Applications of Wavelet-Based Compression to Multidimensional Earth Science Data

Jonathan N. Bradley and Christopher M. Brislawn
Computer Research Group, Los Alamos National Laboratory
Mail Stop B-265, Los Alamos, New Mexico 87545

ABSTRACT

A data compression algorithm involving vector quantization (VQ) and the discrete wavelet transform (DWT) is applied to two different types of multidimensional digital earth-science data. The algorithm (WVQ) is optimized for each particular application through an optimization procedure that assigns VQ parameters to the wavelet transform subbands subject to constraints on compression ratio and encoding complexity.

Preliminary results of compressing global ocean model data generated on a Thinking Machines CM-200 supercomputer are presented. The WVQ scheme is used in both a predictive and nonpredictive mode. Parameters generated by the optimization algorithm are reported, as are signal-to-noise ratio (SNR) measurements of actual quantized data. The problem of extrapolating hydrodynamic variables across the continental landmasses in order to compute the DWT on a rectangular grid is discussed.

Results are also presented for compressing Landsat TM 7-band data using the WVQ scheme.The formulation of the optimization problem is presented along with SNR measurements of actual quantized data. Postprocessing applications are considered in which the seven spectral bands are clustered into 256 clusters using a k-means algorithm and analyzed using the Los Alamos multispectral data analysis program, SPECTRUM, both before and after being compressed using the WVQ program.

I. INTRODUCTION.

This work describes the application of an image compression algorithm involving the discrete wavelet transform and vector quantization to two problems involving earth science data. The coding of outputs of supercomputer-generated global climate model (GCM) ocean simulations and Landsat Thematic Mapper (TM) multispectral imagery is investigated. The compression algorithm has its origins in the coding of gray-scale imagery [1, 2]. A set of vector quantizers is designed (one for each subband in the wavelet decomposition) with parameters selected from the solution of an optimization problem that is formulated to minimize quantization distortion with constraints on the overall bit rate and encoding complexity. Although both data types considered in this work are of dimensionality higher than two, we restrict the discussion to coding implementations based on 2-D transforms.

Compression of the GCM data is approached by both a straightforward two-dimensional extension of the earlier algorithm and a predictive scheme in which two-dimensional prediction residuals are coded. The Landsat images are coded by modifying the bit allocation algorithm to allocate coder resources simultaneously among all of the spectral components. For all scenarios, measurements of quantization distortion are presented as functions of compression ratio and encoding complexity, thus revealing tradeoffs involved in the system design.

II. Multidimensional Wavelet Transform–Vector Quantization.

The data-coding technique used in this work, known as the *wavelet–vector quantization* (WVQ) algorithm, is based on vector quantization of the subbands resulting from a discrete wavelet transform (DWT) decomposition of the data signal. For signals in two or more dimensions, the transform used is based on product filter banks (i.e., tensor products of one-dimensional DWT filters). A d-dimensional signal transformed in this manner with a two-channel filter bank yields 2^d subbands, any of which can be cascaded back through the filter bank to produce a multirate decomposition of the original signal. Although we are currently working on three-dimensional wavelet transforms for use with the three-dimensional climate model data under investigation, the DWT results presented here are restricted to the case of two-dimensional data fields.

Single-level 2-D DWT analysis and synthesis filter banks are depicted in Figures 1 and 2. The analysis filters (H_i) and synthesis filters (F_i) used in this paper are biorthogonal linear phase FIR wavelet filters constructed in [3, 4]. Note the use of binary subscripts on the subbands, a_{ij}, to indicate the filters applied to the rows and columns of the signal, x. Signals obtained from sampling smooth, continuous data fields usually have most of their energy (or variance) concentrated in the low-frequency part of the spectrum, so it is usually most efficient to cascade only the lowpass-lowpass filtered subband, a_{00}, back through the analysis bank in Figure 1; the resulting subband is denoted $a_{00,00}$, or $a_{,00}$ for short. This cascade is typically carried down four levels (or so), to the $a_{,,,00}$ band, which will then contain a large portion of the signal energy concentrated in a heavily downsampled signal component. Note that the downsample factor for an l^{th}-level subband in a d-dimensional scheme using an M-channel coder is $m_i = M^{ld}$, so, e.g., subband $a_{,,,00}$ in the 2-D decomposition has been downsampled by $m_i = 256$-to-1.

A further consideration when applying a DWT to finite-duration signals, like the rows or columns in a digital image, is the handling of boundary conditions. The most straightforward way of dealing with signal boundaries is to regard the signal as a single period of an infinite, periodic input and apply the DWT filter bank by circular convolution and downsampling. This has the disadvantages, however, of introducing a spurious jump discontinuity when the data isn't inherently periodic and of constraining the signal length (i.e., its "period") to be divisible by the downsample factor. For a four-level decomposition using a two-channel filter bank, for instance, this means the input length, N_0, must be divisible by 16. Both of these problems can be avoided by using symmetric extrapolation techniques to extend finite-duration inputs; moreover, this can be done with no increase in the memory allocation needed to transform or store the input signal [5, 6].

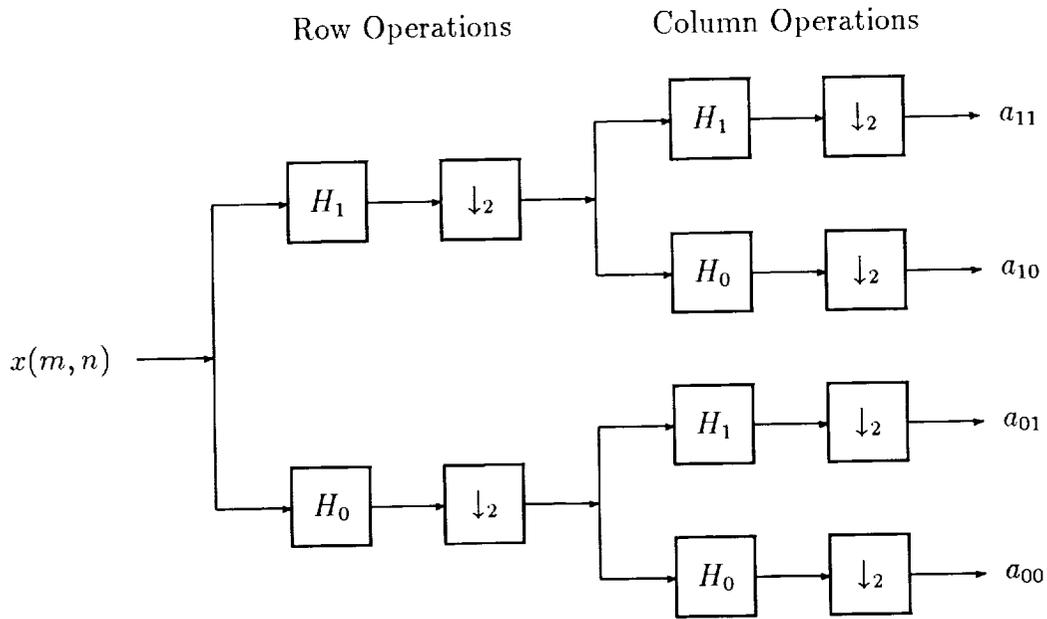The design of vector quantizers for the subbands in a DWT decomposition is based on

14

Row Operations · Column Operations

Figure 1: Single-Level, Two-Dimensional Wavelet Transform Analysis.

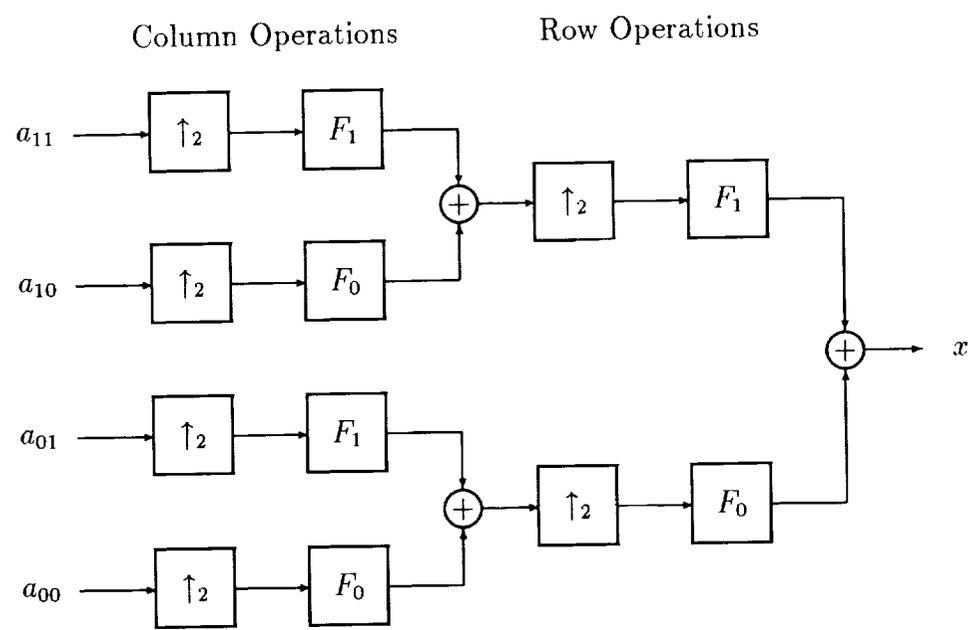Column Operations · Row Operations

Figure 2: Single-level, Two-Dimensional Wavelet Transform Synthesis.

the assumption of exponential VQ rate-distortion characteristics,

$$D_i(k_i, r_i) = \beta_i(k_i)e^{-\gamma_i(k_i)r_i} \qquad (1)$$

$D_i(k_i, r_i)$ is the distortion (mean-square error) between the original and quantized data in the i^{th} subband for a bit rate of r_i bits per pixel (bpp) and a vector dimension of k_i. $\beta_i(k_i)$ and $\gamma_i(k_i)$ are constants that depend on k_i and the probability density function of the data vectors. The motivation for this assumption is based on theoretical VQ rate-distortion modelling [7] and confirmed by empirical data; values for the constants $\beta_i(k_i)$ and $\gamma_i(k_i)$ are determined from a set of training data.

In the case of an orthogonal subband decomposition, the overall distortion can be expressed as a weighted sum of the distortion in each subband,

$$D = \sum_i \frac{1}{m_i} D_i(k_i, r_i) \quad ,$$

where m_i is the downsample factor, the ratio of the number of samples in the original to the number in the i^{th} subband. Since the DWT conserves the number of data samples, m_i satisfies the identity $\sum m_i^{-1} = 1$. By (1), the overall distortion is

$$D = \sum_i \frac{1}{m_i} \beta_i(k_i)e^{-\gamma_i(k_i)r_i} \quad . \qquad (2)$$

Formula (2) is customarily used as a distortion measure with nonorthogonal transforms, too, although it no longer coincides exactly with overall mean-square error. The bit-allocation problem for quantizer design involves using nonlinear optimization techniques to compute the bit rates, r_i, and dimensions, k_i, that minimize (2), subject to constraints on overall bit rate and encoder complexity.

For a target overall bit rate of R bpp, the constraint on subband bit rates is

$$\sum_i \frac{r_i}{m_i} \leq R \quad . \qquad (3)$$

If subband vector dimensions, k_i, are to be optimized, an additional constraint besides (3) is necessary to obtain a well-posed optimization problem for VQ bit allocation. The encoder complexity constraint used here is an upper bound, Q, on the computational cost of performing exhaustive nearest-neighbor searches of k_i-dimensional VQ codebooks containing $N_i = 2^{k_i r_i}$ codevectors:

$$\sum_i \frac{1}{m_i} 2^{k_i r_i} \alpha \leq Q \quad . \qquad (4)$$

The parameter α is a constant corresponding to the arithmetic cost of performing two additions and one multiplication per pixel. With the additional constraints $r_i \geq 0$ we obtain a convex nonlinear optimization problem to solve for the r_i; the k_i are optimized by a heuristic search procedure. Once optimal bit rates and vector dimensions are computed, optimal VQ codebooks are constructed from training data using the Linde-Buzo-Gray method [8, 9]. More details about the WVQ algorithm are given in [10, 2, 1, 11].

16

III. APPLICATION TO OCEAN MODEL DATA.

This section describes the use of the WVQ algorithm on synthetic data generated by a Bryan-Cox-Semtner global ocean circulation model running on the Connection Machine CM-200 at the Los Alamos Advanced Computing Laboratory (ACL) [12, 13]. The model is computed on a 320×768 grid at 20 depth levels; boundary conditions are given on a three-dimensional bottom topography with 80 islands. The data used in the compression experiments was the surface temperature field (no depth components), taken at three-day intervals over a decade's worth of simulation. Time-frames from the first year of the simulation were used for training data, and the resulting WVQ algorithm was then tested on frames from the last year of the simulation, i.e., on data similar but not identical to the training data. We feel this is a valid test since it is similar to the manner in which the algorithm will be used in practice.

The two-dimensional data frames were transformed with a four-level octave-scaled DWT decomposition. Since the model is periodic in the east-west direction, periodic boundary conditions were used along parallels of latitude ("rows"). However, due to the lack of continuity between the north and south edges of the grid, symmetric (i.e., reflected) boundary conditions were used along meridians of longitude ("columns"). Because the DWT is most easily computed on a rectangular grid, the temperature data was extended across the continental landmasses before transforming. A simple approach like zero-padding of the data would be undesirable because it would induce a large jump discontinuity around the coastlines; this would show up as added variance in the highpass-filtered DWT subbands and would therefore reduce the compressibility of the high-frequency signal components in the transform domain. For this reason we used a continuous extension of the data given by linear extrapolation from coast to coast along parallels of latitude. This still leaves a "corner" at the coastlines in the extended data; since the initial data field is extremely smooth, this corner results in a slight increase in energy in highpass-filtered subbands. It is not yet clear whether this added variance is significant alongside the variance naturally present in the data. We are currently looking into using smoother two-dimensional extrapolation schemes for this task.

Two different approaches were taken to quantizing the time-series data generated in the simulation: nonpredictive and predictive coding. In the nonpredictive scheme, each frame is treated as a separate image and compressed accordingly using the WVQ method. In predictive coding, a prediction of each frame is made based on past frames and subtracted from the current frame, resulting in a two-dimensional *residual* image, which is compressed and stored. The image sequence is decoded from the first frame in the sequence and the residuals. We used a simple first-order predictor in this scheme; i.e., the prediction of a given frame is just equal to the quantized value of the previous frame. Block diagrams for the transmitter and receiver in this predictive encoding/decoding system are given in Figures 3 and 4. The experiment assumed that the first frame in the sequence is transmitted with nonpredictive quantization, and the compression ratios reported are those of the subsequent residuals.

For both the nonpredictive and predictive schemes, WVQ coders were designed for bit rates, R, ranging from 2.0 to 0.25 bpp. Since the original data was 32 bpp, the corresponding compression ratios range from 16:1 to 128:1. Encoding complexities, Q, were varied between
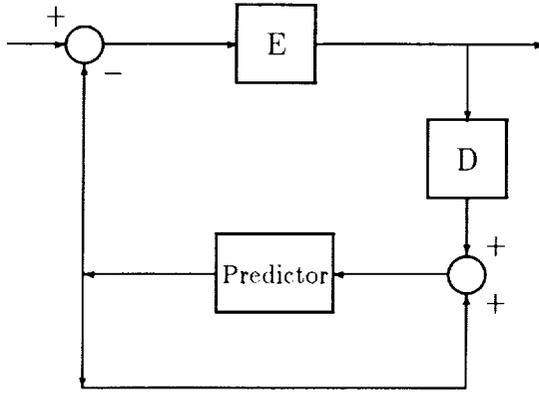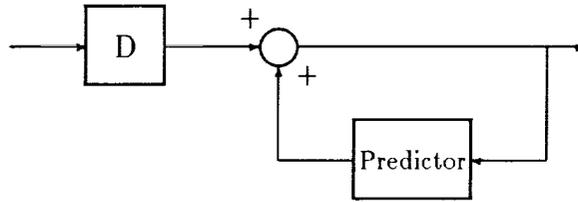
17

Figure 3: Predictive Transmitter.

Figure 4: Predictive Receiver.

16α and 64α. The optimal codebook sizes and vector dimensions for each combination of R and Q were computed by the WVQ design algorithm described in Section II. Bit allocation results for the 13 subbands are presented in Tables I and II for $R = 0.5$ and 0.25 bpp in terms of vector dimensions, k_i, and codebook sizes, N_i. Note that as the bit rate decreases, the highest frequency subbands are quantized more heavily or discarded altogether and remaining high-frequency subband vector dimensions typically increase. Vector dimensions also increase as the upper bound on complexity increases. Bit allocation for the residual subbands in Table II is similar to that for the nonpredictive scheme, although much less of the quantizer resources (i.e., far fewer bits) are allocated to the lower frequency subbands in the predictive scheme. This means that the first-order predictor effectively predicts the low-wavenumber modes of the model, indicating that these modes are evolving slowly compared to the sampling rate for archiving data.

Quantizer performance is measured in terms of signal-to-noise ratio (SNR),

$$\text{SNR} = 10 \log_{10} \frac{\sigma_s^2}{\sigma_e^2} \quad \text{(dB)} \quad ,$$

where σ_s^2 is the signal variance and σ_e^2 is the quantization error variance. The average SNR in dB for the test data is shown in Figures 5 and 6 for various values of R and Q. This diagram illustrates the various tradeoffs involved in the selection of R and Q. At a given bit rate, R, note that a higher SNR is possible using an encoder with higher complexity, Q; i.e., higher subband vector dimensions, k_i, and correspondingly larger codebook sizes, N_i. Note

18

Table I: Vector Dimension and Codebook Size Assignments (k, N) for $R = 0.5$ bpp and $R = 0.25$ bpp, Nonpredictive Coding.

	$R = 0.5$			$R = 0.25$		
	$Q = 64\alpha$	$Q = 32\alpha$	$Q = 16\alpha$	$Q = 64\alpha$	$Q = 32\alpha$	$Q = 16\alpha$
Subband ,,,00	(1,511)	(1,493)	(1,308)	(1,251)	(1,251)	(1,251)
Subband ,,,01	(1,132)	(1,119)	(1,91)	(1,32)	(1,31)	(1,32)
Subband ,,,10	(1,93)	(1,103)	(1,75)	(1,15)	(1,17)	(1,18)
Subband ,,,11	(1,48)	(1,58)	(1,41)	(1,8)	(1,7)	(1,7)
Subband ,,01	(4,723)	(2,247)	(2,133)	(4,377)	(4,382)	(4,295)
Subband ,,10	(4,261)	(2,61)	(2,41)	(4,16)	(4,16)	(4,17)
Subband ,,11	(4,134)	(4,76)	(2,25)	(4,4)	(4,4)	(4,4)
Subband ,01	(8,373)	(8,177)	(4,68)	(8,176)	(8,176)	(8,132)
Subband ,10	(8,19)	(8,13)	(8,8)	—	—	—
Subband ,11	—	—	(8,2)	—	—	—
Subband 01	(8,74)	(8,42)	(8,22)	(8,4)	(8,4)	(8,4)
Subband 10	—	—	—	—	—	—
Subband 11	—	—	—	—	—	—

Table II: Vector Dimension and Codebook Size Assignments (k, N) for $R = 0.5$ bpp and $R = 0.25$ bpp, Predictive Coding.

|	$R = 0.5$			$R = 0.25$		
	$Q = 64\alpha$	$Q = 32\alpha$	$Q = 16\alpha$	$Q = 64\alpha$	$Q = 32\alpha$	$Q = 16\alpha$
Subband ,,,00	(1,14)	(1,10)	(1,16)	(1,2)	(1,2)	(1,4)
Subband ,,,01	(1,21)	(1,15)	(1,23)	(1,3)	(1,4)	(1,6)
Subband ,,,10	(1,26)	(1,20)	(1,28)	(1,5)	(1,6)	(1,9)
Subband ,,,11	(1,35)	(1,28)	(1,35)	(1,8)	(1,10)	(1,14)
Subband ,,01	(4,383)	(4,235)	(2,84)	(4,207)	(4,198)	(4,120)
Subband ,,10	(4,423)	(2,125)	(2,86)	(4,293)	(4,257)	(4,140)
Subband ,,11	(4,305)	(4,185)	(2,64)	(4,93)	(4,109)	(4,87)
Subband ,01	(8,276)	(4,126)	(4,62)	(8,461)	(8,209)	(8,78)
Subband ,10	(8,239)	(8,138)	(8,54)	(8,206)	(8,145)	(8,65)
Subband ,11	(8,97)	(8,51)	(8,24)	—	(8,3)	(8,12)
Subband 01	(8,30)	(8,12)	(8,11)	—	—	—
Subband 10	—	—	—	—	—	—
Subband 11	—	—	—	—	—	—
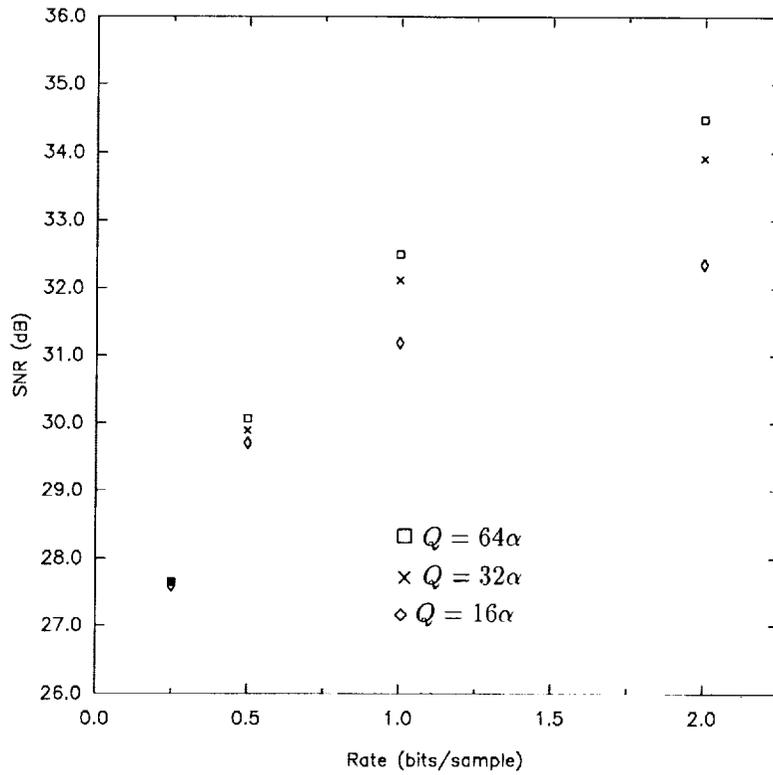
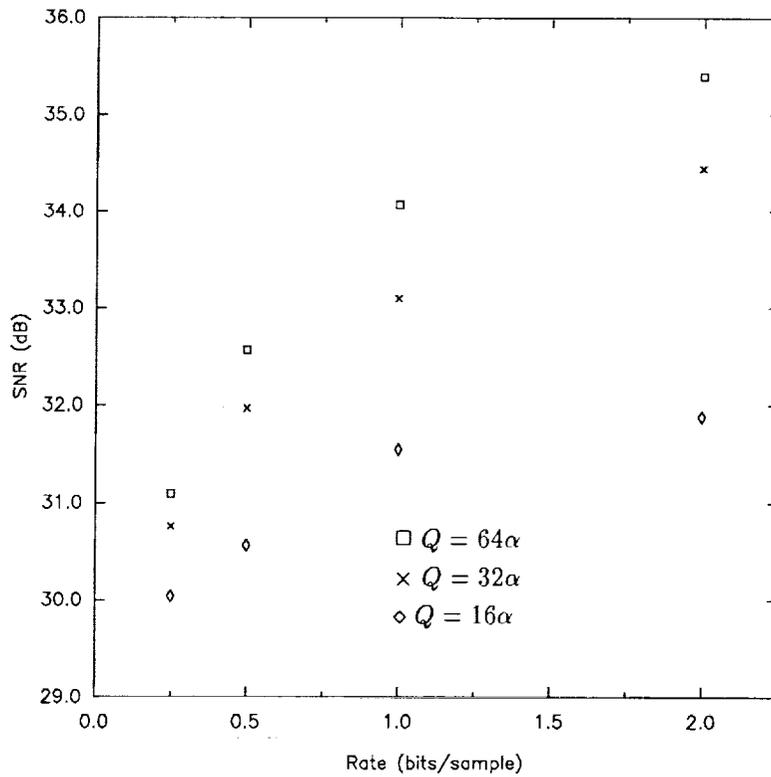Figure 5: SNR Measurements of Quantized Temperature Data, Nonpredictive.

Figure 6: SNR Measurements of Quantized Temperature Data, Predictive Coding.

20

that the complexity can be increased in this manner while the subband bit rates, r_i, remain unchanged since

$$r_i = \frac{\log_2 N_i}{k_i} \quad .$$

Increasing the encoding complexity results in an encoder with a more time-consuming codebook lookup but does not affect decoder performance. As the bit rate increases, we see from Figure 5 that the gain in SNR achieved by increasing the encoding complexity becomes more significant. For a fixed R and Q, comparison of Figures 5 and 6 shows that the predictive scheme achieves a gain on the order of 1–4 dB over nonpredictive coding. The improvement in coding gain is more pronounced at lower bit rates and higher complexities, since the residual coding scheme is better able to exploit higher limits on encoding complexity at low bit rates than the nonpredictive scheme.

IV. APPLICATION TO MULTISPECTRAL DATA.

This section discusses the application of WVQ to the compression of multispectral imagery. Since each spectral band is a separate monochromatic image, the approach is to code each of the bands by two-dimensional WVQ using symmetric boundary conditions. The bit-allocation is performed for the various spectral components simultaneously and hence the coding of each spectral component is not viewed as a separate two-dimensional problem.

The multispectral problem requires a modification to the WVQ design algorithm discussed in Section II since the rate and complexity are expressed in terms of multidimensional pixels. For the case of L spectral components the system design procedure entails minimizing

$$D = \frac{1}{L} \sum_i \frac{1}{m_i} \beta_i(k_i) e^{-\gamma_i(k_i) r_i} \tag{5}$$

over the k_i and r_i subject to

$$\frac{1}{L} \sum_i \frac{r_i}{m_i} \leq R \tag{6}$$

$$\frac{1}{L} \sum_i \frac{1}{m_i} 2^{r_i k_i} \alpha \leq Q \tag{7}$$

$$r_i \geq 0 \tag{8}$$

$$k_i \in K_i \tag{9}$$

where K_i denotes a prespecified set from which k_i must be selected. The optimization is performed over all of the two-dimensional subbands generated from all of the spectral components.

Multispectral image WVQ was considered for the application of compressing Landsat Thematic Mapper (TM) data. Such data consist of seven 8-bit spectral bands (three visible, three infrared, and one thermal) at a ground sample distance of 28.5 meters. Four data sets were used in training the coder: Albuquerque, NM (2984×3356); Cairo, Egypt (2945×3320); Los Alamos, NM (2984×3254); and Mexico City, Mexico (5965×6967). The performance of the coder was evaluated in terms of results obtained by compressing a (2976×3552) scene from the Moscow, Russia, area containing both urban and agricultural areas. The resulting

21

Table III: RMSE Quantizer Performance as a Function of Compression Ratio and Complexity.

	$Q = 8\alpha$	$Q = 16\alpha$	$Q = 32\alpha$	$Q = 64\alpha$
16:1	3.04	2.35	1.92	1.70
32:1	3.06	2.63	2.43	2.26
64:1	3.85	3.51	3.32	3.24
128:1	4.76	4.71	4.58	4.56

root mean-square error (RMSE) for sixteen combinations of bit rate and encoding complexity are tabulated in Table III. The compression ratios reported are relative to 56 bits/pixel in the original data and assume that the bit rates satisfy

$$\frac{1}{L} \sum_i \frac{r_i}{m_i} = R \quad , \quad \text{with } L = 7 \quad . \tag{10}$$

The additional gain available from entropy and run-length coding is not included.

It is interesting to compare these results to those obtained by another wavelet-based compression technique. In [14] Landsat TM images were compressed via a subband decomposition of each spectral component by a 7-tap nonperfect reconstruction filter bank; each subband was coded with uniform scalar quantization followed by Huffman and zero-run-length coding. The experiment was repeated with an image-dependent Karhunen-Loeve transform (KLT) in the interband direction, which provided noticeable coding gain at the expense of computational complexity. The WVQ RMSE results depicted in Table III appear to lie between these two previous approaches, although any comparisons must be qualified by the fact that the numerical results in these two papers were obtained from different imagery (Kuwaiti oil fields in the case of [14]). For instance, Table III shows that 32:1 compression (1.75 bpp) with a complexity of $Q = 64\alpha$ yields an average RMSE per band of 2.26, or a little over 2 bits of error. The closest comparable value for non-KLT coding in [14] is a MSE of 40.02 at 2.51 bpp; dividing the MSE by 7 and taking a square root gives an average RMSE per band of 2.39, which is a slightly greater error at a higher bit rate than our result. With interband KLT coding, [14] reports a MSE of 25.11, or an average RMSE of 1.89, at 1.55 bpp; this is a lower distortion at a lower bit rate than our result. We are currently currently working on incorporating interband KLT coding with the WVQ compression method.

The motivation for our investigation of TM data compression is the need to store and process large amounts of data for postprocessing applications. Using the software package SPECTRUM [15], developed by Los Alamos National Laboratory and the University of New Mexico, we are able to use a desktop workstation running Unix and X-windows to analyze and categorize multispectral data that has been clustered into 256 clusters using a variant of the k-means algorithm. SPECTRUM can manipulate the color map for the computer display using any transformation of the clustered data, and can display cluster position as a two-dimensional scatter plot. Using these features, users are able to categorize data by selecting areas with a known type of land cover, causing all associated pixels in the image to be given the same pseudocolor representation. Of great interest to us is the robustness of SPECTRUM

data clustering when applied to data that has first been compressed by the WVQ algorithm. While visual quality of pseudocolor visualizations remains good after compression by as much as 32:1, it remains to be determined how much quantization distortion SPECTRUM can tolerate for tasks like Level 1 Land Use Categorization. We are attempting to establish quantitative distortion criteria based on the analysis of classification error presented in [15], which is based on computing levels of confidence for classifications done on clustered data.

REFERENCES

[1] J. N. Bradley, T. G. Stockham, Jr., and V. J. Mathews, "An optimal design procedure for intraband vector quantized subband coding," Tech. Rep. LA-UR-90-4372, Los Alamos National Lab., 1990. Submitted to *IEEE Trans. Commun.*

[2] J. N. Bradley and C. M. Brislawn, "Image compression by vector quantization of multiresolution decompositions," *Physica D*, vol. 60, pp. 245–258, 1992.

[3] A. Cohen, I. C. Daubechies, and J.-C. Feauveau, "Biorthogonal bases of compactly supported wavelets," Tech. Rep. 11217-900529-07 TM, AT&T Bell Labs, Murray Hill, NJ, May 1990.

[4] I. C. Daubechies, *Ten Lectures on Wavelets*. No. 61 in CBMS-NSF Regional Conf. Series in Appl. Math., (Univ. of Lowell, Lowell, MA, June 1990), Philadelphia, PA: Soc. Indust. Appl. Math., 1992.

[5] C. M. Brislawn, "Classification of symmetric wavelet transforms," Tech. Rep. LA-UR-92-2823, Los Alamos National Lab., Aug. 1992.

[6] J. N. Bradley, C. M. Brislawn, and V. Faber, "Reflected boundary conditions for multirate filter banks," in *Proc. 1992 IEEE-SP Int. Symp. on Time-Freq. and Time-Scale Analysis*, (Victoria, B.C.), IEEE Signal Processing Soc., Oct. 1992.

[7] A. Gersho, "Asymptotically optimal block quantization," *IEEE Trans. Info. Theory*, vol. 25, pp. 373–380, July 1979.

[8] Y. Linde, A. Buzo, and R. M. Gray, "An algorithm for vector quantizer design," *IEEE Trans. Commun.*, vol. 28, pp. 84–95, Jan. 1980.

[9] A. Gersho and R. M. Gray, *Vector Quantization and Signal Compression*. No. 159 in Int'l. Series in Engineering & Computer Science, Norwell, MA: Kluwer Academic Publishers, 1992.

[10] J. N. Bradley and C. M. Brislawn, "Wavelet transform–vector quantization compression of supercomputer ocean models," in *Proc. 1993 Data Compress. Conf.*, (Snowbird, UT), IEEE Computer Soc., Mar. 1993. To appear.

[11] J. N. Bradley and C. M. Brislawn, "Compression of fingerprint data using the wavelet vector quantization image compression algorithm," Tech. Rep. LA-UR-92-1507, Los Alamos National Lab., Apr. 1992. Progress report to the FBI.

[12] J. K. Dukowicz, R. D. Smith, and R. C. Malone, "A reformulation and implementation of the Bryan-Cox-Semtner ocean model on the Connection Machine," Tech. Rep. LA-UR-91-2864, Los Alamos National Lab., 1991. Submitted to *J. Phys. Oceanogr.*

[13] R. D. Smith, J. K. Dukowicz, and R. C. Malone, "Parallel ocean general circulation modeling," Tech. Rep. LA-UR-92-200, Los Alamos National Lab., 1992.

[14] B. R. Epstein, R. Hingorani, J. M. Shapiro, and M. Czigler, "Multispectral KLT-wavelet data compression for landsat thematic mapper images," in *Proc. 1992 Data Compress. Conf.*, (Snowbird, UT), pp. 200–205, IEEE Computer Soc., Mar. 1992.

[15] P. M. Kelly and J. M. White, "Preprocessing remotely sensed data for efficient analysis and classification," in *Proc. Conf. Knowledge-Based Systems Aerospace Industry*, vol. 1963 of *Proc. SPIE*, (Orlando, FL), Soc. Photo-Opt. Instrument. Engineers, Apr. 1993. To appear.