WAVELET-BASED CODING OF THREE-DIMENSIONAL OCEANOGRAPHIC IMAGES AROUND LAND MASSES

James E. Fowler

Dept. of Electrical & Computer Engineering
NSF Engineering Research Center
Mississippi State University, Starkville, MS

Daniel N. Fox

Naval Research Laboratory
Stennis Space Center, MS

ABSTRACT

We describe an algorithm for the embedded coding of 3D oceanographic images. These images differ from those arising in other applications in that valid data exists only at grid points corresponding to sea; grids points that cover land or lie beyond the bathymetry have no associated data. For these images, we employ a 3D lifting wavelet transform tailored specifically to the potentially sparse nature of the data by processing only the valid sea data points in between land masses. In addition, we introduce successive-approximation runlength (SARL) coding, an embedded-coding procedure which adds successive-approximation properties to the well known stack-run (SR) algorithm. SARL is a general technique applicable to the oceanographic images considered here as well as to other coding tasks in which embedded coding is desired but for which zerotree-techniques are impractical.

1. INTRODUCTION

The US Naval Oceanographic Office (NAVOCEANO) produces a variety of oceanographic datasets which are distributed to users worldwide, many of whom possess only very low-bandwidth communications links (e.g., submarines capable of only low-frequency reception). A typical organization of these datasets is that of a scalar field defined on a 3D rectilinear grid covering a certain rectangular region of sea and land at a certain number of standard ocean depths. As these datasets can be sizeable, compression is a must. In addition, the capability of decoding an initial portion of a longer message would be helpful for low-bandwidth users in the case that the transmission channel is interrupted or prematurely terminated—such needs call for efficient embedded coding.

In recent years, a number of compression algorithms have been proposed for 3D imagery arising in applications such as multispectral/hyperspectral remote sensing, volumetric medical imaging, and video coding. Typically, such techniques extend well known 2D image-coding techniques to the third dimension, either by coupling purely 2D image coding on individual image planes, or "slices," with decorrelation across the third dimension (e.g., motion compensation, with MPEG being the prime example), or by employing fully 3D mechanisms such as 3D transforms (e.g., 3D DCT [1] or 3D wavelet transforms [2]) and 3D zerotree coding ([2, 3]). Initially it may appear that any of these 3D image-coding methods would be suitable, with perhaps minor adaptations, to oceanographic data. However, not all of these techniques are embedded. More importantly though, the 3D oceanographic imagery considered here differs significantly from the imagery arising in multispectral, hyperspectral, medical, or video applications—scalar-field values (i.e., "pixels"), exist only for 3D grid points that correspond to sea, while grid points that cover land

(at the ocean surface), or that lie beyond the bathymetry (at depths below the ocean surface), have no associated data. Although it is a possible that a given dataset has no "land points" (i.e., it covers an area that is deep open ocean), the more general case includes arbitrarily shaped coastlines, islands, and ocean floor which occupy an arbitrarily large portion of the overall 3D rectilinear grid. A successful 3D-oceanographic-image compression algorithm must efficiently code valid data values while skirting around these land masses, regardless of how data values are grouped or how sparsely they occur in the dataset.

To this end, we propose the Wavelets Around Land Masses (WAVAL) algorithm for the coding of 3D oceanographic images. Our WAVAL algorithm employs a separable 3D biorthogonal wavelet transform that is carefully calculated over only sea regions in the dataset. In addition, successive-approximation runlength coding (SARL), a new simple embedded-coding technique, is introduced to provide efficient coding of the 3D array of wavelet coefficients. In the remainder of this paper, we describe the WAVAL technique, including details behind the wavelet transform and the SARL algorithm. We also overview experimental results comparing performance of WAVAL to the technique currently in use by the Navy. We note that, although we will be primarily concerned with oceanographic images consisting of water-temperature values, our WAVAL coder can be applied directly to other scalar quantities, such as salinity, sound speed, vorticity, etc., as well.

2. THE WAVAL ALGORITHM

The WAVAL encoder algorithm consists of the following steps: 1) extract and code a binary land-sea mask indicating where, on the 3D rectilinear grid, valid data values are located; 2) perform a 3D wavelet transform over points identified as sea by the land-sea mask; 3) use SARL to quantize and entropy-code wavelet coefficients in an embedded fashion; and 4) pack the embedded bitstream into a standard ASCII message format for distribution.

2.1. Land-Sea Mask

The land-sea mask is a binary mask that differentiates points in the 3D rectilinear grid corresponding to sea (where valid data values reside) from those corresponding to land. Since the land-sea mask is critical to maintaining synchronization between the WAVAL encoder and decoder, an encoding of the land-sea mask is the initial portion of the bitstream output by the WAVAL encoder. Since the entire land-sea mask must be received before decoding begins, this mask encoding is not embedded.

Encoding of the binary land-sea mask consists of runlength coding of similar consecutive values encountered in a raster-scan traversal of the dataset. An initial bit indicates whether the traversal starts with land or sea. Afterwards, each run of consecutive

land or sea values is represented by the length of the run. The encoding of runlengths for the land-sea mask is nearly identical to that used in the SARL algorithm (see below and Fig. 2). The resulting symbol stream is entropy coded using an arithmetic coder [4] operating on the three-symbol alphabet $\{0, 1, +\}$.

Raster-scan coding of the land-sea mask proceeds so that all mask values at a certain depth are coded before the next deeper values (i.e., depth forms the outer loop of the scan). The oceanographic bathymetry is nondecreasing as depth increases; that is, the set of land points at a certain depth level is a subset of the land points at the next deeper level. Consequently, in the coding of the land-sea mask, those points labeled as land at a given depth need not be coded again at deeper levels. The raster-scan and runlength coding of the land-sea mask thus skips over previously labeled land points as the scan proceeds deeper in the 3D volume.

2.2. 3D Wavelet Transform

The WAVAL coder uses a lifting implementation of the Cohen-Daubechies-Feauveau $(2, 2)$ biorthogonal wavelet [5] (i.e., "linear lifting" [6]), appropriately adapted to transform data values in-between land masses. Below we describe this transform operation for a single dimension. To construct the separable 3D transform, the 1D transform is applied independently in each of the three directions (row, column, depth), producing 8 subbands; further decompositions in this fashion are carried out recursively on the baseband subband.

In the usual implementation (e.g., [6]), of lifting, 1) a lazy wavelet transform (LWT) separates even and odd samples of a baseband signal, 2) odd samples are predicted from even samples with the difference becoming the highpass band of the next scale, and 3) even samples are "updated" from highpass coefficients to produce the lowpass band of the next scale. This same procedure is used in the wavelet transform of the WAVAL coder, except that, rather than transforming an entire row, column, or depth of pixels, only 1D *segments* of consecutive sea data values in between land masses are transformed (see Fig. 1). Each sea segment is transformed individually into lowpass and highpass *subband segments*; subband segments resulting from all the baseband segments in a row, column, or depth are interwoven with subsampled land points as necessary to create the lowpass and highpass bands (see Fig. 1) for the row, column, or depth. The difference between transforming segments and transforming an entire row, column, or depth lies in that segments may start with an odd or an even sample, whereas the usual convention in lifting assumes starting with an even sample. The WAVAL coder keeps track of the starting location of each segment to assure correct subsampling and placement into lowpass and highpass bands as illustrated in Fig. 1. Symmetric extension is used at either end of each sea segment when lifting and updating operations require data values beyond the boundary of the segment (i.e., over a land mass or beyond the bounds of the signal itself); this data extension must be tailored specifically to the cases in which a segment starts/ends on an even/odd sample. Finally, we note that, since the length of a baseband segment may be arbitrarily odd or even, the corresponding lowpass and highpass subband segments are not necessarily the same length and may even be of zero length.[1]

2.3. Successive-Approximation Runlength Coding (SARL)

The embedded zerotree wavelet (EZW) algorithm [7], designed originally for 2D images, is an easily implemented, computation-

---

[1]In this latter case, a baseband segment consists of only one data value, and the lifting decomposition degenerates to a multiplication or division by $\sqrt{2}$ to maintain unitary scaling.

ally efficient, embedded technique with effective coding performance. It and its variants have been extended to 3D datasets (e.g., [2]). However, the coding efficiency of EZW relies on the representing of large regions of zeros with single zerotree symbols. The potentially sparse nature of the oceanographic datasets considered in our application hinders zerotree performance and significantly complicates implementation. On the other hand, stack-run (SR) coding [8] achieves 2D-image performance somewhat superior to EZW without relying on zerotree structure. Because of its lack of zerotrees, SR is suitable for use in applications more general than 2D-image coding; however, SR is not embedded. The SARL algorithm employs successive-approximation embedded coding in the form of EZW's "bit-plane" coding while replacing EZW's reliance on zerotrees with an efficient runlength-coding scheme similar to that of SR. Although SARL was developed for the 3D oceanographic coding considered here, we anticipate that it is equally applicable to other similar coding applications (3D or otherwise) for which embedded coding is desired but, for whatever reason, cannot make use of zerotrees.

The SARL encoding algorithm operates as follows. Each coefficient is compared to a threshold. If a coefficient is greater than or equal to the threshold, it is a *significant* coefficient, otherwise it is *insignificant*. The threshold is successively decreased by dividing by 2 as coding progresses; i.e., SARL employs a "bit-plane" successive approximation. Within the coding of each bit plane, the positions of coefficients that become significant with respect to the threshold for the first time are indicated to the decoder by coding the distances between them as runlengths; the signs of these newly significant coefficients are included in the runlength-encoding scheme (see Fig. 2). After all runlengths are coded, "refinement" symbols are output for all coefficients previously determined to be significant during coding of prior bit planes. Then the threshold is halved, and the algorithm repeats for the next bit plane. Pseudocode for the algorithm follows (assume N wavelet coefficients, $x(n)$, with $mask(n)$ denoting the land-sea mask):

$max\_bits = \lfloor \log_2 (\max_n |x(n)|) \rfloor$
$threshold = \text{pow} (2, max\_bits)$
$significance\_map(n) = 0, \forall n$
while (TRUE) do
  $runlength = 0$
  $difference\_map(n) = 0, \forall n$
  $AC\_context = \textbf{RUNLENGTH}$
  for $n = 1$ to N do
    if $mask(n) = \textbf{SEA}$
      if $significance\_map(n) = 0$
        if $|x(n)| \geq threshold$
          $significance\_map(n) = 1$
          $difference\_map(n) = 1$
          $\texttt{output\_run} (runlength, \text{sign of } x(n))$
          $x(n) = |x(n)|$
          $runlength = 0$
        else
          $runlength = runlength + 1$
  done
  $\texttt{output\_run} (runlength, 1)$
  $AC\_context = \textbf{REFINEMENT}$
  for $n = 1$ to N do
    if $mask(n) = \textbf{SEA}$
      if $significance\_map(n) = 1$
        if $x(n) \geq threshold$
          if $difference\_map(n) = 0$

$$\text{output\_symbol ("1")}$$
$$x(n) = x(n) - threshold$$
else
  if $difference\_map(n) = 0$
    output_symbol ("0")
  done
  $threshold = threshold/2$
done

The output_run subroutine outputs symbols for integer runlengths in a manner similar to the run encodings employed by SR [8]; the SARL runlength-encoding scheme is illustrated in Fig. 2. The stream of symbols resulting from the runlength encodings and the refinement symbols is entropy coded using arithmetic coding [4] with two contexts; a four-symbol alphabet, $\{0, 1, +, -\}$, is used for the **RUNLENGTH** context while a two-symbol alphabet, $\{0, 1\}$, is used in the **REFINEMENT** context.

2.4. Channel Packing

The channel employed by the Navy for distribution of oceanographic data employs a specific message format defined on a restricted ASCII-based alphabet. The encoded data is transmitted via a number of message lines, each of which is limited to 68 characters plus a carriage-return. To this message body, twelve lines of fixed header/trailer information is attached. A set of 40 ASCII characters constitute the valid symbols for the message body. The WAVAL coder packs bits emerging from the arithmetic coder into 5-bit characters which are in turn mapped into the 40-symbol alphabet; thus, only 32 of the 40 permitted symbols are used by the encoder.

3. EXPERIMENTAL RESULTS

In this section, we compare the performance of the WAVAL coder to that of the technique currently employed by the Navy for the transmission of 3D oceanographic images. The current Navy standard is known as the Empirical Orthogonal Function (EOF) technique and is, in essence, a Karhunen-Loève transform (KLT) followed by uniform scalar quantization and a form of runlength encoding. In the EOF scheme, a separable 2D KLT is applied to image slices independently; no decorrelation is performed along the third dimension. Points corresponding to land in the dataset are "filled" with a special extrapolation based on the sea data values prior to taking the KLT, and the entire resulting volume is subject to quantization and runlength coding. We note that the EOF technique is not embedded.

Figs. 3 through 5 show distortion versus rate for three ocean-temperature datasets. Rate is expressed as the total number of message lines produced, included the fixed 12-line header/trailer. The distortion measure is the maximum absolute error between the original and reconstruction over the sea values of the dataset. The hawai dataset (around Hawaii) features simple land-sea mask information as most of the dataset is sea; the adrtc dataset (Adriatic Sea) is of moderate land-sea complexity; and the ylsoj dataset (Yellow Sea, Sea of Japan) has a relatively complex land-sea mask. Table 1 presents the sizes of the datasets and the composition of the bitstream produced by the WAVAL coder.

4. CONCLUSIONS

Figs. 3 through 5 show that WAVAL routinely outperforms the EOF algorithm, sometimes by a gain of an order-of-magnitude or more in distortion performance. The only case in which EOF outperforms WAVAL is when very short (< 175 lines) messages are sent for the ylsoj dataset. This dataset comprises the complex
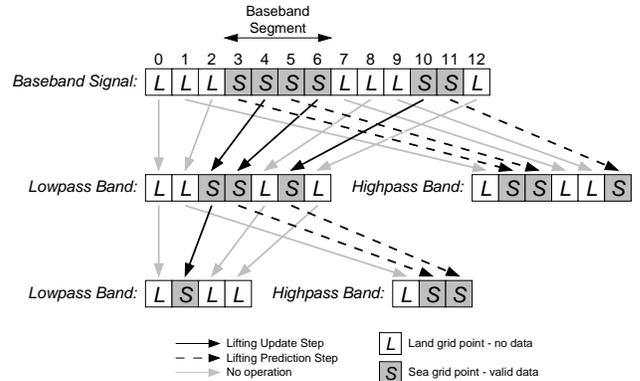
Figure 1: An example of two stages of 1D wavelet decomposition. Each baseband segment of consecutive sea values is decomposed into a lowpass and a highpass subband segment; subband segments are assembled into lowpass and highpass bands maintaining appropriate subsampling pattern and positioning within the band.

Runlength	0	1	2	3	4	5	6	7	8	...
Encoding	$\pm$	0	1	0	1	0	1	0	1	...
		$\pm$	$\pm$	0	0	1	1	0	0	...
				$\pm$	$\pm$	$\pm$	$\pm$	0	0	...
								$\pm$	$\pm$	...

Figure 2: Symbol encodings of runlengths used in the coding of the land-sea mask and within the SARL algorithm (symbols in order first to last listed from top down). For the SARL algorithm, $\pm$ is either $+$ or $-$ depending on the sign of the subsequent wavelet coefficient; for the land-sea mask, $\pm$ is always $+$.

coastline of the Sea of Japan area. The land-sea mask for this dataset (nonembeddedly coded by WAVAL at 113 lines) occupies a majority of the lines of short messages with very few, if any, lines remaining to refine wavelet coefficients, resulting in poor distortion performance. We note that, for such short encodings of the ylsoj dataset, superior EOF performance is not significant since the distortion achieved by both algorithms is too large to be useful in the Navy's oceanographic applications. In the region of acceptable distortion performance (on the order of $0.1°C$), WAVAL consistently outperforms EOF on all datasets tested.

In this paper, we have described the WAVAL algorithm for the embedded coding of 3D oceanographic images. To this end, we have introduced the SARL embedded-coding procedure which adds EZW-like successive-approximation properties to SR. We employ a 3D wavelet transform based on lifting tailored specifically to the potentially sparse nature of oceanographic image data. Our transform processes only the valid sea data points in between land masses—a task difficult to achieve with traditional block-based transforms such as the DCT. Although SARL has been developed for the specific application considered here, it is a general technique applicable to other coding tasks, particularly those in which embedded coding is desired but for which zerotree-techniques are impractical (such as our application). As an example, for the 2D coding of the grayscale Lenna image at 0.5bpp, SARL produced a PSNR on the order of only 0.7dB below that of SR (both algorithms using the $(2, 2)$ biorthogonal wavelet used here); however, the perceptual-quality performance of both algorithms was identical. It appears that there is little practical performance cost in adding embedded coding to SR in the form of SARL.

	Original Dataset			WAVAL Output	
Name	Size	Percent Land	Header	Land-Sea Mask	Wavelet Coeffs.
hawai	$46 \times 41 \times 34$	1.3%	0.1%	0.6%	99.3%
adrtc	$81 \times 61 \times 34$	87.6%	0.1%	3.2%	96.7%
ylsoj	$161 \times 101 \times 34$	44.2%	0.1%	20.7%	79.2%

Table 1: Original dataset sizes and composition of WAVAL output for a message size of 500 lines.
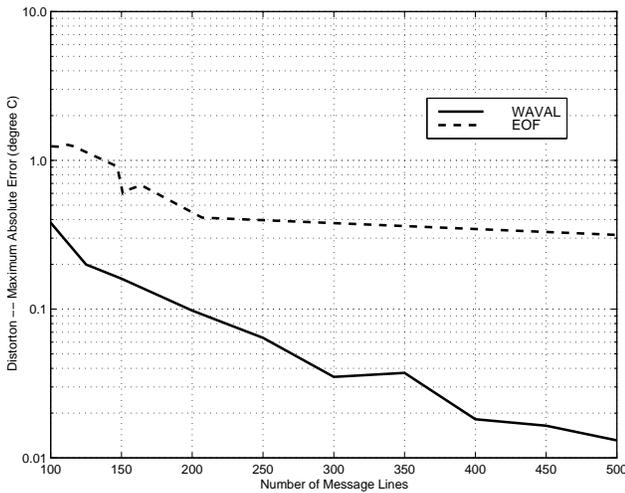
Figure 3: Rate-distortion performance for the `hawai` dataset.
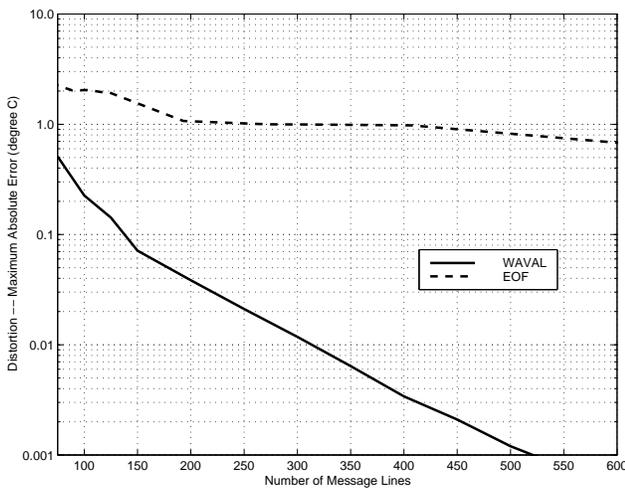
Figure 4: Rate-distortion performance for the `adrtc` dataset.
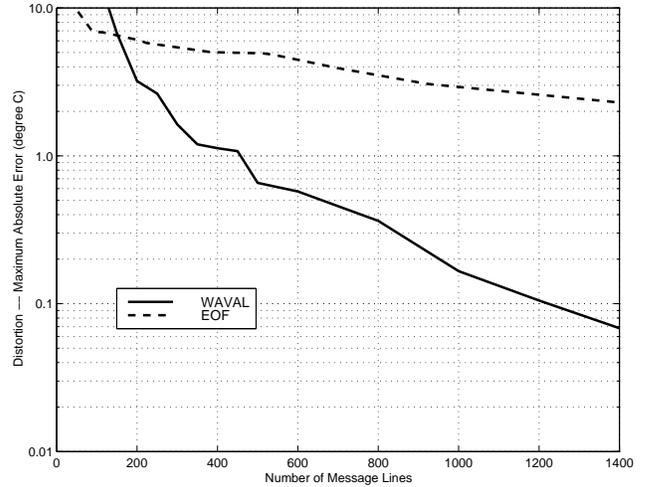
Figure 5: Rate-distortion performance for the `ylsoj` dataset.

5. REFERENCES

[1] K. K. Chan, C. C. Lau, S. L. Lou, A. Hayrapetian, B. K. T. Ho, and H. K. Huang, "Three-dimensional Transform Compression of Images from Dynamic Studies," in *Medical Imaging IV*, Y. Kim, Ed. Proc. SPIE 1232, 1990, pp. 322–326.

[2] B.-J. Kim and W. A. Pearlman, "An Embedded Wavelet Video Coder Using Three-Dimensional Set Partitioning in Hierarchical Trees (SPIHT)," in *Proc. DCC*, Mar. 1997, pp. 251–257.

[3] F. Amato, C. Galdi, and G. Poggi, "Embedded Zerotree Wavelet Coding of Multispectral Images," in *Proc. ICIP*, Oct. 1997, pp. 612–615.

[4] I. H. Witten, R. M. Neal, and J. G. Cleary, "Arithmetic Coding for Data Compression," *Comm. ACM*, vol. 30, no. 6, pp. 520–540, Jun. 1987.

[5] A. Cohen, I. Daubechies, and J.-C. Feauveau, "Biorthogonal Bases of Compactly Supported Wavelets," *Comm. Pure Appl. Math.*, vol. 45, no. 5, pp. 485–560, May 1992.

[6] I. Daubechies and W. Sweldens, "Factoring Wavelet Transforms into Lifting Steps," *Jour. Four. Anal. Appl.*, vol. 4, no. 3, pp. 245–267, 1998.

[7] J. M. Shapiro, "Embedded Image Coding Using Zerotrees of Wavelet Coefficients," *IEEE Trans. Signal Proc.*, vol. 41, no. 12, pp. 3445–3462, Dec. 1993.

[8] M.-J. Tsai, J. D. Villasenor, and F. Chen, "Stack-Run Image Coding," *IEEE Trans. Circ. Sys. Video Tech.*, vol. 6, no. 5, pp. 519–521, Oct. 1996.