Embedded Wavelet-Based Coding of Three-Dimensional Oceanographic Images With Land Masses

James E. Fowler, *Member, IEEE*, and Daniel N. Fox

Abstract— **We describe the Wavelets Around Land Masses (WAVAL) system 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, the WAVAL system employs 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. 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 employed to code wavelet coefficients resulting from the 3D transform in the WAVAL system; however, it is a general technique applicable to other coding tasks in which embedded coding is desired but for which zerotree-techniques are impractical. Experimental results show that the WAVAL system achieves substantial improvement in rate-distortion performance over the technique currently used by the US Navy for compression of oceanographic imagery.**

I. 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 communication 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. This grid describes a rectangular region of sea and land at the ocean surface and extends downwards at a certain number of standard ocean depths. As these oceanographic datasets can be sizeable, efficient compression is a must for low-bandwidth transmission. However, since low-frequency communications entail long transmission times, it is often the case that the transmission channel is interrupted or the communication is otherwise prematurely terminated. Thus, the capability of decoding an initial portion of a longer message would be extremely helpful for many users.

One of the recent advances in the maturing field of data compression is ideally suited to the task at hand—*embedded coding*. An embedded-coding system produces compressed representations of an original dataset at a variety of amounts, or rates, of compression, where each of these representations are included ("embedded") at the beginning of a single output bitstream. That is, any prefix of an embedded bitstream can be decoded to produce a valid representation of the original dataset; the longer the prefix, the more accurate is the representation. Although it was initially thought that the imposition of such an embedded-structure constraint would result in inefficient coding performance when compared to traditional nonembedded coders [1], recent wavelet-based coders (e.g., [1, 2]) for 2D images have been shown to produce embedded bitstreams at efficiencies that rival or surpass performance of traditional nonembedded image-compression algorithms.

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. Some, but not all, of these proposals involve embedded algorithms. Typically, 3D coding techniques extend well known 2D image-coding techniques to the third dimension, either by 1) coupling purely 2D image coding on individual image planes, or "slices," with decorrelation across the third dimension, or by 2) employing fully 3D mechanisms such as 3D transforms. In the first category of algorithms, the prime example is the MPEG [3] coding standards for motion video which couple a 2D block-based discrete cosine transform (DCT) with motion compensation for decorrelation in time. Other proposals following this same fundamental approach include multispectral/hyperspectral coders using various forms of linear/nonlinear prediction for decorrelation along the third dimension and 2D slice-based processing based on a DCT [4] or vector quantization [5].

More recent algorithms typically fall in the second category with many of these techniques employing 3D transforms. Many of these coders, particularly those designed for multispectral imagery, recognize that often the extent of data in one of the dimensions (e.g., the spectral dimension for multispectral images) is significantly shorter than that in the other two directions. In these applications, the 3D transform typically employs the statistically optimal yet data-dependent Karhunen-Loève transform (KLT) for decorrelation and energy compaction along the shorter direction while 2D DCT [6] or wavelet transforms [7–9] are used in the other two dimensions. In other applications wherein the data has significant size in all three dimensions, the KLT is usually too computationally complex and 3D DCT [4, 10] or 3D wavelet transforms [11, 12] are used. In either case, a variety of quantization and entropy-coding strategies deriving from 2D image-coding algorithms are extended to 3D and employed following the 3D transform. Examples include uniform scalar quantization/runlength Huffman coding [7, 11], trellis-coded quantization [4], and zerotree coding [8, 9, 12].

Initially it may appear that any of the above 3D image-coding methods would be suitable, with perhaps minor adaptations, to

Portions of the work described here were presented at the 2000 IEEE International Conference on Image Processing.

J. E. Fowler is with the Department of Electrical & Computer Engineering and the NSF Engineering Research Center, Mississippi State University, Starkville, MS 39762 USA.

D. N. Fox is with the Naval Research Laboratory, Stennis Space Center, MS 39539 USA.

This work was funded in part under PE63207N by the Oceanographer of the Navy via PMW185 of the Space and Naval Warfare System Command.

the present application of coding of oceanographic data. However, only a few [8, 9, 12] of these techniques produce embedded bitstreams. Additionally, as it is pointed out in [9], implementation problems may arise when extending 2D algorithms to sizeable 3D data (as in [8, 9, 12]) since the internal data structures (in particular, the "significance lists" of zerotree-based algorithms such as [2]) employed by the algorithm may grow to prohibitive sizes. Finally, and more importantly for the present application, 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. The presence of land points in our data effectively eliminates the possibility of employing the KLT as is done in [6–9] and increases the implementation complexity of zerotree-based approaches such as [8, 9, 12] to the point of infeasibility.

We circumvent the problems associated with prior techniques by introducing the Wavelets Around Land Masses (WAVAL) system for the coding of 3D oceanographic images. Our WAVAL system 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. Our SARL algorithm is inspired by stack-run (SR) coding [13], a computationally simple method for coding wavelet coefficients with performance surprising competitive with more complicated zerotree methods [1, 2]. As in SR, our SARL coder relies on simpler runlength coding, rather than zerotrees, for efficient coding of large regions of insignificant coefficients; however, contrary to SR, quantization of significant coefficients takes place in an embedded manner.

In the remainder of this paper, we elaborate upon the WAVAL system and investigate its performance. Specifically, in the next section, we present an in-depth description of WAVAL, including details behind the wavelet transform and the SARL algorithm. Then in Section III, we overview experimental results comparing performance of WAVAL to the technique currently in use by the Navy to transmit oceanographic image data. Finally, we make some concluding remarks in Section IV. We note that, although in this paper we will be concerned with the coding of oceanographic images consisting of primarily water-temperature values, our WAVAL system can be applied directly to other scalar quantities, such as salinity, sound speed, vorticity, etc., as well.

II. The WAVAL System

The encoder in the WAVAL system 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. We describe each of these components below.

A. 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 encoder and decoder, an encoding of the land-sea mask is the initial portion of the bitstream output by the encoder. The entire land-sea mask must be received before decoding begins; therefore the mask encoding is not embedded.

Encoding of the binary land-sea mask consists of runlength coding of consecutive like values encountered in a raster-scan traversal of the mask. 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, so we will delay further discussion of this encoding until Section II-C (see below and also Fig. 2). The symbol stream resulting from the runlength encoder is entropy coded using an arithmetic coder [14] 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 on deeper levels directly below points labeled as land at a some shallower depth need not be coded—they are already known to be land. The raster-scan and runlength coding of the land-sea mask thus skips such land points as the scan proceeds deeper in the 3D volume.

B. 3D Wavelet Transform

The WAVAL system uses a lifting implementation of the Cohen-Daubechies-Feauveau $(2, 2)$ biorthogonal wavelet [15] (i.e., "linear lifting" [16]), 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 subbands. Our strategy for transforming data in between land masses is similar to other transforms (e.g., [17–19]) that have been proposed recently for the coding of arbitrarily shaped objects arising in modern video coders. We note that these prior techniques center around filter-bank transforms, whereas our lifting-based approach is signifi-

cantly faster and somewhat simpler to implement in practice.

In the usual implementation (e.g., [16]), 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. For linear lifting, we have (see [16] for details)

$$
\begin{aligned}
c'_j[k] &= c_{j+1}[2k] && \text{even samples} \\
d'_j[k] &= c_{j+1}[2k+1] && \text{odd samples} \\
d''_j[k] &= d'_j[k] - \frac{1}{2}\left(c'_j[k] + c'_j[k+1]\right) && \text{lifting prediction} \\
c''_j[k] &= c'_j[k] + \frac{1}{4}\left(d''_j[k-1] + d''_j[k]\right) && \text{lifting update} \\
d_j[k] &= \frac{1}{\sqrt{2}} d''_j[k] && \text{highpass coefficients} \\
c_j[k] &= \sqrt{2} c''_j[k] && \text{lowpass coefficients}
\end{aligned}
$$

where $c_{j+1}[k]$ is the lowpass subband at scale $j+1$, and $c_j[k]$ and $d_j[k]$ are the lowpass and highpass subbands, respectively, at the next coarser scale, j.

Let us consider a single row, column, or depth passing through our 3D dataset and call this 1D data structure a *signal*. We observe that each 1D signal is composed of one or more 1D *segments* of consecutive sea data values separated by land. In our transform, each segment is transformed individually into lowpass and highpass *segment subbands*. The lowpass segment subbands from all the segments of the signal are collected together with land points placed in between to form the *lowpass band* of the signal; the highpass segment subbands likewise form the *highpass band*. Our transform enforces a global subsampling pattern that ensures that each lowpass or highpass segment subband is placed in the corresponding lowpass or highpass band, respectively, so as to maintain the relative spatial positioning of the segments in the original signal (a process called "phase alignment" in [17]).

Symmetric extension is used at either end of each 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 four cases which arise from the fact that a segment may start with an even- or odd-indexed sample and end with an even- or odd-indexed sample. Additionally, since the length of a segment may be arbitrarily odd or even, the corresponding lowpass and highpass segment subbands are not necessarily of the same length and may even be of zero length, this latter case occurring when a segment consists of only one data value so that its lifting decomposition degenerates to a multiplication or division by $\sqrt{2}$ to maintain unitary scaling. Finally, we note that, due to the robust manner in which the transform must handle segments of arbitrary lengths, no restriction is placed on the size of the overall 3D data volume; i.e., the WAVAL system is under no "power-of-2" size restrictions as is common for implementations of other wavelet-based coders.

C. Successive-Approximation Runlength Coding (SARL)

The embedded zerotree wavelet (EZW) algorithm [1], designed originally for 2D images, is an easily implemented, computationally efficient, embedded technique with effective coding performance. Recent refinements of the EZW technique, such as the set partitioning in hierarchical trees (SPIHT) [2], have been widely employed in 2D image-coding applications. As expected given their proven performance in 2D applications, zerotree techniques have been extended to 3D datasets (e.g., [8, 9, 12]). However, the coding efficiency of EZW and related algorithms relies on the representing of large regions of zeros with single zerotree symbols. In our application, the potentially sparse nature of the oceanographic datasets hinders zerotree performance and significantly complicates implementation. Additionally the "significance lists" of algorithms such as [2] may grow to prohibitive sizes quickly as the dimensions 3D datasets increase [9]. On the other hand, SR coding [13] achieves 2D-image performance competitive with zerotree techniques without relying on the zerotree structure or associated significance-list complexity. Its implementation is extremely simple and is thus amenable to applications that are more general than 2D-image coding—the only drawback lies in that SR is not embedded.

For our WAVAL system, we developed the SARL algorithm which 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
  $runlength = 0$
  $difference\_map(n) = 0, \forall n$
  $AC\_context = \mathbf{RUNLENGTH}$
  for n = 1 to N do

```
  if mask(n) = SEA
    if significance_map(n) = 0
      if |x(n)| ≥ threshold
        significance_map(n) = 1
        difference_map(n) = 1
        output_run (runlength, sign of x(n))
        x(n) = |x(n)|
        runlength = 0
      else
        runlength = runlength + 1
done
AC_context = REFINEMENT
for n = 1 to N do
  if mask(n) = SEA
    if significance_map(n) = 1
      if x(n) ≥ threshold
        if difference_map(n) = 0
          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 [13]. The SARL runlength-encoding scheme is illustrated in Fig. 2; note that the land-sea mask is encoded using this same procedure. The runlengths are coded by outputting the binary representation of the runlength integer, using symbols "0" and "1," starting from the least-significant bit. A "punctuation" symbol, "+" or "-," is output between consecutive runlengths. For SARL, the punctuation symbol serves not only to delineate the start of one run encoding from another, but also to transmit the sign of the wavelet coefficient (for the land-sea mask, the punctuation symbol is always "+"). Finally, since the most-significant bit of each runlength is always a "1," this symbol is omitted. The stream of symbols resulting from the runlength encodings and the refinement symbols is entropy coded using arithmetic coding [14] 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.

D. Channel Packing

The channel employed by the Navy for distribution of oceanographic data uses 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 system 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.

III. Experimental Results

In this section, we compare the performance of the proposed WAVAL system 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 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, the entire resulting volume is subject to scalar quantization and runlength coding, and a nonembedded bitstream is output. Being slice-based rather than 3D, the transform used by EOF requires significantly less memory than that of WAVAL; however, this fact has not hindered the use WAVAL in practice since the largest of the Navy's datasets (on the order of $500 \times 500 \times 34$, or 35Mb) fits easily into the spacious memory of modern computers.

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 in these plots 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 I summarizes the sizes of the datasets and the composition of the bitstream produced by the WAVAL system. Fig. 6 shows the difference images formed by subtracting the original `ylsoj` dataset from the reconstructed datasets for both approaches.

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 coastline of the Sea of Japan area. The land-sea mask for this dataset (nonembeddedly coded by WAVAL at 99 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, the observed superior EOF performance is not really meaningful since the distortion achieved by both algorithms is too large to be useful in the Navy's oceanographic applications. These applications require near-lossless performance, i.e., maximum absolute error on the order of $0.1°$C, the precision of the original datasets. For such near-lossless coding, WAVAL consistently outperforms EOF on all datasets tested.

IV. Conclusions

In this paper, we have described the WAVAL system for the embedded coding of 3D oceanographic images. To this end, we have employed a 3D lifting wavelet transform tailored specifically to the potentially sparse nature of oceanographic images

and have introduced the SARL embedded-coding procedure which adds EZW-like successive-approximation properties to SR. Our 3D 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 and impossible with the KLT. Additionally, our SARL coder avoids zerotree structures whose performance is hindered by the presence of land points and significance-list implementations whose complexity quickly becomes prohibitive for large 3D datasets.

As a final remark, we note that, 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. Our investigations have revealed that, although there is some performance penalty associated with the embedded nature of SARL as compared to nonembedded SR, it appears that it is minimal. For example, when applied to the coding of the ubiquitous 2D grayscale Lenna image at 0.5 bits/pixel, SARL produces a PSNR on the order of only 0.7 dB below that of SR (both algorithms using the $(2, 2)$ biorthogonal wavelet used here); however, the perceptual-quality performance of both algorithms is identical for this example.

Future work includes extending the WAVAL system to other datasets arising in geoscience applications, including atmospheric measurement data and vector-valued data such as ocean currents. As part of this planned work, we anticipate extending both the wavelet transform and the SARL algorithm to vector-valued data; we are currently investigating vector-valued transforms based on multiwavelets for the first task and a variety of vector-quantization techniques for the second. A final future-work item is the incorporation of an error-control mechanism that will allow the user to specify a desired error tolerance in advance of coding. For the WAVAL coder, such error control is complicated by factors such as the use of a biorthogonal wavelet and maximum absolute error (i.e., Parseval's theorem does not apply and an inverse transform must be performed to know distortion). However, our preliminary investigations indicate that a small number of inverse-transform calculations plus an intelligent search strategy can narrow in on the desired error performance in a timeframe acceptable for practical applications. We are presently implementing and refining this error-control strategy.

REFERENCES

[1] J. M. Shapiro, "Embedded Image Coding Using Zerotrees of Wavelet Coefficients," *IEEE Transactions on Signal Processing*, vol. 41, no. 12, pp. 3445–3462, December 1993.

[2] A. Said and W. A. Pearlman, "A New, Fast, and Efficient Image Codec Based on Set Partitioning in Hierarchical Trees," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 6, no. 3, pp. 243–250, June 1996.

[3] ISO/IEC 13818-2, *Information Technology—Generic Coding of Moving Pictures and Associated Audio Information: Video*, 1995, MPEG-2 Video Coding Standard.

[4] G. P. Abousleman, M. W. Marcellin, and B. R. Hunt, "Compression of Hyperspectral Imagery Using the 3-D DCT and Hybird DPCM/DCT," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 33, no. 1, pp. 26–34, January 1995.

[5] S. Gupta and A. Gersho, "Feature Predictive Vector Quantization of Multispectral Images," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 30, no. 3, pp. 491–501, May 1992.

[6] J. A. Saghri, A. G. Tescher, and J. T. Reagan, "Practical Transform Coding of Multispectral Imagery," *IEEE Signal Processing Magazine*, vol. 12, no. 1, pp. 32–43, January 1995.

[7] B. R. Epstein, R. Hingorani, J. M. Shapiro, and M. Czigler, "Multispectral KLT-Wavelet Data Compression for Landsat Thematic Mapper Images," in *Proceedings of the IEEE Data Compression Conference*, J. A. Storer and M. Cohn, Eds., Snowbird, UT, March 1992, pp. 200–208.

[8] P. L. Dragotti, G. Poggi, and A. R. P. Ragozini, "Compression of Multispectral Images by Three-Dimensional SPIHT Algorithm," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 38, no. 1, pp. 416–428, January 2000.

[9] F. Amato, C. Galdi, and G. Poggi, "Embedded Zerotree Wavelet Coding of Multispectral Images," in *Proceedings of the International Conference on Image Processing*, Santa Barbara, CA, October 1997, pp. 612–615.

[10] 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.

[11] J. Wang and H. K. Huang, "Medical Image Compression by Using Three-Dimensional Wavelet Transformation," *IEEE Transactions on Medical Imaging*, vol. 15, no. 4, pp. 547–554, August 1996.

[12] B.-J. Kim and W. A. Pearlman, "An Embedded Wavelet Video Coder Using Three-Dimensional Set Partitioning in Hierarchical Trees (SPIHT)," in *Proceedings of the IEEE Data Compression Conference*, J. A. Storer and M. Cohn, Eds., Snowbird, UT, March 1997, pp. 251–257.

[13] M.-J. Tsai, J. D. Villasenor, and F. Chen, "Stack-Run Image Coding," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 6, no. 5, pp. 519–521, October 1996.

[14] I. H. Witten, R. M. Neal, and J. G. Cleary, "Arithmetic Coding for Data Compression," *Communications of the ACM*, vol. 30, no. 6, pp. 520–540, June 1987.

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

[16] I. Daubechies and W. Sweldens, "Factoring Wavelet Transforms into Lifting Steps," *The Journal of Fourier Analysis and Applications*, vol. 4, no. 3, pp. 245–267, 1998.

[17] J. Li and S. Lei, "Arbitrary Shape Wavelet Transform with Phase Alignment," in *Proceedings of the International Conference on Image Processing*, Chicago, IL, October 1998, pp. 683–686.

[18] A. Mertins, "Optimized Biorthogonal Shape Adaptive Wavelets," in *Proceedings of the International Conference on Image Processing*, Chicago, IL, October 1998, pp. 673–677.

[19] B.-F. Wu and C.-Y. Su, "Arbitrarily Shaped Image Coding by Using Translation Invariant Wavelet Transforms," *Signal Processing*, vol. 79, no. 3, pp. 309–314, December 1999.

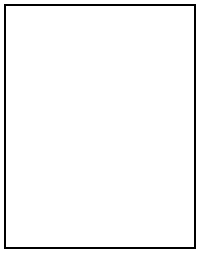

James E. Fowler received the B.S. degree in computer and information science engineering and the M.S. and Ph.D. degrees in electrical engineering in 1990, 1992, and 1996, respectively, all from the Ohio State University. In 1995, he was an intern researcher at AT&T Labs in Holmdel, NJ, and, from January to July of 1997, he held an NSF-sponsored postdoctoral assignment at the Université de Nice-Sophia Antipolis, France. He is currently an assistant professor in the Department of Electrical and Computer Engineering at Mississippi State University in Starkville, MS, and a researcher at the NSF Engineering Research Center for Computational Field Simulation at Mississippi State. His research interest include image and video coding, data compression, and wavelets.

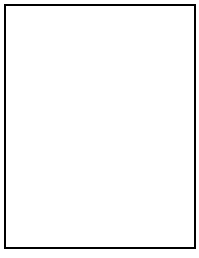

Daniel N. Fox received the B.S. degree in physics in 1975 from the Illinois Institute of Technology, and the M.S. degree in physics in 1976 from the University of Maryland. He is presently completing the Ph.D. in scientific computing from the University of Southern Mississippi. From 1977 to 1987, he was a geophysicist with Shell Oil in New Orleans, LA, and Houston, TX. Since 1987 he has been a research physicist at the Naval Research Laboratory at Stennis Space Center, MS, where he is the principal investigator of the On-Scene Tactical Ocean Forecast Capability and the Relocatable Coastal Ocean Modeling projects. His research interests include data assimilation and rapidly relocatable ocean analysis and forecast systems.

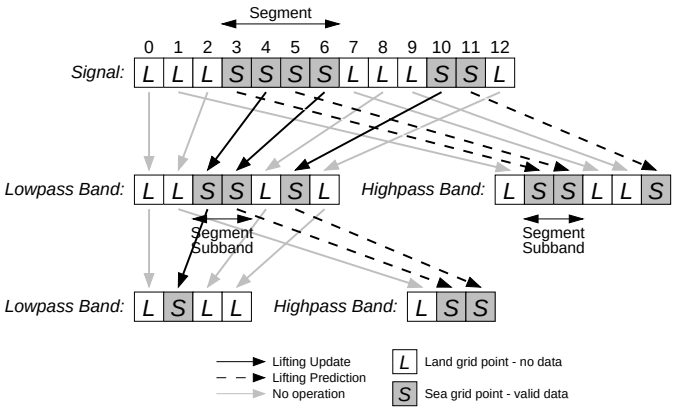

Fig. 1. An example of two stages of 1D wavelet decomposition. Each segment is decomposed into lowpass and highpass segment subbands; segment subbands are assembled into lowpass and highpass bands maintaining appropriate subsampling pattern and positioning within the band.

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	...
								$\pm$	$\pm$	...

Fig. 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 sign of wavelet coefficient; for the land-sea mask, $\pm$ is always $+$.

TABLE I

ORIGINAL DATASET SIZES AND COMPOSITION OF WAVAL OUTPUT FOR A MESSAGE SIZE OF 500 LINES.

Original Dataset			*WAVAL Output*		
Name	*Size*	*Percent Land*	*Header*	*Land-Sea Mask*	*Wavelet Coeffs.*
`hawai`	$91\times81\times34$	1.6%	0.1%	2.0%	97.9%
`adrtc`	$81\times61\times34$	87.6%	0.1%	3.1%	96.8%
`ylsoj`	$151\times101\times34$	53.1%	0.1%	20.4%	79.5%

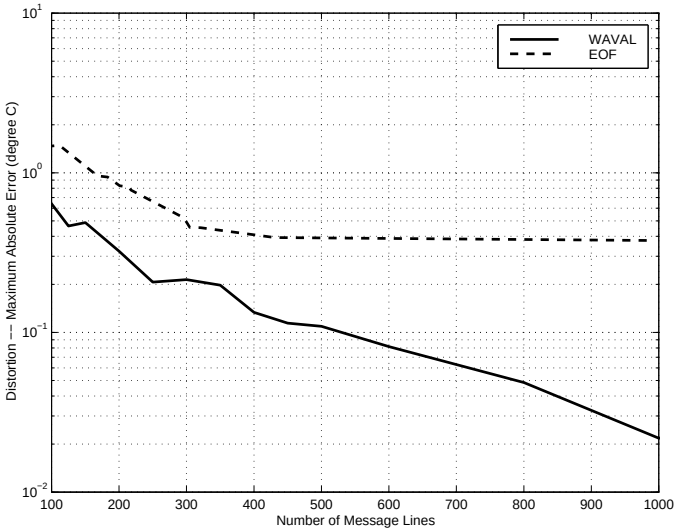

Fig. 3. Rate-distortion performance for the `hawai` dataset.

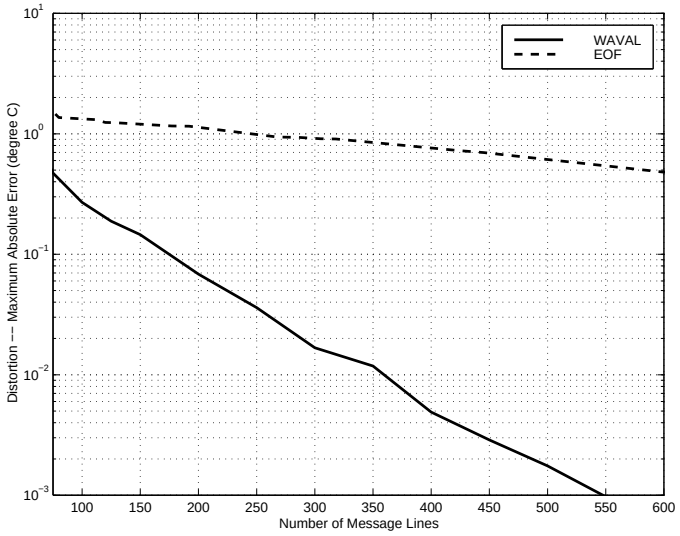

Fig. 4. Rate-distortion performance for the `adrtc` dataset.

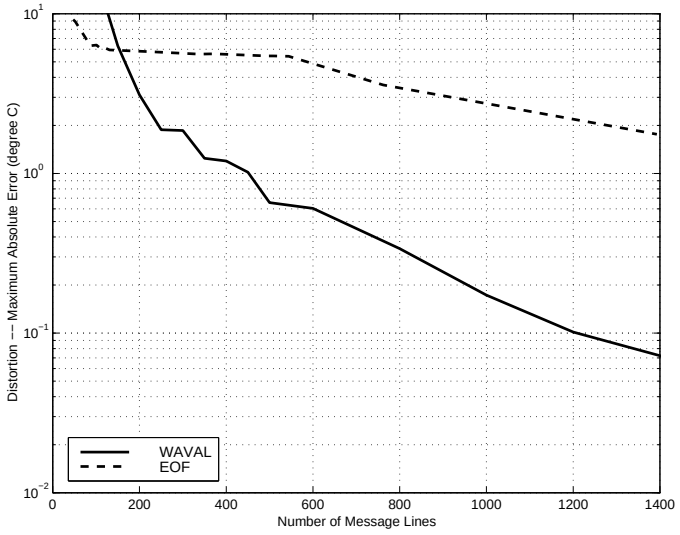

Fig. 5. Rate-distortion performance for the `ylsoj` dataset.

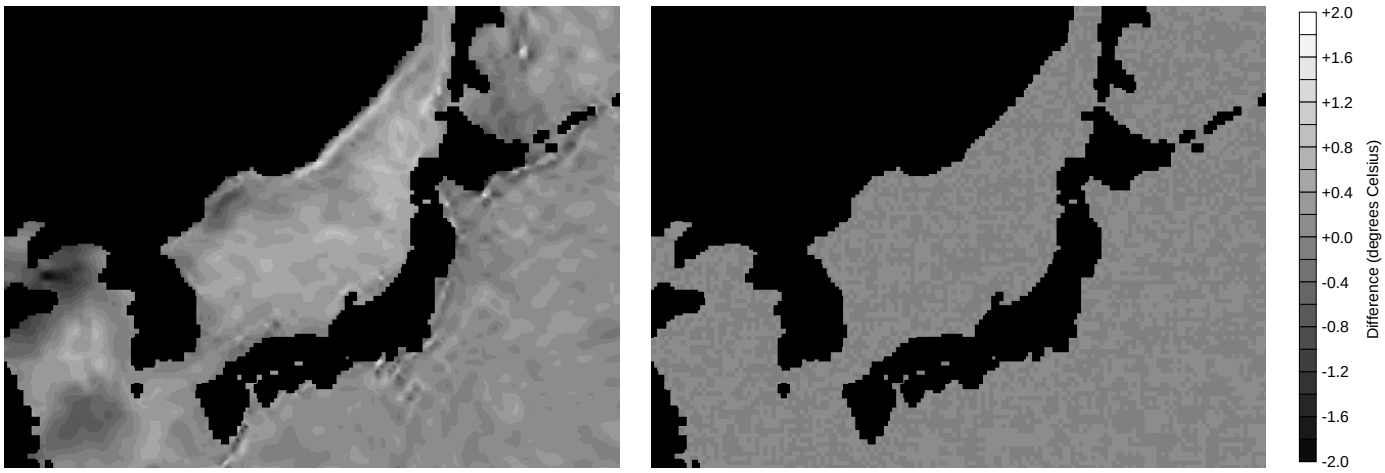

(a) (b)

Fig. 6. Image of differences between original and reconstructed `ylsoj` datasets at the ocean-surface level; reconstructed images produced from 762-line messages; solid black regions are land masses. (a) EOF, (b) WAVAL.