

Wavelet-Based Coding of Time-Varying Vector Fields of Ocean-Surface Winds

Li Hua, *Member, IEEE*, and James E. Fowler, *Senior Member, IEEE*

Abstract—Geoscience applications often produce sizable datasets which are vector-valued and increasingly in need of compression algorithms to reduce storage and transmission burdens, particularly when the data is time-varying. In this paper, several advanced interframe-compression techniques are extended from the traditional realm of natural video to the coding of time-varying vector fields. Although similar to natural video in some respects, time-varying vector-field sequences often possess complex temporal evolution of vector-valued features that are important to the analytic quality of the data yet defy the simple motion models widely employed for natural video. To improve coding performance, motion compensation with reduced resolution is proposed such that motion compensation is applied only at low spatial resolution while high-resolution information, for which the motion model fails, is intraframe coded with no temporal decorrelation. In empirical results on datasets of ocean-surface winds, this reduced-resolution motion-compensation technique results in significant performance improvement and greater feature preservation.

Keywords—vector-valued data compression, time-varying vector fields, vector wavelet transform, motion compensation

I. INTRODUCTION

Datasets arising in geoscience applications are increasingly of enormous size and often represent physical phenomena, such as atmospheric winds and ocean currents, that naturally take on vector values. The storage and transmission of such vector-valued data is a challenging problem which is further complicated in the case of time-varying vector-field sequences since many timesteps may need to be stored or transmitted. However, vector-valued datasets in general have received very little attention within the data-compression community.

When considering the issue of the compression of time-varying vector fields, one quickly notices many similarities to the traditional problem of the coding of natural video, which has been the focus of decades of research and the topic of several international standards (e.g., H.263 [1], MPEG-2 [2]). For example, in both natural video and vector-field sequences, there exists strong correlation spatially within each frame (spatial redundancy), as well as correlation between successive frames (temporal redundancy). Advanced video-compression techniques have been developed for natural video to eliminate these spatial and temporal redundancies to produce compression. These techniques include methods for intraframe coding, which exploit spatial correlations, as well as for interframe coding, which reduce temporal redundancy. For an in-depth overview of such intraframe and interframe techniques, see [3, 4].

L. Hua was with the Department of Electrical & Computer Engineering, Mississippi State, Starkville, MS, and the Visualization, Analysis, and Imaging Laboratory (VAIL) within the GeoResources Institute (GRI) at the Mississippi State ERC and is currently at Microsoft, Redmond, WA.

J. E. Fowler is with the Department of Electrical & Computer Engineering, Mississippi State, Starkville, MS, and the Visualization, Analysis, and Imaging Laboratory (VAIL) within the GeoResources Institute (GRI) at the Mississippi State ERC.

This work was funded in part by the National Science Foundation under the Large Data and Scientific Software Visualization Program, Grant No. ACI-9982344.

In our previous work [5, 6], we constructed an intraframe coder for static vector fields. In doing so, we designed a vector wavelet transform (VWT) to exploit the spatial correlation inherent in vector fields. We then coupled this VWT to vector-valued successive-approximation quantization [7] to produce an efficient, practical coder for static vector fields. In this paper, we extend this prior intraframe work by studying interframe redundancy-reduction methods for the coding of time-varying vector fields. We focus on the most widely employed form of temporal-redundancy reduction, motion estimation and compensation (ME/MC), by producing several ME/MC-based coders for vector-valued fields and experimentally evaluating their performance.

Although we expect that the resulting coders should be amenable to the compression of vector-field sequences arising in a number of geoscience applications, we conduct experimental evaluation specifically on datasets of rectangularly-gridded ocean-surface winds. Ocean-surface winds exhibit, as do other fluid-flow data, complex temporal evolutions of vector-valued features such as swirls and shear lines. Unfortunately, we find that the complex nature of these feature evolutions often exceeds the motion-compensating capability of the simple translational approaches to ME/MC commonly used for natural video. Consequently, in an effort to improve ME/MC performance for vector-field sequences, we introduce in this paper reduced-resolution motion-compensation (RRMC), a technique which uses a VWT to partition features in resolution such that low-resolution feature components are motion compensated as usual, whereas high-resolution components are merely intraframe coded with no temporal decorrelation. As a result, the limited abilities of the translational motion model are applied to only those feature components that can benefit from it—the components with low spatial resolution—while the model is not applied to the high-spatial-resolution components to which it is ill suited. Upon adoption of the proposed reduced-resolution approach to ME/MC, we observe a significant performance improvement for our ME/MC-based coders.

The remainder of this paper is organized as follows. In Sec. II, we first briefly overview the concepts of vector fields and VWTs and then consider the complex temporal evolution of vector features typical to time-varying vector fields. Next, in Sec. III, we extend several interframe coding schemes to the coding of vector-field sequences. In Sec. IV, we discuss the problem of the inefficiency of ME/MC methods in capturing vector-feature evolutions and propose the RRMC solution to remedy this problem. Finally, we present experimental results in Sec. V and make some concluding remarks in Sec. VI.

II. BACKGROUND

A. Vector Fields and Vector Wavelet Transforms

We define a vector-valued signal composed of N -dimensional vectors to be a collection of vectors $u[m] \in \mathcal{R}^N$ indexed by integer index m . A vector field is defined similarly such that $u[m, n] \in \mathcal{R}^N$ is indexed by both integer row and column indices m and n , respectively.

A VWT can be viewed as a generalization to vector signals of the common discrete wavelet transform (DWT) which transforms scalar-valued signals. In this generalization, the transform coefficients are vector-valued while the transform itself is implemented via a filtering operation using matrix-valued filters and vector-matrix convolution. Specifically, a VWT recursively decomposes vector signal $u[m]$ into vector-valued coefficient signals as

$$c_{j+1}[k] = \sum_m \tilde{H}[m - 2k]c_j[m], \quad (1)$$

$$d_{j+1}[k] = \sum_m \tilde{G}[m - 2k]c_j[m], \quad (2)$$

where c_j and d_j are the scaling and wavelet coefficient vector signals, respectively, at scale j with $c_0[m] = u[m]$; and \tilde{H} and \tilde{G} are matrix-valued analysis filters, each a sequence of $N \times N$ matrices. The corresponding synthesis operation is used to recursively reconstruct $u[m]$ from its coefficients; i.e.,

$$c_j[k] = \sum_m H^T[k - 2m]c_{j+1}[m] + \sum_m G^T[k - 2m]d_{j+1}[m], \quad (3)$$

where H and G are matrix-valued synthesis filters. The simplest approach to creating a VWT is to use scalar-valued filters from a traditional DWT to create diagonal VWT matrix filters; e.g., $H[m] = \text{diag}(h[m], h[m], \dots, h[m])$, where $h[m]$ is a scalar-valued DWT filter. In [5, 6], we designed more general, non-diagonal VWT filters which also have good signal-processing performance.

A 2D VWT for a vector field $u[m, n]$ can be implemented as is usual using a separable transform by first performing a 1D VWT as described above along each row of the field, and then repeating along each column. In this case, the VWT partitions the vector field into a baseband vector field consisting of information with low spatial resolution, and a collection of highpass subband fields with higher resolution.

B. Vector-Field Sequences

We define a vector-field sequence to be a collection of vector fields indexed by an integer time value t , i.e., $u[m, n, t]$. As we have stated above, the coding of time-varying vector fields has many similarities to the classical video-coding problem, and popular video-coding schemes can be directly extended to construct compression systems for sequences of vector fields. We propose a number of such direct extensions in Sec. III.

However, video signals and time-varying vector datasets have many significant differences. Most important of these is that, in video-coding applications, the salient characteristics to be preserved during compression are areas such as object edges and regions with high-intensity gradients which are important

to the perceptual quality of the video. On the other hand, vector datasets often have special vector-valued features—such as swirls, critical points, shock waves, and shear lines—which need to survive the compression since they are important to the analytic quality of the data. However, such features are characterized by phenomena more complex than edges or gradients. Furthermore, the evolutions of these features through the sequence are typically much more complex than in video sequences which often capture simpler phenomena such as rigid-body translations at slow speed. We illustrate the complexities associated with vector features in the next section by considering the special case of swirl features.

C. Swirl Features

In the ocean-surface-wind datasets we consider in experimental evaluations to come, swirl features are of particular interest. Although we know of no formal definition for a swirl vortex, in [8], swirling flow is considered to be any fluid motion that suggests rotation of fluid particles about a common center, either fixed or moving, and vortices are identified based on the eigenvalues and eigenvectors of the velocity-gradient tensor. Following this methodology, we locate regions of swirling flow in an ocean-surface-wind vector field by detecting areas with complex eigenvalues in the velocity-gradient-tensor matrix. This process results in a classification of each vector in a vector field as either “swirl” or “no swirl;” we call the resulting binary classification for the field a “swirl map.” Fig. 1 gives an example visualization of a vector field and its corresponding swirl map. Such swirl maps have been used extensively [9, 10] to detect and segment swirl features in computational fluid flows; below, we evaluate the performance of compression algorithms based on how well the swirl maps for a vector-field sequence are preserved. It has been observed [9] that the swirl map is somewhat prone to false positives since it can incorrectly classify merely locally curved flows as closed swirling flow due to the local nature of the velocity gradient tensor. As a consequence, avoiding the introduction of false positives is crucial to the swirl-preservation performance of compression systems.

Fig. 2 illustrates a complex temporal evolution that is typical of swirl features. In comparing the evolution of the swirl in Fig. 2(a) into the subsequent frame depicted in Fig. 2(b), we see that the swirl roughly maintains its “strength;” i.e., the vectors that constitute the swirl do not change much in magnitude. However, we see that not only does the swirl shift in spatial position, but the shape of the swirling region changes, becoming more elongated. The elongation of the swirl feature results in a shift in orientation of many of the flow vectors that constitute the swirl. Thus, the swirl exhibits not only simple translational motion, but also more complex evolutions involving the vector orientation of the flow about the swirl center.

III. CODING SYSTEMS FOR VECTOR-FIELD SEQUENCES

Popular methods of identifying and removing interframe redundancy in natural video sequences are based on the fundamental idea of reusing one or more previously coded frames to predict the current frame and coding the difference. Although such interframe redundancy-reduction techniques have a long history of use in the video-coding community, we know of no

prior efforts to employ them on time-varying vector fields. Thus, in this section, we extend several interframe coding systems to the coding of vector-field sequences.

A. Frame Difference (FD)

The simplest approach to interframe redundancy reduction is perhaps difference coding, the coding of the difference between successive frames. Specifically, in frame-difference (FD) coding, the current frame is predicted using the corresponding values in the immediately preceding frame, called the reference frame. The difference between the current frame and its prediction, rather than the current frame itself, is coded.

Our FD coding system for vector-field sequences is shown in Fig. 3. The reference frame is subtracted from the current frame to produce a residual frame which is then subjected to frame-by-frame intraframe coding. For intraframe coding, we use the vector-valued successive-approximation runlength (VSARL) coder developed in [5,6], in which we first apply a full-frame VWT to the entire residual frame. Subsequently, vector quantization (VQ) using the successive-approximation VQ (SAVQ) of [7] is applied to the VWT coefficient vectors. Finally, we finish with runlength coding of all insignificant vectors and arithmetic coding with multiple contexts [11]. In whole, the VSARL coder, which produces an embedded bitstream, is roughly an extension to vector data of a coder we developed recently for scalar-valued oceanographic imagery [12, 13].

B. Simple Block-Based Motion Compensation (BMC)

In real sources, there can exist substantial motion throughout a sequence. Difference coding is not usually effective when there is significant motion since only regions that remain stationary between the frames will be effectively coded through simple frame subtraction. For natural video signals, prediction errors due to motion can be significantly reduced by ME/MC.

The most widely used form of ME/MC, block-based MC (BMC), is based on the observation that complex motion of large objects over a short time interval can often be well-approximated by simple translations of small blocks. In BMC, the current frame is first segmented into square blocks, and then each block in the current frame is matched against a corresponding block in the reference frame within a search window. The offset of the best-match block from the current-frame block is the motion vector for the current-frame block [3].

Our simple BMC system for vector-field sequences is shown in Fig. 4. As in the FD system described above, the VSARL [5, 6] algorithm is employed as an intraframe coder on each motion-compensated residual frame. In this system, we set the block size for BMC to $M = 16$, and the window size to $\omega = 15$. For vector-field sequence $u[m, n, t]$, we use the mean absolute error (MAE) as the block-matching criterion,

$$\text{MAE}[i, j] = \frac{1}{M^2} \sum_{k=0}^{M-1} \sum_{l=0}^{M-1} \|u[m+k, n+l, t] - u[m+k+i, n+l+j, t-1]\|_1, \quad (4)$$

where $\|\cdot\|_1$ is the N -dimensional ℓ_1 norm, $[m, n]$ is the location

of the current block, and $[i, j]$ is the offset of the search block, $-\omega \leq i, j \leq \omega$.

Due to its simplicity and effectiveness, BMC has been widely adopted by most commercial, and all standard, video-compression systems. However, BMC is subject to the well-known drawback of blocking artifacts. Conventional block-based ME/MC is based on an implicit assumption that each block undergoes a uniform translational motion, which does not always occur. Blocking artifacts occur when neighboring blocks receive different motion vectors during the ME process, resulting in discontinuities along the block boundaries during MC. These discontinuities yield large energy in the high-frequency domain in the transform coefficients, thus the coding efficiency is significantly reduced [14]. For vector fields, blocking artifacts are especially problematic since a single vector feature is often segmented into different blocks by ME/MC, and the resulting block artifacts can destroy the very existence of a feature from an analytic perspective.

C. Overlapped-Block Motion Compensation (OBMC)

A popular method for reducing block artifacts is overlapped-block motion compensation (OBMC) [14] which provides a linear estimator of each pixel by repositioning overlapped blocks of pixels from the reference frame, each weighted by some smooth window. OBMC generalizes conventional BMC through a weighted prediction of the current block using multiple motion vectors drawn from neighboring blocks. For example, the OBMC method specified in H.263 [1] uses three predictions for each current-frame block—one using the current block's own motion vector and two using motion vectors from its adjacent blocks. The final approximation for the current block is a weighted sum of these three prediction values. For vector-field sequences, we implement OBMC using the same prediction and weighting procedure as in H.263 [1] applied to each vector component individually. Fig. 5 illustrates the resulting vector-field OBMC system.

IV. REDUCED-RESOLUTION MOTION COMPENSATION

The effectiveness of block-based MC as deployed in the systems described above ultimately depends on the extent to which the assumption of a simple translational block model holds. In vector-field sequences, feature motions are much more complex than in a natural video sequence, involving translational motion in conjunction with other complex evolutionary events. Under such conditions, neither BMC nor OBMC will not be able to match temporal evolutions in time-varying vector fields as efficiently as it does for natural video sources.

For example, recall the temporal evolution of the swirl feature that was illustrated in Fig. 2, in which we saw a swirl exhibit translational motion while maintaining swirl strength (vector magnitude). However, the swirling region underwent elongation causing a shift in orientation of many of the flow vectors that constituted the swirl. Although it is likely that BMC can correctly track the translational motion of this swirl, the residual vectors produced by BMC may still have relatively significant signal energy due to changes in vector orientation.

Fig. 2 illustrated the temporal evolution of a swirl feature in the original spatial domain of the vector field. If, instead, we ob-

serve the swirl in the domain of a VWT, we observe that most of the complexities associated with the temporal evolution of the swirl, namely, the elongation and resulting shift in vector orientation, are phenomena that occur strictly at high spatial resolution. Specifically, Figs. 6(a) and (b) give the VWT baseband fields of the frames from Figs. 2(a) and (b), respectively. We see that, at low spatial resolution, the swirl features in Figs. 6(a) and (b) are quite similar in strength, shape, and vector orientation although displaced one from the other spatially. We argue that BMC should be better able to compensate for this spatial displacement here in the baseband field rather than in the original vector fields since the majority of the complexities of the feature evolution are absent at low spatial resolution.

Thus, we propose the following approach for the use of MC in the coding of vector-field sequences. Specifically, we partition the vector field into high-resolution and low-resolution components. In the low-resolution component, the MC process operates as usual—translational motion for the low-resolution information in the frame is predicted and compensated, and the low-resolution residual is coded. On the other hand, the high-resolution information in the frame, corresponding largely to the complex evolutionary phenomena not captured adequately by the translational block motion model, is neither predicted nor compensated. Instead, the high-resolution information is fed directly into the intraframe coder.

This reduced-resolution approach to MC, RRMC, can be applied to each of the ME/MC-based systems considered above in Sec. III. Using the OBMC system as an example, the reduced-resolution OBMC (RR-OBMC) system is depicted in Fig. 7. A reduced-resolution of the simple BMC system (RR-BMC) can be constructed similarly.

The operation of the system of Fig. 7 is as follows. First the block-based ME process determines motion for each block of the current frame. Using OBMC (or BMC in the case of the RR-BMC system), a prediction of the current frame is assembled from blocks culled from the reference frame as indicated by the motion vectors resulting from the ME process. Once the entire predicted frame is assembled, it is partitioned into low-resolution and high-resolution subbands by a full-frame VWT. Subsequently, one or more of the high-resolution subbands of the predicted frame are set to zero. This “reduced-resolution” predicted frame is then subtracted from the VWT of the current frame, resulting in MC for only low-resolution information; the high-resolution subbands of the current frame, whose corresponding subbands in the reduced-resolution predicted frame have been zeroed, essentially pass into the intraframe coder unmodified.

In RRMC, we typically zero all subbands at all transform levels other than that with the lowest resolution. The efficacy of RRMC thus depends on the number of resolutions, J , used in the VWT. It is well known that, for an intraframe coder, increasing J typically results in increased rate-distortion performance, with the rate of increase dropping off sharply after the first several levels. However, for RRMC, if J is small, little resolution reduction is applied to the motion model, and little rate-distortion improvement is thus seen. On the other hand, if J is large, the size of the lower-resolution subbands is small, and, consequently, very few wavelet coefficients in the frame ac-

tually get motion-compensated. As a result, the performance of an RRMC system, which combines RRMC with an intraframe coder, increases with increasing J to a point, and increasing J beyond this optimal value results in decreased rate-distortion performance. In the next section, we present experimental results that show that the RRMC approach is capable of significant performance gains. Although the value of J leading to the best rate-distortion performance is likely to be somewhat data dependent, for the sea-wind data we use in the next section, we have determined that $J = 3$ results in the best RRMC performance. This optimal J was determined by exhaustively evaluating performance over the range $J = 2, \dots, 6$.

V. EXPERIMENTAL RESULTS

A. Datasets and Implementation Specifics

Experiments were performed on four different vector-field sequences (*seawind1*, *seawind2*, *seawind3*, and *seawind4*) each consisting of 64 frames of data obtained from the NASA/NOAA Seaflux system [15] which measures the global ocean-surface wind. These sequences consist of 128×128 vectors of dimension $N = 2$ on a rectangular grid in each frame, and the temporal resolution is 12 hours per frame. Vector-component values are 32-bit floating point numbers.

All coders use the diagonal CDF9-7 VWT described in [5, 6] and $J = 3$ levels of decomposition. The RRMC systems zero all highpass VWT subbands of the predicted frame in the highest two resolution levels. To evaluate system performance, we compute a rate and a fidelity measure. For rate, we use bits per vector (bpv) calculated as the total number of bits produced divided by the total number of vectors encoded over the course of the vector-field sequence; motion-vector information is included in bpv figures. As a fidelity measure, we consider both a signal-to-noise ratio (SNR), intended to give an indication of application-independent performance, as well as an application-specific measure based on the swirl-map classification described in Sec. II-C.

B. SNR Performance Results

For an application-independent fidelity measure, we compute an SNR as follows. Given a vector-field sequence $u[m, n, t]$, the mean square error (MSE) between field u and some approximation \hat{u} to u at time t is

$$\text{MSE}[t] = \frac{1}{M_1 M_2} \sum_{m=0}^{M_1-1} \sum_{n=0}^{M_2-1} \|u[m, n, t] - \hat{u}[m, n, t]\|^2, \quad (5)$$

where $\|\cdot\|$ denotes the usual N -dimensional Euclidean norm (ℓ_2 norm), and both u and \hat{u} have size $M_1 \times M_2$. The SNR is then defined as

$$\text{SNR}[t] = 10 \log_{10} \frac{\sigma^2[t]}{\text{MSE}[t]}, \quad (6)$$

where the variance $\sigma^2[t]$ is

$$\sigma^2[t] = \frac{1}{M_1 M_2} \sum_{m=0}^{M_1-1} \sum_{n=0}^{M_2-1} \|u[m, n, t] - \bar{u}[t]\|^2, \quad (7)$$

and the mean $\bar{u}[t]$ is

$$\bar{u}[t] = \frac{1}{M_1 M_2} \sum_{m=0}^{M_1-1} \sum_{n=0}^{M_2-1} u[m, n, t]. \quad (8)$$

The average SNR is $\text{SNR}[t]$ averaged over all frames t .

Average-SNR performance of each of the coders considered above for each of the four sea-wind datasets is tabulated in Table I for a rate of 2.0 bpv. We observe a somewhat unexpected result in that the simple FD coding system outperforms both the BMC and OBMC systems. From this result, we conclude that ME/MC in the BMC and OBMC systems is unable to reduce the magnitude of the residual vectors substantially over those produced by FD, such that any advantages in distortion performance gained due to ME/MC are lost in the transmission overhead needed to send motion information to the decoder.

However, we notice in Table I a significant improvement in SNR performance following adoption of the RRMC approach. That is, the RR-BMC and RR-OBMC systems achieve a significant performance improvement over their full-resolution counterparts. The RR-BMC system benefits from the larger improvement because the low-resolution RRMC filter also acts as a de-blocking filter, which reduces blocking artifacts in the predicted frame and smooths the residual field. As a consequence, the RR-BMC and RR-OBMC systems achieve roughly the same performance, and both represent a significant improvement over both the full-resolution BMC and OBMC systems as well as the simple FD system.

In order to simplify presentation of the remaining results, we focus on the FD, BMC, and RR-OBMC coders. Fig. 8 shows frame-by-frame SNR profiles for these coders, while Fig. 9 shows average SNR performance over a range of rates. We conclude from these figures that the observations drawn from Table I hold consistently throughout the vector-field sequence as well as over a range of rates.

C. Application-Specific Performance Results

As we have argued above, the ability of coders to preserve application-specific features is of paramount importance to the analytic quality of vector-valued datasets. However, the SNR results of the previous section do not directly address feature-preserving performance. As a consequence, in this section, we develop an application-specific performance measure in order to gauge feature preservation and evaluate our coders according to this measure.

Specifically, swirl features are of particular interest in fluid-flow datasets such as our sea-wind sequences, while the swirl map described in Sec. II-C provides a binary classification of each vector in a vector field according to the presence or absence of swirl at that location. We use swirl maps calculated for each vector field of an original vector-field sequence as “ground truth” and compare them to the corresponding swirl maps calculated on the reconstructed fields resulting from the compression systems under consideration. Our swirl-based fidelity measure is then the total number of vectors that are misclassified in the swirl maps for the reconstructed frames over the entire sequence. Results are tabulated in Table II for the RR-OBMC, FD, and BMC coders. We see that the FD and BMC coders

typically result in on the order of 8% and 20% greater misclassifications, respectively, than the RR-OBMC coder, indicating that the RRMC approach not only increases SNR performance but also achieves greater feature preservation for swirls.

Reconstructed vector fields are depicted in Figs. 10(a), 11(a), and 12(a) for Frame 10 of the *seawind1* sequence. Although it is difficult to visually distinguish these vector fields from one another and from the original field (Fig. 1(a)), a rather pronounced difference is observed in the corresponding swirl maps as shown in Figs. 10(b), 11(b), and 12(b). In particular, we see that, although all three coders do largely preserve the prominent swirl in the middle right of the field, the FD and BMC coders produce a sizable false-positive region incorrectly classified as swirl in the lower left of the field.

VI. CONCLUSIONS

In this paper, we investigated effective algorithms for the compression of time-varying vector-valued datasets. We invoked several established interframe techniques arising in the relatively mature field of video coding and extended them to the application of vector-valued data compression. Inspired by advanced interframe-coding approaches employed in modern coders for natural video, we considered several techniques for the removal of temporal redundancy in time-varying vector fields. Specifically, we examined the simple differencing of FD, the simple block-based ME/MC approach of BMC, and the artifact-removing OBMC technique. However, contrary to expectations, the performance of the ME/MC-based systems was not as efficient as that for natural video sources due to vector-valued features that evolve in time in a manner more complex than is usual in natural video. In consideration of this characteristic of vector data, we proposed a novel methodology for the ME/MC of the complex temporal evolutions common to vector-field sequences. In this RRMC paradigm, only low-resolution information is motion-compensated while high-resolution information, wherein a majority of the complexities of the feature evolutions take place, is simply intraframe-coded. In experimental evaluations of the proposed vector-field-sequence coders, the proposed RRMC approach to motion-compensation significantly improved not only SNR performance but also application-specific performance at preserving features of swirling flow.

As one final observation, we note that the RRMC approach proposed here bears some resemblance to the loop filters that have been used in the H.261 standard [16] and the upcoming H.264 standard [17]. However, whereas the H.261 and H.264 loop filters are a form of lowpass filtering intended specifically to remove blocking artifacts, our RRMC filter does this as well as partitions vector features according to resolution. As a consequence, improved performance is seen even when OBMC, which is much less prone to blocking artifacts than simple BMC, is used in our coders for time-varying vector fields.

ACKNOWLEDGMENT

The sea-wind data is from the NASA/NOAA sponsored data system Seaflux at JPL through the courtesy of W. T. Liu and W. Tang.

REFERENCES

- [1] ITU-T, *Video Coding for Low Bitrate Communication*, January 1998, ITU-T Recommendation H.263, Version 2.
- [2] ISO/IEC 13818-2, *Information Technology—Generic Coding of Moving Pictures and Associated Audio Information: Video*, 1995, MPEG-2 Video Coding Standard.
- [3] M. Ghanbari, *Video Coding: An Introduction to Standard Codecs*, The Institution of Electrical Engineers, London, United Kingdom, 1999.
- [4] R. J. Clarke, *Digital Compression of Still Images and Video*, Academic Press, London, 1995.
- [5] J. E. Fowler and L. Hua, "Wavelet transforms for vector fields using omnidirectionally balanced multiwavelets," *IEEE Transactions on Signal Processing*, vol. 50, no. 12, pp. 3018–3027, December 2002.
- [6] J. E. Fowler and L. Hua, "Omnidirectionally balanced multiwavelets for vector wavelet transforms," in *Proceedings of the IEEE Data Compression Conference*, J. A. Storer and M. Cohn, Eds., Snowbird, UT, April 2002, pp. 422–431.
- [7] E. A. B. da Silva, D. G. Sampson, and M. Ghanbari, "Super high definition image coding using wavelet vector quantization," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 6, no. 4, pp. 399–406, August 1996.
- [8] C. H. Berdahl and D. S. Thompson, "Eduction of swirling structure using the velocity gradient tensor," *AIAA Journal*, vol. 31, no. 1, pp. 97–103, January 1993.
- [9] D. S. Thompson, R. Machiraju, M. Jiang, J. S. Nair, G. Craciun, and S. S. D. Venkata, "Physics-based feature mining for large data exploration," *IEEE Computing in Science and Engineering*, vol. 4, no. 4, pp. 22–30, July–August 2002.
- [10] M. Jiang, R. Machiraju, and D. S. Thompson, "Geometric verification of swirling features in flow fields," in *Proceedings of IEEE Visualization*, Boston, MA, October 2002, pp. 307–314.
- [11] 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.
- [12] J. E. Fowler and D. N. Fox, "Embedded wavelet-based coding of three-dimensional oceanographic images with land masses," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 39, no. 2, pp. 284–290, February 2001.
- [13] J. E. Fowler and D. N. Fox, "Wavelet-based coding of three-dimensional oceanographic images around land masses," in *Proceedings of the International Conference on Image Processing*, Vancouver, Canada, September 2000, vol. 2, pp. 431–434.
- [14] S. Nogaki and M. Ohta, "An overlapped block motion compensation for high quality motion picture coding," in *Proceedings of the IEEE International Symposium on Circuits and Systems*, San Diego, CA, May 1992, vol. 1, pp. 184–187.
- [15] J. Graf, C. Sasaki, C. Winn, W. T. Liu, W. Tsai, M. Freilich, and D. Long, "NASA scatterometer experiment," *Acta Astronautica*, vol. 43, no. 7-8, pp. 397–407, October 1998.
- [16] ITU-T, *Video Coding for Audiovisual Services at $p \times 64$ kbit/s*, March 1993, ITU-T Recommendation H.261.
- [17] T. Wiegand, G. J. Sullivan, G. Bjøntegaard, and A. Luthra, "Overview of the H.264/AVC video coding standard," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 13, no. 7, pp. 560–576, July 2003.




Fig. 1. Frame 10 of *seawind1*. (a) A flow-arrow visualization of the vector field, (b) the corresponding swirl map (white indicates areas classified as "swirl," black is "no swirl").






Fig. 2. Evolution of a swirl feature (vortex core indicated with "★"). (a) Zoomed in view of a swirl feature in frame 10 of *seawind1*; (b) the corresponding region in Frame 11.



wavelets.

Li Hua received the B.S. and the M.S. degrees in electrical engineering in 1997 and 1999, respectively, from Nanjing University, Nanjing, China. In August 2003, she received the Ph.D. degree in electrical engineering from Mississippi State University, Starkville, MS. From 2000 to 2003, she was a research assistant in the Visualization, Analysis, and Imaging Laboratory (VAIL) within the GeoResources Institute (GRI) at the Mississippi State ERC. She is currently with Microsoft, Redmond, WA. Her research interests include image and signal processing, video coding, and



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 associate professor in the Department of Electrical & Computer Engineering at Mississippi State University in Starkville, MS. He is also a researcher in the Visualization, Analysis, and Imaging Laboratory (VAIL) within the GeoResources Institute (GRI) at the Mississippi State ERC. His research interests include image and video coding, data compression, and wavelets.




Fig. 3. The FD system. Δ indicates a delay in time by one frame and the corresponding frame buffer.




Fig. 4. The BMC system.




Fig. 5. The OBMC system.




Fig. 6. Evolution of a swirl feature in the baseband field of a VWT with 3 levels of decomposition (vortex core indicated with “★”). (a) Baseband of VWT of frame 10 of *seawind1*; (b) baseband of VWT of frame 11.




Fig. 7. The RR-OBMC system.

TABLE I
AVERAGE SNR AT RATE = 2.0 BPV.

Coder	SNR (dB)			
	<i>seawind1</i>	<i>seawind2</i>	<i>seawind3</i>	<i>seawind4</i>
FD	16.61	15.84	16.31	16.45
BMC	13.70	12.79	13.40	13.91
OBMC	15.64	14.81	15.38	15.71
RR-BMC	17.57	16.66	17.33	17.31
RR-OBMC	17.62	16.72	17.38	17.36




Fig. 8. Frame-by-frame SNR profiles for test sequence *seawind1* coded at 2.0 bpv.

TABLE II
NUMBER OF MISCLASSIFIED VECTORS OVER THE ENTIRE TEST SEQUENCE.
RATE = 4.0 BPV. PERCENT FIGURES INDICATE THE PERCENTAGE OF
ADDITIONAL MISCLASSIFICATION RELATIVE TO THE RR-OBMC CODER.

Dataset	RR-OBMC	FD	BMC
<i>seawind1</i>	34708	37644 +8.5%	44774 +29.0%
<i>seawind2</i>	33450	35652 +6.5%	37958 +13.5%
<i>seawind3</i>	32820	37120 +13.1%	43516 +32.6%
<i>seawind4</i>	28836	30094 +4.4%	33764 +17.1%




Fig. 9. SNR performance over a range of rates. SNR figures are averaged over all frames of the test sequence *seawind1*.




Fig. 10. Frame 10 of *seawind1* coded with RR-OBMC at 4.0 bpv. (a) Vector field, (b) swirl map.




Fig. 11. Frame 10 of *seawind1* coded with FD at 4.0 bpv. (a) Vector field, (b) swirl map.




Fig. 12. Frame 10 of *seawind1* coded with BMC at 4.0 bpv. (a) Vector field, (b) swirl map.