrypted, which is not offered by the ECB and CBC encryption modes. Information about the exact specification of the cryptographic techniques used (e.g. key exchange) may be inserted into the JPEG2000 bitstream taking advantage of so-called termination markers. Parts of the bitstream bounded by termination markers are automatically ignored during bitstream processing and do not interfere with the decoding process. Note that a JPEG2000 bitstream which is selectively encrypted in the described way is fully compliant to the standard and can therefore be decoded by any codec which adheres to the JPEG2000 specification.

We want to investigate whether resolution progressive order or layer progressive order is more appropriate for selective JPEG2000 bitstream encryption. We therefore arrange the packet data in either of the two progression orders, encrypt an increasing number of packet data bytes, reconstruct the images and measure the corresponding quality.

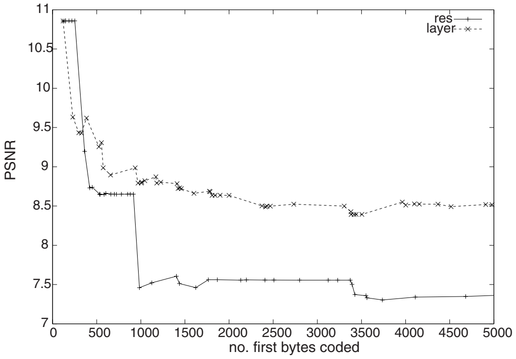

Fig. 3. Angiogram: Comparison of selective encryption (PSNR of reconstructed images) using resolution or layer progressive encoding - part 1.

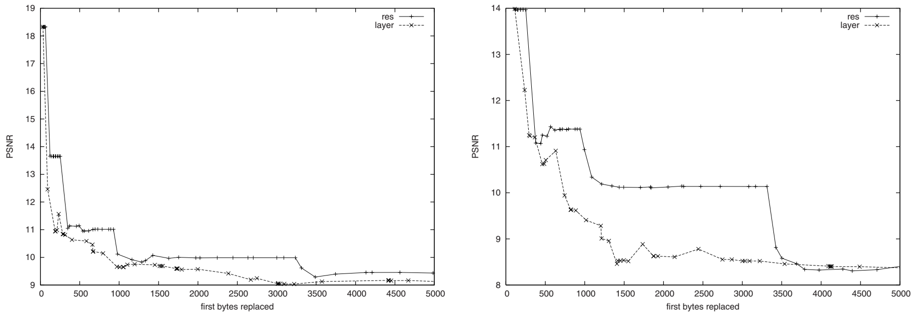

(a) Lena 256 × 256 pixels

(b) Lena 512 × 512 pixels

Fig. 4. Comparison of selective encryption (PSNR of reconstructed images) using resolution or layer progressive encoding - part 2.

Resolution progression is more suited for selectively encrypting the angiogram image at higher rates of encrypted data (see Fig. 3, where a lower PSNR means that it is more suited for selective encryption). In order to relate the obtained numerical values to visual appearance, two reconstructed versions of the angiogram, corresponding to the two progression orders, are displayed in Fig. 5. In both cases, 1% of the entire packet data has been encrypted.

Whereas no details are visible using layer progression (Fig. 5.a at 8.79 dB), only very high frequency visual information (right lower corner) is visible using resolution progression (Fig. 5.b at 7.45 dB).

When considering the Lena image in Fig. 4, we observe that resolution progression shows superior PSNR results for both tested image dimensions as compared to layer progression. Two reconstructed versions of the Lena image with 512 × 512 pixels, corresponding to the two progression orders, are displayed in Fig. 6. In each case, 1% of the entire packet data has been encrypted. Whereas only very high frequency information is visible in the reconstructed image using layer progression (Fig. 6.a at 8.51 dB), important visual features are visible using resolution progression (Fig. 6.b at 10.11 dB). In this case, the visible high frequency information is enough to reveal sensible data. At 2 percent encrypted packet data, this information is destroyed fully in the resolution progressive case.

The lena image at lower resolution (256 × 256 pixels) performs equally, and the results are therefore only given for the 512^2 pixels version.

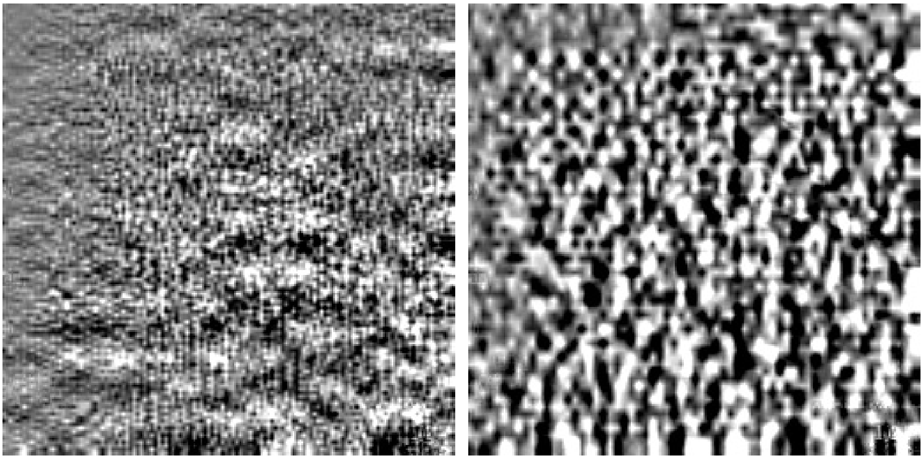

(a) layer progressive, 8.79 dB (b) resolution progressive, 7.45 dB

Fig. 5. Comparison of selective encryption (visual quality of reconstructed Angiogram where 1% of the bitstream data have been encrypted) using resolution or layer progressive encoding.

Please note also the difference in coarseness of the noise pattern resulting from encryption between resolution and layer progression. Since in resolution

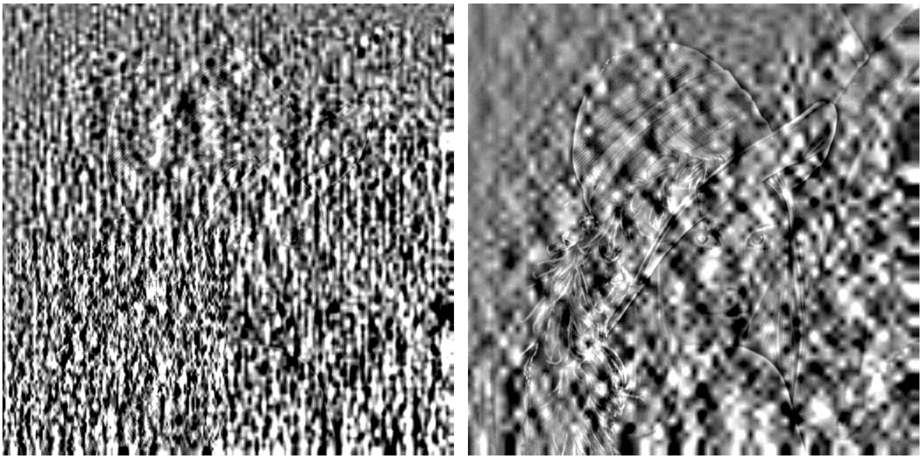

(a) layer progressive, 8.51 dB

(b) resolution progressive, 10.11 dB

Fig. 6. Comparison of selective encryption (visual quality of reconstructed Lena (512 pixels) where 1% of the bitstream data have been encrypted) using resolution or layer progressive encoding.

progression data corresponding to the higher levels of the wavelet transform is encrypted, the noise introduced by the cipher is propagated by the repeated inverse transform and thereby magnified resulting in a much coarser pattern as compared to layer progression. When summarizing the obtained numerical and visual results, it seems that encrypting 1-2% of the packet data in layer progressive mode is sufficient to provide confidentiality for the JPEG2000 bitstream. This is a very surprising result of course.

3 Security Evaluation

We want to assess the security of the presented selective encryption scheme by conducting a simple ciphertext-only attack. Therefore, an attacker would replace the encrypted parts of the bitstream by artificial data mimicking typical images ("replacement attack", see also [7]). This attack is usually performed by replacing encrypted data by some constant bits (i.e. in selective bitplane encryption). In encrypting the JPEG2000-bitstream, this attack does not have the desired effect, since bitstream values are arithmetically decoded and the corresponding model depends on earlier results and corrupts the subsequently required states. Therefore, the reconstruction result is a noise-like pattern similar as obtained by directly reconstructing the encrypted bitstream. We exploit a built-in error resilience functionality in JJ2000 to simulate a bitstream-based replacement attack. An error resilience segmentation symbol in the codewords at the end of each bit-plane can be inserted. Decoders can use this information to detect and conceal errors. This method is invoked in JJ2000 encoding using the option `-Cseg_symbol on`.

If an error is detected during decoding (which is of course the case if data is encrypted) it means that the bit stream contains some erroneous bits that have led to the decoding of incorrect data. This data affects the whole last decoded bit-plane. Subsequently, the affected data is concealed and no more passes should be decoded for this code-block's bit stream. The concealment resets the state of the decoded data to what it was before the decoding of the affected bit-plane started. Therefore, the encrypted packets are simply ignored during decoding.

Using this technique, we again compare selective JPEG2000 encryption using resolution and layer progressive mode by reconstructing images with a different amount of encrypted packets. Decoding is done using error concealment. In Fig. 7 and 8 we immediately recognize that the PSNR values are significantly higher as compared to directly reconstructed images (see Fig. 3 and 4). Layer progression is more suited for selectively encrypting the angiogram image. For the lena test images, the situation differs slightly: When encrypting only minor parts of the overall bitstream, layer progression is superior, at higher rates of encryption, the resolution progression scheme shows superior results.

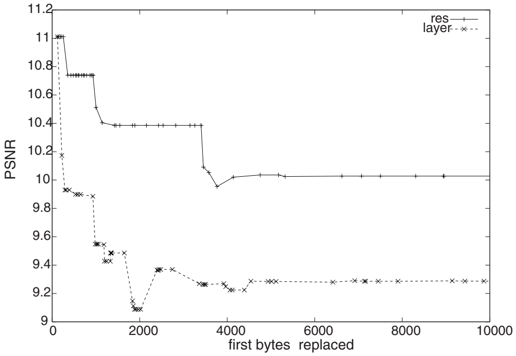

Fig. 7. Angiogram: PSNR of reconstructed images after replacement attack using resolution or layer progressive encoding - part 1.

Again, the numerical values have to be related to visual inspection. Fig. 9.a shows a reconstruction of the selectively compressed angiogram image, where the first 1% of the packets in resolution progressive mode have been encrypted and the reconstruction is done using the error concealment technique. In this case, this leads to a PSNR value of 10.51 dB, whereas the directly reconstructed image has a value of 7.45 dB (see Fig. 5.b). The text in the right corner is clearly readable and even the structure of the blood vessels is exhibited. The lena performs similarly (see Fig. 10.a), all important visual features are reconstructed at 1% encrypted. Here, we have a resulting PSNR of about 11.31 db, whereas the directly reconstructed image has a value of 10.11 dB (see Fig. 6.b).

When increasing the percentage of encrypted packet data steadily, we finally result in 20% percent of the packet data encrypted where neither useful visual nor textual information remains in the image (see Fig. 9.b and 10.b). This result is confirmed also with other images including other angiograms and other still

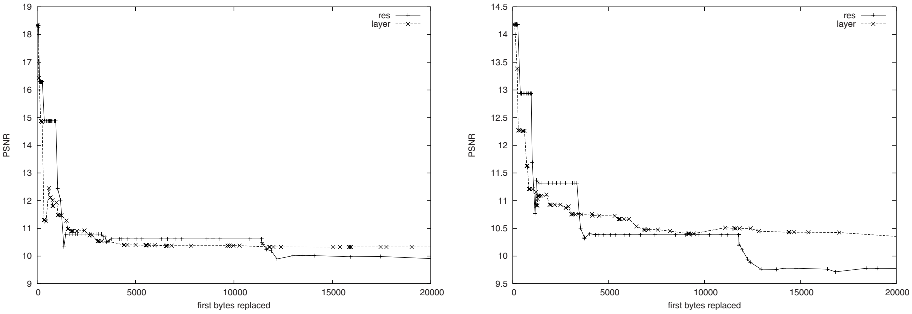

(a) Lena 256 × 256 pixels

(b) Lena 512 × 512 pixels

Fig. 8. PSNR of reconstructed images after replacement attack using resolution or layer progressive encoding - part 2.

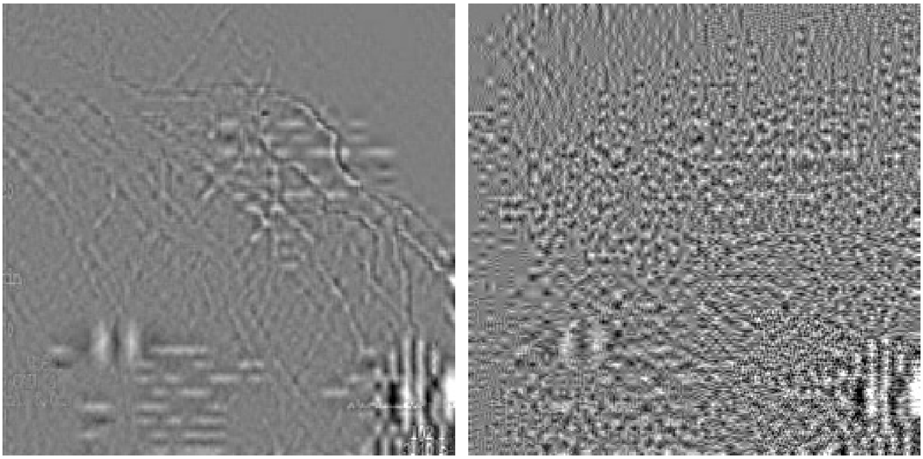

(a) 1% encrypted, 10.51 dB

(b) 20% encrypted, 9.90 dB

Fig. 9. Visual quality of reconstructed Angiogram after replacement attack using resolution encoding.

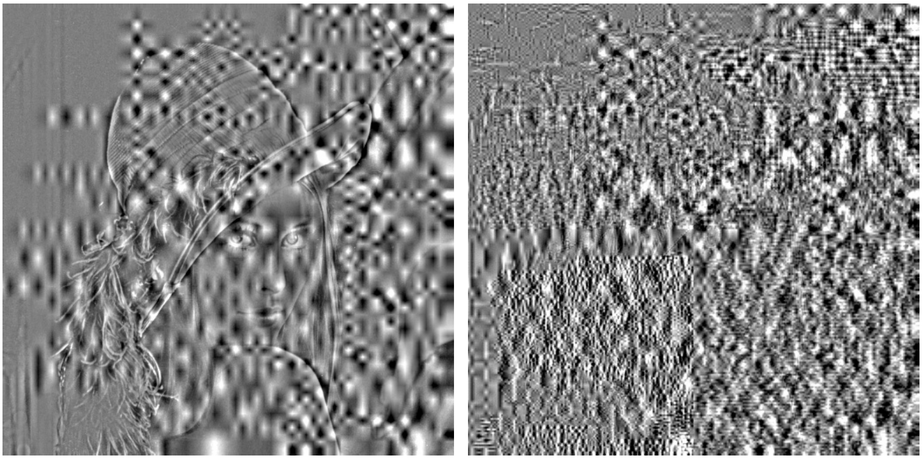

(a) 1% encrypted, 11.31 dB

(b) 20% encrypted, 9.77 dB

Fig. 10. Visual quality of reconstructed Lena (512 pixels) after replacement attack using resolution encoding.

images and can be used as a rule of thumb for a secure use of selective encryption of the JPEG2000 bitstream.

4 Conclusion

A computationally efficient technique for the confidential storage and transmission of digital image data has been discussed. In detail, we propose a partial encryption technique based on AES where parts of the JPEG2000 bitstream are encrypted. The percentage of data subjected to encryption while maintaining high confidentiality is significantly reduced as compared to full encryption, the encryption of 20% data already delivers a satisfying secure result. This is due to the fact that important visual features are concentrated at the begin of the embedded JPEG2000 bitstream and may therefore be protected effectively.

References

1. A. M. Alattar, G. I. Al-Regib, and S. A. Al-Semari. Improved selective encryption techniques for secure transmission of MPEG video bit-streams. In *Proceedings of the 1999 IEEE International Conference on Image Processing (ICIP'99)*. IEEE Signal Processing Society, 1999.
2. H. Cheng and X. Li. Partial encryption of compressed images and videos. *IEEE Transactions on Signal Processing*, 48(8):2439–2451, 2000.

3. J. Daemen and V. Rijmen. *The Design of Rijndael: AES - the advanced encryption standard*. Springer Verlag, 2002.
4. 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*, San Diego, CA, USA, July 2001.
5. Thomas Kunkelmann. Applying encryption to video communication. In *Proceedings of the Multimedia and Security Workshop at ACM Multimedia '98*, pages 41–47, Bristol, England, September 1998.
6. Benoit M. Macq and Jean-Jacques Quisquater. Cryptology for digital TV broadcasting. *Proceedings of the IEEE*, 83(6):944–957, June 1995.
7. M. Podesser, H.-P. Schmidt, and A. Uhl. Selective bitplane encryption for secure transmission of image data in mobile environments. In *CD-ROM Proceedings of the 5th IEEE Nordic Signal Processing Symposium (NORSIG 2002)*, Tromso-Trondheim, Norway, October 2002. IEEE Norway Section. file cr1037.pdf.
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 subband structures for obscured transmission of visual data. In *Proceedings of the 3rd IEEE Benelux Signal Processing Symposium (SPS 2002)*, pages 25–28, Leuven, Belgium, March 2002. IEEE Benelux Signal Processing Chapter.
10. Lintian Qiao and Klara Nahrstedt. Comparison of MPEG encryption algorithms. *International Journal on Computers and Graphics (Special Issue on Data Security in Image Communication and Networks)*, 22(3):437–444, 1998.
11. C. Shi and B. Bhargava. A fast MPEG video encryption algorithm. In *Proceedings of the ACM Multimedia 1998*, pages 81–88, Boston, USA, 1998.
12. Iraj Sodagar, H. J. Lee, P. Hatrack, and Ya-Qin Zhang. Scalable wavelet coding for synthetic/natural hybrid coding. *IEEE Transactions on Circuits and Systems for Video Technology*, 9(2):244–254, 1999.
13. L. Tang. Methods for encrypting and decrypting MPEG video data efficiently. In *Proceedings of the ACM Multimedia 1996*, pages 219–229, Boston, USA, November 1996.
14. D. Taubman. High performance scalable image compression with EBCOT. *IEEE Transactions on Image Processing*, 9(7):1158 – 1170, 2000.
15. D. Taubman and M.W. Marcellin. *JPEG2000 — Image Compression Fundamentals, Standards and Practice*. Kluwer Academic Publishers, 2002.
16. P.N. Topiwala, editor. *Wavelet Image and Video Compression*. Kluwer Academic Publishers Group, Boston, 1998.
17. P. Trunk, B. Gersak, and R. Trobec. Topical cardiac cooling - computer simulation of myocardial temperature changes. *Computers in Biology and Medicine*, 2003. To appear in this issue.
18. T. Uehara, R. Safavi-Naini, and P. Ogunbona. Securing wavelet compression with random permutations. In *Proceedings of the 2000 IEEE Pacific Rim Conference on Multimedia*, pages 332–335, Sydney, December 2000. IEEE Signal Processing Society.
19. Wenjun Zeng and Shawmin Lei. Efficient frequency domain video scrambling for content access control. In *Proceedings of ACM Multimedia 1999*, pages 285–293, Orlando, FL, USA, November 1999.