

Simultaneous Compression and Denoising of Side Scan Sonar Images Using the Discrete Wavelet Transform¹

R. A. Cunha², M. T. Figueiredo³ and C. J. Silvestre²

²Instituto de Sistemas e Robótica, Instituto Superior Técnico, Av. Rovisco Pais, 1049-001 Lisboa, Portugal

³Instituto de Telecomunicações, Instituto Superior Técnico, Av. Rovisco Pais, 1049-001 Lisboa, Portugal

Email: rita@isr.ist.utl.pt, mtf@lx.it.pt, cjs@isr.ist.utl.pt

Abstract- There has been an increasing interest in the use of autonomous underwater vehicles for ocean exploration. Side scan sonars are often installed onboard these vehicles to survey the sea floor. In this context, the need to transmit side scan sonar images over an acoustic low-bandwidth channel calls for the use of compression techniques. This paper proposes a wavelet-based method for compressing side scan sonar images. Wavelets are an adequate choice because their intrinsic properties suit side scan sonar images. In fact, achieved results confirm that wavelets yield high compression rates. Furthermore, wavelet-based compression enhances side scan sonar images by reducing noise without smoothing important details.

I. INTRODUCTION

A. Motivation

Over the last few years, there has been an increasing interest in the use of *autonomous underwater vehicles* (AUVs) to expand the ability to survey ocean areas. See, for example, [1, 2] and the references therein, for an overview of existing prototype autonomous underwater vehicles and their applications. Envisioned AUV missions at sea include environmental monitoring, underwater inspection of estuaries, harbors, and pipelines, and geological and biological surveys of the sea floor.

The side scan sonar is a fundamental component of the AUV surveying equipment. In its normal operating mode, the sonar follows a straight path, towed by the support craft, sending acoustic beams to the left and to the right. The signal propagates through the water until it hits the seafloor or any other obstacle. Some of the energy is reflected back to the

sonar that converts it into electrical energy and stores its intensity and time of arrival. This procedure generates a gray-scale line in which bright sections correspond to strong reflections and dark sections represent lack of return energy. Stacking up these lines forms a side scan sonar image, such as the ones shown in Fig. 1. They are the records of two distinct shipwrecks, Fig.1 (a) and (b), and of a large rock, Fig.1 (c). These images exhibit a considerable and expected amount of noise. Sources of noise range from biological and environmental sources to propeller cavitation [3].

Side scan sonars give scientists the ability to perform large scale geological and biological surveys of the sea floor in search for the most interesting areas of the seabed. During a typical mission, sending sonar images back to the surface ship, rather than simply storing them on the AUV, allows end-users to follow the progress of the survey and redirect the vehicle, if required. This gives rise to the question of efficiently sending data from the sonar, back to the surface ship, over a low bandwidth (typically less than 10 Kbps) acoustic communication channel. In this scenario, compressing side scan sonar images becomes an essential task. The system devised to tackle this problem is standard and consists of a transmission unit installed onboard the AUV and a reception unit onboard the support ship. The AUV unit starts by storing a certain number of sensor lines produced by the side scan sonar. These lines are then fed into the compression algorithm. The result is transmitted over an acoustic channel using a transport layer communication protocol. After receiving the compressed data, the support ship unit reconstructs the image and displays it.

B. Why Wavelets?

The nature of side scan images suggests wavelet-based compression as an appropriate choice. Referring to our description above, these images usually present large uniform

¹ The work of R. A. Cunha and C. Silvestre was partially supported by the MAST Programme of the EU under contract no. MAS3-CT97-0092 (ASIMOV). The first author benefited from a PRAXIS XXI Graduate Scholarship (reference PRAXIS_XXI/BM/20745/99).




Fig. 1. Original side scan sonar images. (a) record of a shipwreck; (b) another shipwreck; (c) rocks.

areas (corresponding to the seabed) which is occasionally disrupted by rocks or other objects such as shipwrecks or pipelines (see Fig. 1). In addition, side scan images are inherently noisy. Given their multiresolution structure and ability to locally represent high-frequency variations [4], wavelets deal with sonar images far better than other compression techniques. For example, JPEG, the widespread Fourier transform-based compression scheme, proved to be quite inadequate when compared with wavelets. Figures 2 and 3 illustrate the difference between the two. The original side scan sonar image in Fig. 1 (a) was compressed using JPEG (Fig. 2) and using wavelets (Fig. 3). Although compression ratios are very similar, image denoising improves considerably with the wavelet-based compression.

In this paper we present preliminary results concerning the use of wavelet-based methods to compress side scan sonar images; an important byproduct of compression is a noticeable enhancement of image quality (less noise without loss of important details).




Fig. 2. JPEG compression (compression ratio: 16.4).



Fig. 3. Wavelet-based compression (Daubechies-2 wavelet, level 6, threshold 40, compression ratio: 15.9).

II. PROPOSED SCHEME

The adopted compression scheme includes the standard steps of transform coding: transform application followed by coefficient thresholding and encoding.

A. Wavelet Decomposition

Wavelets underlie many recent advances in key areas of signal and image processing; namely, approximation (or representation), estimation, and compression (for example, see Mallat's recent book [4] and the many references therein). In these applications, two important properties of the discrete wavelet transform (DWT) of real-world signals and images are exploited: **(a)** it is *sparse*, meaning that a few large coefficients dominate the representation, and **(b)** the coefficients tend to be much less correlated than the original data. These properties, together with the existence of fast implementations, make the DWT an excellent tool for many signal/image processing tasks [4].

The basic approach to DWT-based signal/image processing consists in manipulating the DWT coefficients, rather than the signal samples themselves. The DWT decomposes a signal (in this case, an image) into a linear combination of shifted and scaled versions of a so-called "mother-wavelet" [4]. Consequently, the DWT has two adjustable attributes: wavelet decomposition base and number of levels of decomposition. Testing more than one decomposition base is a relevant issue because the wavelet shape plays a significant role on DWT's ability to "approximate" original data. In this application, since we are working with noisy images, the wavelet should fit the ideal (unknown) noise-free image.

We considered four wavelets, from the vast available set, to be tested as decomposition bases: Haar, Daubechies-2, Biorthogonal-1.3 and Coiflets-1 [4]. The Haar wavelet is the simplest one and therefore was chosen to set the reference - if another basis brings negligible improvement, it should be discarded since it also brings added complexity. The other three were chosen because they are the simplest representatives of three important families: Daubechies wavelets, biorthogonal wavelets and Coiflets [4].

Wavelet decomposition produces a multiscale/multilevel hierarchy. Each level of decomposition has two sets of coefficients: the approximation and the details (residuals of the approximation). These are built by decomposing the approximation of the previous level, which, in the case of level one, is the image itself. As the level is increased, the scale or resolution of the approximation decreases. This suggests that a high number of decomposition levels may improve both compression and denoising.

B. Coefficient Thresholding and Encoding

The DWT of most real-world signals and images tend to be dominated by a few large coefficients (see [5, 6]). This is the so-called *sparseness* property which, in probabilistic terms, corresponds to a wavelet coefficient density function with a marked peak at zero and heavy tails; that is, a strongly super-Gaussian density. A natural denoising/compression criterion then results from this statistical characterization of DWT coefficients: if the magnitude of an observed wavelet coefficient is large, its signal component is probably much larger than the noise and it should be kept; conversely, if a coefficient has small absolute value, it is probably due to

noise and it should be attenuated or even removed. This (together with the decorrelation property that suggests processing coefficients independently of each other) is the rationale underlying the now classical thresholding methods pioneered by Donoho and Johnstone [5] and all their descendants (see [4, 5, 6, 7] and references therein).

In the case of side scan sonar images, by simply inspecting the reconstruction of the approximation and details, we decided that the level one details could be discarded. Omitting these details improves the compression ratio considerably, since they make up approximately three-quarters of the total number of coefficients. To reduce higher level details, we applied three different types of thresholding nonlinearities, expressed by the following equations

$$y = \begin{cases} 0, & |x| < a \\ x, & |x| \geq a \end{cases}, \quad (1)$$

$$y = \text{sign}(x) / (|x| - a)_+, \quad (2)$$

$$y = (x^2 - a^2)_+ \checkmark x, \quad (3)$$

where a is a threshold value and $(\)_+$ denotes the “positive part”, defined as $y_+ = y$, if $y \geq 0$, and $y_+ = 0$, if $y < 0$. The nonlinearities described in (1) and (2) are called *hard* and *soft* thresholding, respectively [4, 5]. The difference between them is clear. In (1) there are discontinuities at a and $-a$, which make this rule somewhat unstable [7]. This does not occur in (2), but instead coefficients above the threshold do not keep their values. The thresholding expressed in (3) tries to keep the advantages and dispose of the disadvantages of both the other two. For details of its derivation, refer to [6, 7].

We chose a very simple implementation for the coding step. The coefficients are rounded to integers and the zero-valued coefficients are run-length encoded. Naturally, there may exist other more efficient alternatives; this simple method sets the lower bound for other compression strategies.

III. EXPERIMENTAL RESULTS

The previous section describes a scheme with several degrees of freedom. These consist of wavelet family, number of decomposition levels, nonlinearity, and threshold. In this section, we describe experimental results where we try to establish choices leading to high compression ratios and noise removal performance, without reducing image information.

A. Wavelet Shape

The compression ratios obtained with the four wavelets referred in II.A., at different decomposition levels, are summarized in Fig. 4. The Daubechies-2 and the Coiflets-1 wavelets generate very similar results. The Haar consistently produces better compression ratios, while the Biorthogonal-1.3 line is always below the others. On the other hand, quality of the compressed images is slightly worse for the Haar wavelet, but fairly equivalent for the others. This confirms that wavelet shape influences the results. In this case, Daubechies-2 and Coiflets-1 are the most adequate choices.

B. Number of Levels of Decomposition

Returning to Fig. 4, it is possible to draw conclusions about another parameter. It is clear that the number of levels of decomposition has a great influence on compression.




Fig. 4. Comparison between compression ratios acquired with different wavelets (threshold 30).

Increasing from two to six levels of decomposition increases compression ratios from values under 9 to values above 17. Moreover, the effect on image definition is even more significant. Figures 5 and 6 show two images generated by applying Daubechies-2 wavelet with two and six levels of decomposition, respectively. Noise reduction in Fig. 5 is already considerable, but in Fig. 6 it is impressive. Important elements such as shadows and high reflection zones (produced by objects) remain well defined. Even relatively smooth relief changes are still present in the compressed image.

Comparing Fig. 6 with both Fig. 5 and Fig. 1 (b), it is clear that no relevant information was lost and only noise was removed. However, there seems to be a limit for the number of levels of decomposition. Decomposing more than six times does not bring much benefit in terms of compression, while image definition starts to suffer. It is likely that up to six levels, the decrease in resolution only “filters” out noise, while over that level it starts to degrade the image.




Fig. 5. Daubechies-2 wavelet decomposition at level 2, threshold 40 (compression ratio: 8.7).



Fig. 6. Daubechies-2 wavelet decomposition at level 6, threshold 40 (compression ratio: 17.2).

C. Thresholding

We still have not mentioned thresholds in this section but it is implicit that some kind of thresholding was applied to generate the described results. For instance, Figs. 5 and 6 are the outcome of thresholding DWT coefficients using the nonlinearity in (3) with $a=40$. Although image quality is highly dictated by wavelet decomposition, thresholding can be tuned to produce the highest compression without causing loss of information. In Fig. 7 we plot compression ratios for threshold values between 10 and 70. Table I lists the corresponding bit rates. Naturally, compression increases with the threshold value and we have reached compression ratios as high as 46 (threshold 70, decomposition level 6). However, at this level of compression, image quality starts to suffer as the finest details are lost (see Fig. 8).

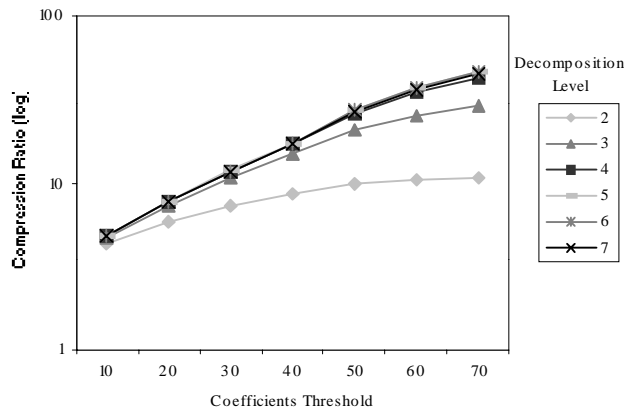


Fig. 7. Compression ratios (logarithmic scale) for different threshold values.

TABLE I
BIT RATES FOR DIFFERENT DWT LEVELS AND THRESHOLDS (bits/pixel).

DWT level	Coefficients threshold					
	20	30	40	50	60	70
2	1.37	1.11	0.92	0.81	0.77	0.75
3	1.09	0.75	0.54	0.38	0.32	0.28
4	1.04	0.69	0.47	0.31	0.23	0.19
5	1.03	0.68	0.46	0.29	0.22	0.17
6	1.04	0.68	0.47	0.30	0.22	0.17
7	1.04	0.68	0.47	0.30	0.22	0.18

Finally, we analyzed the difference between the three types of thresholding nonlinearity referred in II.B. The outcome shown in Fig. 9 was as expected – hard-thresholding generates small undesirable discontinuities (in Fig. 9 (a)), soft-thresholding “smooths” the image too much (see Fig. 9 (b)), while the function in equation (3) produces the best result (shown in Fig. 9 (c)).

These threshold dependent results propose a basis for building a configurable application adequate for AUV missions. For example, as the mission progresses, the threshold value could be adapted to meet the changing requirements in image accuracy. This implies managing the trade-off between compression ratio (that must be kept high due to transmission bandwidth limitations) and image quality. A 70-valued threshold generates high compression ratios (above 40) and images have enough definition to allow a large-scale survey (see Fig. 8). Whenever a target is identified, more detailed information can be obtained by lowering the threshold value. Future work may focus on determining adequate criteria for tuning the process, based on compression and accuracy specifications.




Fig. 8. Daubechies-2 wavelet decomposition at level 6, threshold 70 (compression ratio: 46.0).




Fig. 9. Comparison between types of thresholding. (a) hard-thresholding; (b) soft-thresholding; (c) thresholding in equation (3); (Daubechies-2 wavelet decomposition at level 6, threshold 40, compression ratios: 19.9, 20.7, 20.2).

IV. CONCLUSIONS

This paper proposes a new method for compression of underwater side scan sonar images. The problem is motivated by the need to transmit these images over an acoustic link from autonomous underwater vehicles.

Analyzing the results, we concluded that wavelets are highly suited to compress side scan sonar images. Transmission rates below 0.5 bit/pixel (compression ratios above 15) are achieved using the simplest encoding. In addition, the quality of the images is considerably improved through the “detail-preserving denoising” inherent to wavelet-based compression.

The merit of this scheme lies in the use of wavelet decomposition which has the ability to rearrange information such that both compression and denoising are highly effective.

ACKNOWLEDGMENT

We would like to thank Mr. Marcus Cardew from System Technologies-UK for providing the sonar images used to develop this work.

REFERENCES

- [1] A. Pascoal, P. Oliveira, C. Silvestre, A. Bjerrum, J.-P. Pignon, G. Ayela, S. Bruun, and C. Petzelt. “MARIUS, An autonomous underwater vehicle for coastal oceanography”, *IEEE Robotics and Automation Magazine*, pp. 46-59, December 1997.
- [2] K. Valavanis, G. Saridis, A. Pascoal, P. Lima, and F-L. Pereira, editors, *Proceedings of the Joint U.S./Portugal Workshop on Undersea Robotics and Intelligent Control*, Lisboa, Portugal, March 1995.
- [3] R. Urick, *Principles of Underwater Sound*, Peninsula Pub, 1996.
- [4] S. Mallat, *A wavelet tour of signal processing*, Academic Press, 1998.
- [5] D. Donoho and I. Johnstone, “Ideal adaptation via wavelet shrinkage,” *Biometrika*, vol. 81, pp. 425-455, 1994.
- [6] M. Figueiredo and R. Nowak, “Bayesian wavelet-based image estimation using non-informative priors”, in *Mathematical Modeling, Bayesian Estimation, and Inverse Problems*, SPIE vol. 3816, pp. 97-108, 1999.
- [7] H. Gao, “Wavelet shrinkage denoising using the non-negative garrote,” in *Journal of Computational and Graphical Statistics*, vol. 7, pp. 469-488, 1998.