Wavelet Compression Technique for High-Resolution Global Model Data on an Icosahedral Grid

NING WANG, JIAN-WEN BAO, JIN-LUEN LEE, FANTHUNE MOENG, AND CLIFF MATSUMOTO

NOAA/Earth System Research Laboratory, and Cooperative Institute for Research in the Atmosphere, Colorado State University, Boulder, Colorado

(Manuscript received 21 November 2014, in final form 16 June 2015)

ABSTRACT

Modern Earth modeling systems often use high-resolution unstructured grids to discretize their horizontal domains. One of the major challenges in working with these high-resolution models is to efficiently transmit and store large volumes of model data for operational forecasts and for modeling research.

A newly developed compression technique is presented that significantly reduces the size of datasets produced by high-resolution global models that are discretized on an icosahedral grid. The compression technique is based on the wavelet transform together with a grid rearrangement algorithm and precision-controlled quantization technology. The grid rearrangement algorithm converts an icosahedral grid to a set of 10 rhombus grids that retain the spatial correlation of model data so that a three-dimensional wavelet transform can be effectively applied. The precision-controlled quantization scheme guarantees specified precision of compressed datasets.

The technique is applied to the output of a global weather prediction model, the Flow-Following, Finite-Volume Icosahedral Model (FIM) developed by NOAA's Earth System Research Laboratory. Experiments show that model data at 30-km resolution can be compressed up to 50:1 without noticeable visual differences; at specified precision requirements, the proposed compression technique achieves better compression compared to a state-of-the-art compression format [Gridded Binary (GRIB) with JPEG 2000 packing option]. In addition, model forecasts initialized with original and compressed initial conditions are compared and assessed. The assessment indicates that it is promising to use the technique to compress model data for those applications demanding high fidelity of compressed datasets.

1. Introduction

Rapid increases in computing power and steady advances in numerical simulation techniques have made it possible to create global high-resolution atmospheric models for operational use. Several atmospheric models with sub-10-km horizontal resolutions are currently under development and are expected to be put into operational use within a decade (Satoh et al. 2008; Jung et al. 2012; Skamarock et al. 2012; Lee and MacDonald 2009). One of the major challenges to the archival and distribution of numerical forecasts from these high-resolution global models is their large data sizes. Datasets produced by these models typically have sizes in tens to hundreds of gigabytes. Efficient transmission and storage of these datasets poses a practical and important problem for both operational and research communities.

Efforts have been made to compress high-resolution model data on Cartesian grids in two of the most widely used data formats for geoscience data—Network Common Data Form (netCDF) and Gridded Binary (GRIB). The latest netCDF format, version 4 (netCDF4), uses LZ77 lossless compression (Ziv and Lempel 1977); and the GRIB format allows users to combine a conservative rounding of field data with a selectable lossless compression (packing) scheme in its new, second, edition (GRIB2; Dey 2007). The LZ77 compression available in netCDF4 achieves moderate compression with no precision loss to the original datasets, while rounding plus a lossless data compression used in GRIB2 achieves more significant compression at the expense of some precision loss.

In this article, we present a new lossy wavelet data compression technique with precision control to compress

Corresponding author address: Ning Wang, NOAA/ESRL, R/GSD1, 325 Broadway, Boulder, CO 80305.
E-mail: ning.wang@noaa.gov

DOI: 10.1175/JTECH-D-14-00217.1

© 2015 American Meteorological Society

datasets generated by numerical models discretized on an icosahedral grid.

Wavelet transform–based data compression techniques developed in the past two decades have proven to be effective for image and video data compression (DeVore et al. 1992; Shapiro 1993; Said and Pearlman 1996; Tian and Wells 1996; Topiwala 1998). A typical wavelet compression scheme involves a wavelet transform, a corresponding quantization procedure, and an optional lossless entropy encoding. Through a careful selection of the algorithm and implementation of each component, wavelet data compression achieves superior performance over most other compression techniques (Strang and Nguyen 1998; Daubechies 1992; Donoho 1993). The technique has been extended and applied to general datasets on Cartesian lattices, particularly numerical datasets from scientific computations. One important extension to the lossy wavelet compression technique for scientific datasets is its ability to guarantee the precision of the reconstructed dataset, that is, to control the maximum error of compressed data at any individual data point. Since the late 1990s, the wavelet compression technique has been used in an Internet-based meteorological workstation to transmit satellite images and model products over communication channels of limited bandwidth (Wang and Madine 1998). The precision-controlled wavelet compression technique has been proposed to compress high-resolution regional model data for transmission and archival purposes (Wang and Brummer 2003; Lucero et al. 2004). Recently, the GRIB2 format added an option to encode Cartesian datasets with rounding plus a lossless wavelet compression (JPEG 2000), which is a type of precision-controlled wavelet data compression as well.

Modern high-resolution global gridpoint models are often based on unstructured grids, such as icosahedral or cubed-sphere grids, for computational efficiency and better numerical treatment of high-latitude regions (Majewski et al. 2002; Tomita et al. 2004; Putman and Lin 2009). To compress model data on an unstructured spherical grid, one would naturally consider using a spherical wavelet (Schröder and Sweldens 1995; Lessig and Fiume 2008) for the horizontal transform and a classic one-dimensional wavelet for the vertical transform. However, the numerical properties of the existing spherical wavelets constructed on the icosahedral grid are not ideal for data compression. Therefore, we propose an alternative approach to conduct the horizontal transform. We arrange the unstructured icosahedral grid into a set of 10 rhombus subgrids that are isomorphic to a Cartesian grid and conduct the wavelet transform in three-dimensional space for each subgrid, using a three-dimensional separable transform (Mallat 1989). Details concerning this choice are discussed in section 2.

Precision-controlled compression, with or without wavelet transform, provides adequate fidelity of compressed datasets for most meteorological applications, as evidenced by the popular use of the GRIB, edition 1 (GRIB1) format and the recent introduction of the GRIB2 format. Those applications include visualizations of satellite and model products at local scales and numerical computations that generate derived products. However, one important measure of the fidelity of compressed model datasets is their ability to reproduce numerical forecasts when the compressed datasets are used as initial model states. This is a more stringent criterion, because during numerical simulations, small differences at the initial time can be amplified and can potentially cause the numerical forecasts to differ noticeably, especially after an extended period of simulation. On the other hand, if consistent forecasts can be produced from both original and compressed initial conditions, and these forecasts show no noticeable discrepancy, then it can be further ensured that the compressed datasets maintain high fidelity and have an inconsequential loss of information content. Therefore, these compressed datasets can potentially be used in even the most fidelity-demanding numerical computations.

Preliminary numerical experiments show that the proposed data compression technique is effective for model data on icosahedral grids. It reduces data volume dramatically for forecast datasets used by visualization applications, since these applications can tolerate more precision loss, as long as no noticeable visual differences appear in the rendered products. The technique achieves significant compression with commonly used precision specifications for meteorological datasets, and it has a better compression performance compared to GRIB2 (with JPEG2000 packing option). When used in a global icosahedral grid model, the Flow-Following, Finite-Volume Icosahedral Model (FIM), the compressed initial conditions produce no noticeable visual differences in 120-h forecasts compared to the forecasts produced with the original uncompressed initial conditions. These experiments demonstrate that the proposed data compression technique provides a practical solution to the compression of high-resolution icosahedral grid data. They also indicate that the proposed data compression scheme offers a promising alternative to the currently adopted compression schemes that use lossless compression or rounding plus lossless compression.

The article is presented as follows. Section 2 describes the compression technique for icosahedral grid data, including discussions of the wavelet transform of model data on an icosahedral grid and the algorithm for precision-controlled

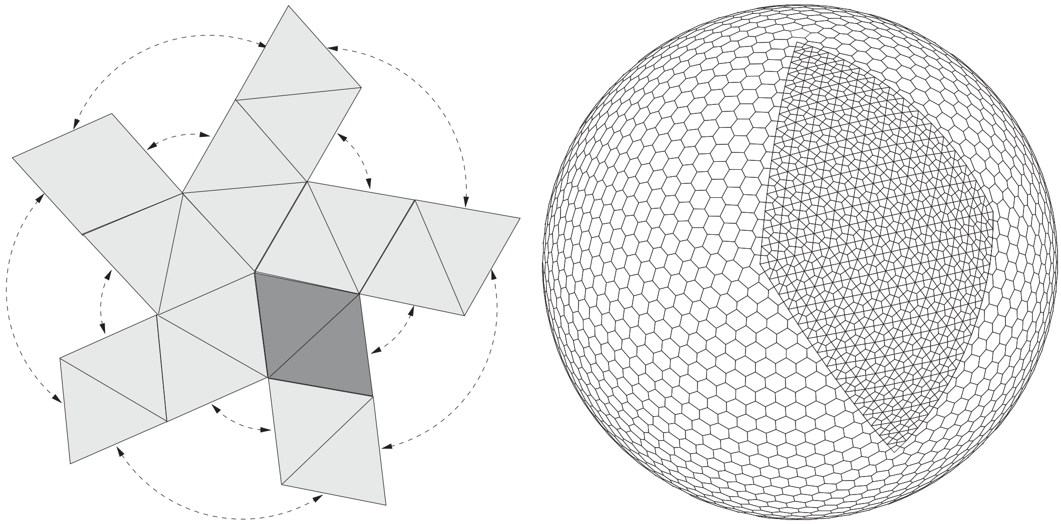

FIG. 1. Icosahedral grid. (left) The 10 rhombi flattened to the plane, where each pair of arrows points to two coincident geodesic edges on the sphere. (right) The icosahedral grid on the sphere. The rhombus with additional triangle mesh in the right plot corresponds to the rhombus highlighted by the darker shade in the left plot.

lossy compression. Section 3 presents a series of experiments that evaluate the compression performance of the technique. In section 4, two cases of weather forecasts initialized with the original and compressed initial conditions are examined and compared to assess the impact of data compression of initial states on the numerical simulations. In section 5, a summary of the compression technique and its experiment results is presented, followed by a brief discussion of future work.

2. Wavelet compression technique for icosahedral grid data

In this section, we present the compression technique developed for model data defined on an icosahedral grid.

a. Wavelet transform

A discrete wavelet transform (DWT) is the foundation for wavelet data compression. Like any transform-based data compression, a wavelet transform decorrelates spatially correlated data points to obtain better energy compaction in its coefficient representation, which leads to more efficient quantization, and therefore better data compression. For datasets defined on a Cartesian lattice, a separable wavelet transform is very effective and often used. The separable wavelet transform uses the wavelet defined on the real line and applies it repeatedly to each dimension (Mallat 1989). However, for datasets defined on an unstructured grid on the sphere, such as an icosahedral grid, one cannot use this approach directly.

A wavelet can also be constructed on a general manifold, such as a sphere using a numerical technique called the ‘‘lifting scheme’’ (Sweldens 1997; Schröder and Sweldens 1995). Discrete spherical wavelets have been constructed with this scheme and are used in computer graphics to code three-dimensional geometric meshes of given objects. The construction of spherical wavelets usually takes advantage of the fact that spherical meshes are often created with recursive subdivisions of faces of inscribing regular polygons. Taking icosahedral grids (Baumgardner and Frederickson 1985) for example, spherical wavelets of different scales can be defined and located at the grid points of different refinement levels. From this perspective, it seems that a spherical wavelet could be a natural choice to conduct wavelet transforms for icosahedral grid data. However, spherical wavelets constructed with the lifting scheme have several limitations for data compression. These spherical wavelets often lack some of the important mathematical properties for data compression—regularity, a reasonable number of vanishing moments, symmetry, and unconditional basis in $L^2(\mathcal{S}^2, d\omega)$ function space (Donoho 1993; Sweldens 1997; Lessig and Fiume 2008). Furthermore, to improve some of these properties, symmetry, for example, requires changing the gridpoint locations, which is not feasible because gridpoint locations are determined and optimized for model numerics (Heikes and Randall 1995; Tomita et al. 2001; Du and Ju 2005; Wang and Lee 2011). Finally, these spherical wavelets do not have an efficient and performance-proven quantization scheme for their coefficients.

On the other hand, compactly supported orthogonal or biorthogonal wavelets on the real line (and on multidimensional Cartesian space) have been well studied

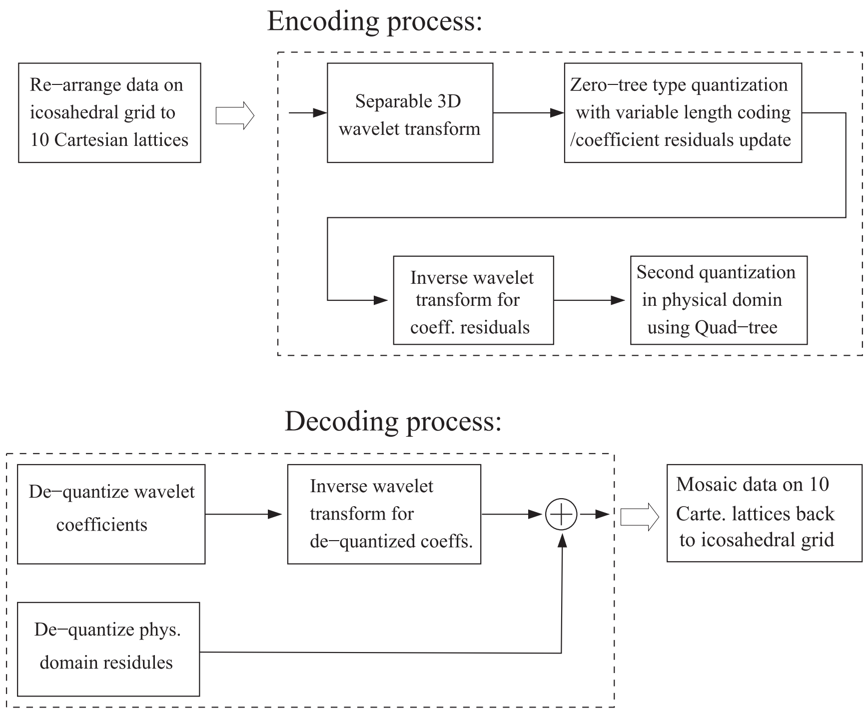

FIG. 2. Compression and decompression schemes.

theoretically and experimented numerically (Daubechies 1992; Antonini et al. 1992; Donoho 1993). The Cohen–Daubechies–Feaveau biorthogonal wavelet with (9, 7) taps (CDF97) (Cohen et al. 1992), for example, has proven to be an effective wavelet to use for separable wavelet transforms in data compression. It has most of the desired mathematical properties for data compression and has been used in data compression for two-dimensional imagery, three-dimensional video streams, and high-dimensional numerical datasets. In addition, there are several mature quantization schemes developed for wavelet coefficients on Cartesian lattices that can be adopted directly.

Therefore, we propose to rearrange the icosahedral grid points into a set of 10 rhombus subgrids on the sphere and to apply a separable three-dimensional wavelet transform to the grid data on these grids using the CDF97 wavelet. It is true that grid points within each subgrid are not perfectly regular and that grid lines are skewed (not orthogonal) compared to a regular lattice on a two-dimensional plane. Nonetheless, these subgrids are sufficiently regular in terms of the geodesic distances between adjacent grid points (section 2b). Therefore, a separable wavelet transform can still be effectively applied to the data on these subgrids.

b. *Grid rearrangement*

An icosahedral grid with l levels of subdivision refinements (commonly referred to as "grid level l icosahedral grid," or "Gl icosahedral grid," for short) has $10 \times 2^{2l} + 2$ grid points (Baumgardner and Frederickson 1985).[1] Excluding the two pole grid points, one can partition the entire grid into exactly 10 subgrids, each containing 2^{2l} grid points. The natural regular subgrids of an icosahedral grid are its initial 20 spherical triangular grids, which can be combined into 10 spherical rhombus grids (Fig. 1). Since each spherical rhombus grid mesh covers an adequately large area (one-tenth of the sphere), it could be a good partition for a two-dimensional separable wavelet transform.

Figure 1 illustrates how an icosahedral grid is partitioned into 10 rhombi. Each rhombus mesh, disregarding the edges that are parallel to the triangle side shared by two initial triangles, is isomorphic to a planar graph regular square lattice. Thus, it can be treated as a Cartesian grid on the plane for data compression purposes. In addition, these lattices are quite regular and rather uniform. In fact, depending on the construction algorithm for the icosahedral grid, the ratios between the maximum and minimum geodesic distances between adjacent grid points in the grid is bound by a constant

[1] Here a bisection subdivision for all refinement levels is assumed. When an arbitrary m section is used during subdivision refinement, the formula should be modified accordingly.

TABLE 1. Mean L^2 Errors for variables at various compression ratios.

Compression ratio	Temp (°C)	Wind u (m s^{-1})	Wind v (m s^{-1})	RH (%)
50:1	0.074 10	0.104 09	0.124 74	0.532 35
30:1	0.043 75	0.054 77	0.069 92	0.305 23
20:1	0.025 11	0.031 95	0.040 69	0.167 16
10:1	0.008 86	0.009 45	0.011 30	0.038 36

ranging from 1.17 to 1.19 (Wang and Lee 2011), which is a desirable property for separable wavelet transforms.

In the proposed compression scheme, we arrange the entire icosahedral grid into 10 rhombus lattices of 2^{2l} grid points, to which a traditional wavelet compression for Cartesian datasets is applied. For the remaining two polar grid points, a lossless data compression is used to encode them separately.

c. *Precision control of the compressed dataset*

In lossy data compression, the fidelity of a compressed dataset is usually measured with the L^p norm in an error vector space:

$$\|\mathbf{E}\|_p = \left[\sum_{i=1}^{n} |f(\mathbf{x}_i) - \tilde{f}(\mathbf{x}_i)|^p\right]^{1/p},$$

with

$$\|\mathbf{E}\|_\infty = \max_{1\le i\le n} |f(\mathbf{x}_i) - \tilde{f}(\mathbf{x}_i)|,$$

where $f(\mathbf{x}_i)$ is the function value at discretized location $\mathbf{x}_i$, $\tilde{f}(\mathbf{x}_i)$ is the lossy compressed function value, and n is the number of discrete data points. The norm $\|\mathbf{E}\|_p$ is referred to as the p norm error metric.

Commonly, an L^2 norm error metric $\|\mathbf{E}\|_2$ is used in image data compression. It measures an average or overall fidelity of the compressed image data. Physically, it also measures the degree of energy conservation. In signal processing $\|\mathbf{E}\|_2$ is directly related to the peak signal-to-noise ratio (PSNR), which is commonly used to indicate the quality of a signal. For scientific data, in addition to $\|\mathbf{E}\|_2$, we would like to retain a specified precision for every data point, that is, to have a bound on the maximum error for compressed datasets. Thus, the L^∞ norm error metric $\|\mathbf{E}\|_\infty$ is used as well to ensure the fidelity of the compressed scientific dataset.

Lossy data compression schemes based on (bi)orthogonal transform are naturally designed to minimize the distortion (error) in terms of $\|\mathbf{E}\|_2$. After transform, a critical procedure called "quantization" is applied to the coefficients to reduce bitrate and achieve data compression. This procedure introduces quantization error to the compressed datasets. To control the physical domain $\|\mathbf{E}\|_2$ in the transformed domain is algorithmically straightforward, because $\|\mathbf{E}\|_2$ is invariant under an orthogonal transform (and it is practically near invariant under a biorthogonal transform), and $\|\mathbf{E}\|_2$ is additive in both physical and transformed domains. For a given function f,

$$\sum_{i=1}^{n} [f(\mathbf{x}_i) - \tilde{f}_q(\mathbf{x}_i)]^2 = C_{\mathcal{R}} \sum_{i=1}^{n} [\hat{f}(\omega_i) - \widetilde{\hat{f}}_q(\omega_i)]^2, \qquad (1)$$

and

$$A \le C_{\mathcal{R}} \le B,$$

where $\hat{f}$ denotes the transform of function f, ω_i denotes the wavelet index to the scale-spatial lattice, $\tilde{f}_q$ denotes approximation of function f by quantization, and A, B are frame bounds (Riesz bounds). For orthogonal transform we have the equality $A = C_{\mathcal{R}} = B = 1.0$, and for biorthogonal transform using popular CDF wavelets, we have a $C_{\mathcal{R}} \in [A, B]$ that is usually close to 1.0. (de Saint-Martin et al. 1999; Cohen et al. 1992).

Therefore, in L^2 norm, we can directly measure and reduce the distortion of $\tilde{f}(\mathbf{x})$ in the transformed domain by measuring and reducing the distortion of $\widetilde{\hat{f}}(\omega)$. However, assessing and controlling distortion in $\|\mathbf{E}\|_\infty$, in the transformed domain, is not so simple. The error metric $\|\mathbf{E}\|_\infty$ is not invariant under (bi)orthogonal transform, and it is not additive in either domain. A few attempts have been made to solve the problem theoretically or practically (Karray et al. 1998; Marpe et al. 2000; Yea and Pearlman 2006). One can derive the minimum size of the quantization step in the transformed domain for a particular wavelet such that a specified $\|\mathbf{E}\|_\infty$ is satisfied. The problem of that approach is "overcoding," meaning that compression uses a much higher bitrate than necessary due to the conservative size of the derived minimum quantization step. The difficulty in applying quantization in the transformed domain alone leads to a two-domain (transform domain plus physical domain) quantization approach (Ansari et al. 1998; Yea and Pearlman 2006). The idea of the approach is to take advantage of the quantization efficiency in both transformed and physical domains to achieve a better bitrate for a given $\|\mathbf{E}\|_\infty$ requirement. Research on two-domain quantization been has focused on the selection of quantization schemes and bit allocation algorithms for the two domains. For the data compression task for high-resolution model data, an important consideration in algorithmic design is the computational efficiency.

We propose to adopt a simple but efficient procedure, previously proposed in Wang and Brummer (2003) for Cartesian grid data, to compress icosahedral grid data with a specified $\|\mathbf{E}\|_\infty$ bound. This procedure is a simple two-domain quantization scheme for floating point datasets, without bit allocation optimization. First, an empirical minimum quantization step size that indirectly specifies the mean L^2 norm error bound is given, and a classic scalar quantization of coefficients is performed in the transformed domain; the magnitude of each coefficient is reduced and updated as the quantization proceeds. When this transformed domain quantization finishes, the coefficient residuals are inverse transformed back to the physical domain. A second quantization is conducted on these physical domain residuals, and the quantization process continues until no residual exceeds the specified $\|\mathbf{E}\|_\infty$ bound. The minimum quantization step size for each meteorological parameter, in transformed domain, is determined empirically from a statistical estimate of the relation between the sizes of the quantization steps in the two domains.

Note in the above-mentioned procedure, the errors in the physical domain are not obtained through the computation of $f(\mathbf{x}) - \tilde{f}_q(\mathbf{x})$. Rather, they are obtained by applying the inverse wavelet transform to the coefficient residual $\hat{f}(\omega) - \tilde{\hat{f}}_q(\omega)$. Since the wavelet transform (and its inverse transform) is linear, we have

$$\begin{aligned} f(\mathbf{x}) - \tilde{f}_q(\mathbf{x}) &= \mathcal{W}^{-1}[\hat{f}(\omega)] - \mathcal{W}^{-1}[\tilde{\hat{f}}_q(\omega)] \\ &= \mathcal{W}^{-1}[\hat{f}(\omega) - \tilde{\hat{f}}_q(\omega)], \end{aligned} \tag{2}$$

where $\mathcal{W}^{-1}[\,.\,]$ is the inverse wavelet transform. This is a more computationally efficient way to compute the physical domain residual errors. It saves the memory space required to keep the original data $f(\mathbf{x})$, and it saves the CPU time required to decompress and create $\tilde{f}_q(\mathbf{x})$.

d. *Description of the entire compression scheme*

Figure 2 illustrates the main components of the entire compression and decompression processes. In the following paragraphs we describe the two processes and the implementation of their main components.

The compression process starts with a grid rearrangement procedure that partitions an icosahedral grid into 10 rhombi and two individual points. Two columns of data points at the poles are losslessly compressed, and this trivial step is omitted from the figure. For the three-dimensional datasets on 10 rhombus grids, we apply the following compression procedure to each of them.

First, an in-place, separable wavelet transform using the CDF97 wavelet is applied to the three-dimensional dataset. Then a two-domain quantization procedure, described in section 2c, is conducted. An enhanced zerotree-based quantization (Shapiro 1993; Said and Pearlman 1996) is used to quantize the coefficients in the transformed domain. This quantization approximates the function $\hat{f}(\omega)$ and reduces both L^2 and L^∞ norm errors according to the specified minimum quantization step size. The coefficient residuals are saved in the place of the coefficients at the end of transformed domain quantization. After an in-place inverse wavelet transform is applied to the coefficient residuals, a quadtree (Finkel and Bentley 1974)-based quantization is applied to the residuals in the physical domain to further reduce the L^∞ norm error, until a specified precision requirement is satisfied. For computational efficiency, we integrated an efficient variable length coding scheme into the quantization procedure, which allows us to skip the computationally intensive entropy encoding step.

As a result of the efficient implementation of the encoder components described above (in-place wavelet transform, efficient physical domain residual computation, and the integrated variable length coding), the proposed encoder runs efficiently for high-resolution icosahedral grid datasets. For example, it takes less than 80 s to compress a model variable on a G9 (15 km) icosahedral grid (~2.6 million grid points) of 64 vertical levels on an i7 processor Linux machine that translates to an encoder throughput of ~8.0 MiB s^{-1} (or ~8.4 MB s^{-1}). In addition, the memory-conserving design and implementation also makes it possible to create a multithreaded version of the software that encodes multiple rhombi in parallell. Depending on hardware configuration, a multithreaded compression software can potentially increase the present throughput significantly.

As mentioned in the previous subsection, the encoding scheme does not include any rate-distortion optimization procedure, which finds the optimal bit allocation for the two-domain quantization procedure. This optimization process is avoided due to its high computational expense. In anticipation of an ever-increasing data volume for high-resolution global models, we decided not to conduct this optimization in the compression process. Instead, empirical thresholds are used to conduct bit allocation for quantizations in the two domains during the encoding process.

The decompression process is largely symmetric to the compression one. For each rhombus grid, first a zerotree dequantization is conducted to restore the quantized wavelet coefficients and an inverse wavelet transform is applied to finish wavelet decompression of the dataset; then a quadtree dequantization is conducted to restore the physical domain residuals and the residuals are added back to the data decompressed in the first step.

Finally, we mosaic the 10 decompressed rhombus datasets back to the icosahedral grid, together with

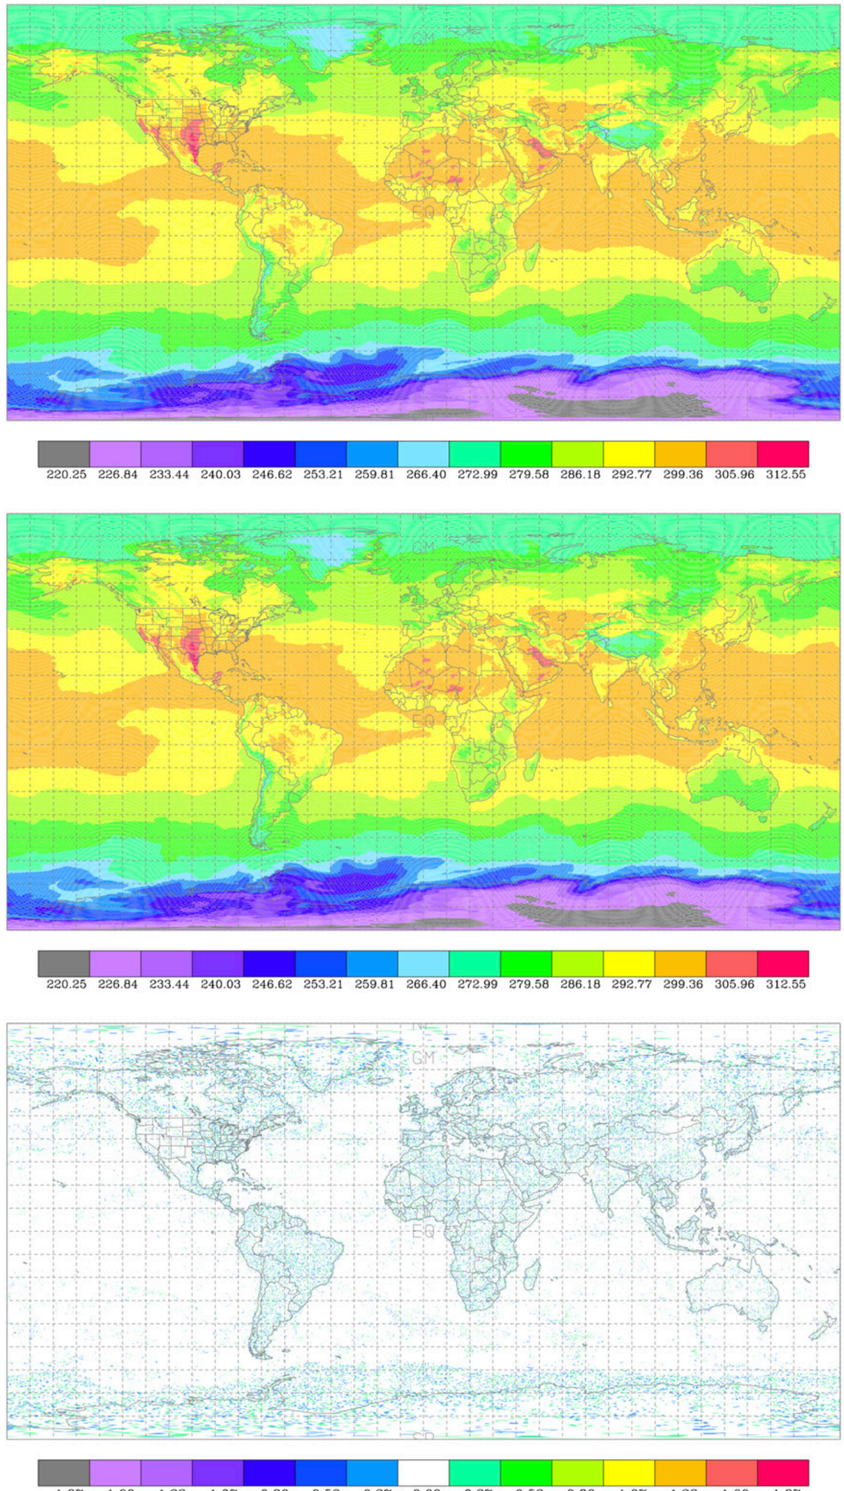

FIG. 3. (top) Original temperature field at surface (K), (middle) 50:1 compressed same field, and (bottom) difference image.

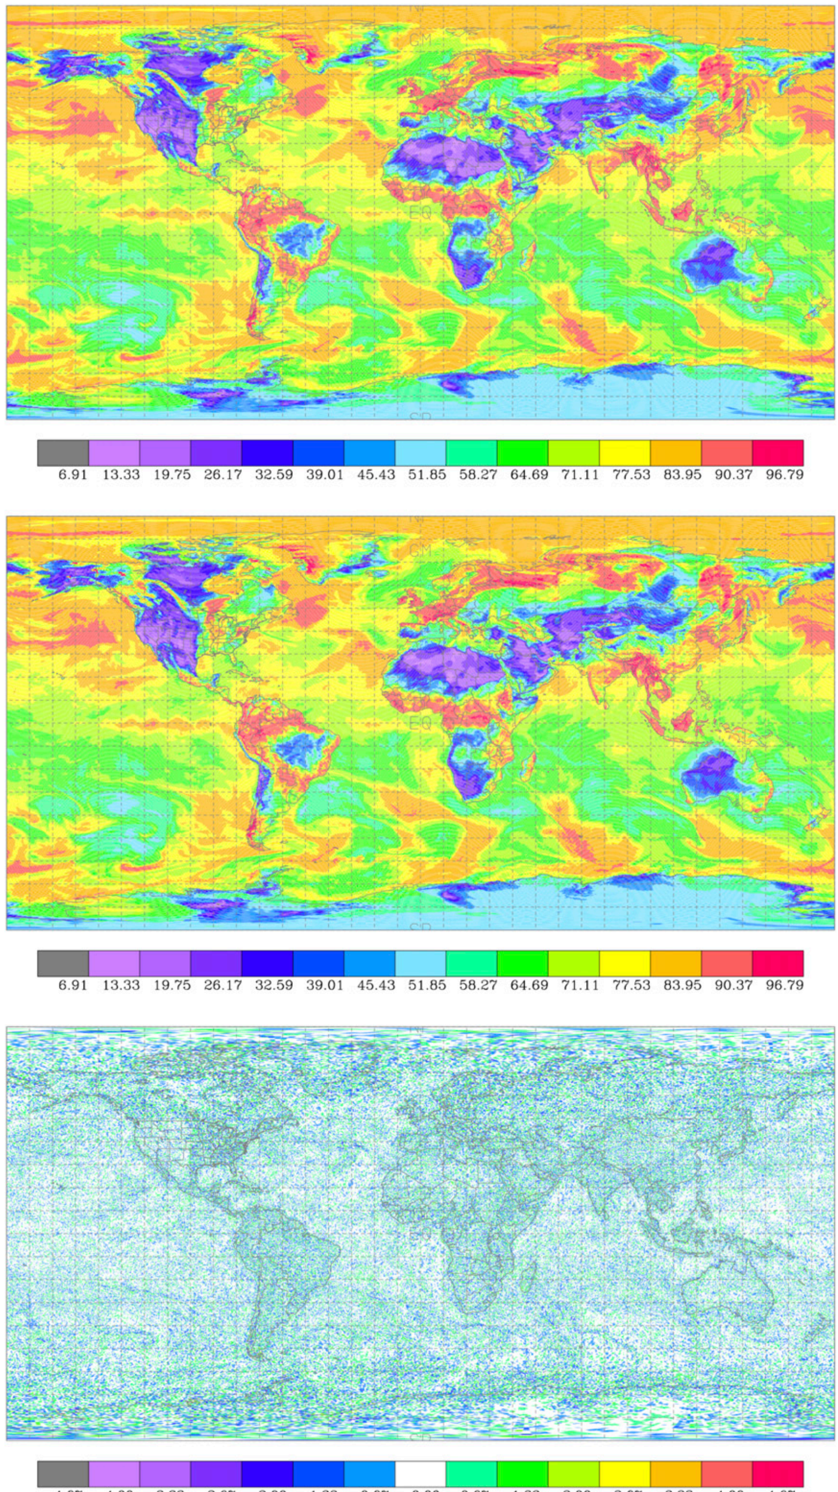

FIG. 4. (top) Original relative humidity field at surface (%), (middle) 50:1 compressed same field, and (bottom) difference image.

TABLE 2. Compression ratios for variables at specified precisions (L^∞ errors).

Precision	Temp (°C)	Wind u ($\mathrm{m\,s^{-1}}$)	Wind v ($\mathrm{m\,s^{-1}}$)	RH (%)
0.0625	15.1:1	14.5:1	12.1:1	7.9:1
0.125	22.9:1	21.3:1	17.2:1	9.7:1
0.25	37.2:1	33.2:1	25.8:1	12.6:1
0.5	73.9:1	55.2:1	42.6:1	18.6:1
1.0	137.6:1	96.5:1	72.2:1	28.1:1

TABLE 3. Mean L^2 errors for variables at specified L^∞ errors.

L^∞ errors	Temp (°C)	Wind u ($\mathrm{m\,s^{-1}}$)	Wind v ($\mathrm{m\,s^{-1}}$)	RH (%)
0.0625	0.017 96	0.019 16	0.015 86	0.017 90
0.125	0.032 92	0.035 43	0.030 04	0.036 25
0.25	0.057 62	0.063 46	0.054 99	0.072 43
0.5	0.108 60	0.108 93	0.102 45	0.159 43
1.0	0.173 91	0.178 30	0.171 60	0.296 39

losslessly decompressed data at the two pole points, to reconstruct the entire dataset.

The decompression process takes a slightly less amount of time than the compression process. The compression and decompression software are implemented in C++ and are compiled with GNU g++ compiler, version 4.2, with an optimization flag −O2.

3. Data compression assessment

We apply the compression technique described above to the model output generated by the FIM weather forecast model. FIM is a flow-following, finite-volume global model horizontally discretized on an icosahedral grid. It was developed by NOAA's Earth System Research Laboratory (ESRL) and is described in Lee and MacDonald (2009), Bleck et al. (2010), Lee et al. (2010), and Bleck et al. (2015). FIM is currently running daily at ESRL at a 30-km horizontal resolution and 64 vertical levels. During hurricane season, FIM is also run at higher horizontal resolutions, up to 10 km in horizontal spacing.

In this data compression experiment, we use the 30-km datasets on an icosahedral grid at grid level $l = 8$.

Three experiments are conducted to assess the performance of data compression: 1) evaluating overall fidelity of the compressed datasets at various compression ratios, 2) assessing compression ratios for given precision requirements, and 3) comparing the compression performance of the proposed method to that of the GRIB2 (with JPEG200 packing option) at the same given precisions.

The model variables involved in the experiments are temperature t; wind components u, v; and relative humidity r.

a. *Fidelity of the compressed datasets at given compression ratios*

First, for given compression ratios, we assess the quality of the compressed datasets using the mean L^2 norm error metric. For data visualization purposes, the mean L^2 error gives a good estimate of the overall fidelity of compressed datasets.

Table 1 lists the mean L^2 norm error for various datasets at different compression ratios. As expected, at the same compression ratio, smoother fields have smaller errors. In general, the mean L^2 norm errors are small for all variables at significant compressions.

Higher overall fidelity is required to display global model data on high-resolution large size displays. To verify and confirm the fidelity of the compressed datasets, we display the compressed FIM datasets on a special global spherical display—Science on a Sphere (SOS; http://sos.noaa.gov), which has a display area close to 15 000 in.^{2}. On the spherical display, we can compare each cell for any differences.

All images in the following presentation are rendered directly from icosahedral grid data to a cylindrical equidistant projection with no interpolations involved. The Voronoi cell mesh is not drawn, since at this high resolution (655 362 cells over the globe), the mesh will be too dense to allow visualization of the data beneath it.

Figures 3 and 4 show images of surface temperature t and relative humidity r before and after compression and their differences. These images are what an SOS system uses to project global images to its spherical display.

From these images, we observe that, even after significant compressions (50:1 compression ratio), model

TABLE 4. Compression ratios and mean L^2 errors of GRIB2 and the proposed method ($\|E\|_\infty = 0.05$).

Compression scheme	Temp (°C)	Wind u ($\mathrm{m\,s^{-1}}$)	Wind v ($\mathrm{m\,s^{-1}}$)	RH (%)
		Compression ratio		
Proposed method	13.55:1	13.01:1	11.84:1	7.82:1
GRIB2 JPEG2000	10.12:1	8.52:1	8.38:1	7.05:1
		Mean L^2 error		
Proposed method	0.013 576	0.015 814	0.015 507	0.019 193
GRIB2 JPEG2000	0.028 864	0.028 866	0.028 867	0.026 520

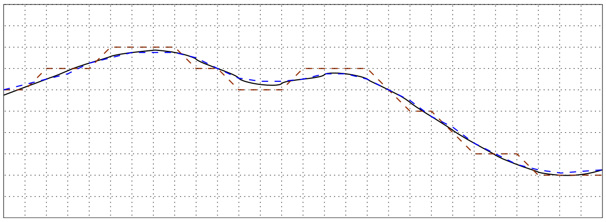

FIG. 5. Illustration of reconstructed functions with different compression schemes. Black solid curve represents the original function, blue dashed curve represents the two-domain quantization scheme, and brown dashed curve represents the rounding plus lossless coding scheme.

visualization products show little visual differences. Examining the difference plots, we notice that differences are mostly in the form of low-magnitude noise. These compressed datasets retain adequate fidelity for many applications, such as visualization of numerical models on meteorological workstations and dissemination of severe weather information to the public via the Internet, to mobile devices, and to the television network.

b. *Compression performance with guaranteed precisions*

For most datasets that need to be further processed and diagnosed numerically, it is required that the compressed data maintain certain precisions. The proposed compression technique can encode a dataset to a specified precision; that is, no grid data within the dataset will have a compression error that exceeds a specified precision threshold value.

Table 2 shows the compression ratios for different model variables at various specified precisions (L^∞ errors).

The mean L^2 errors are usually much smaller compared with L^∞ errors, especially for a compression scheme using two-domain (transform and physical domains) quantizations. Table 3 shows the mean L^2 errors of the datasets after precision-controlled compression.

Overall, significant compressions are achieved with reasonable precision requirements, especially for those variables that are smooth and stationary in space. Even for the variables that are less smooth and significantly less stationary, such as relative humidity, good compressions are achieved with some practical precision requirements.

c. *Comparison to GRIB2 compression*

GRIB is a compact data format commonly used in meteorology to store and transmit weather data. Based on its first edition (GRIB1), a newer version (GRIB2) was introduced by the end of last century. GRIB2 allows a more general and flexible format to encode metadata and, more importantly, it allows several modern compression schemes to be used in packing (lossless encoding) precision-specified binary data.

The most advanced packing option available in the GRIB2 format is the lossless wavelet compression in JPEG2000. In the following, we compare the compression performance of the proposed compression scheme to that of the GRIB2 format with JPEG2000 packing scheme.

Since GRIB2 is not available for unstructured grid, we first interpolate the icosahedral grid data to a global Cartesian grid with similar 30-km resolution, then apply GRIB encoding to the data. To be fair, any errors caused by interpolation will not be accounted for in the comparison.

Table 4 compares the compression ratios and mean L^2 errors of the proposed compression method and GRIB2 at the same specified precisions. This brief comparison shows that the proposed method achieves better compression, more noticeable for temperature and wind components. Table 4 also shows that the proposed method that uses the two-domain quantization scheme has smaller mean L^2 errors compared to GRIB2compression, which uses the rounding plus lossless wavelet compression scheme (JPEG2000), at the same specified L^∞ errors (precision). This result is expected and could be explained statistically. Assuming uniform distribution of function values within each rounding interval, it is straightforward to estimate the mean L^2 error for the rounding plus lossless compression scheme, by the following mathematical expectation:

$$\left(\frac{1}{0.5\tau}\int_0^{0.5\tau} x^2\,dx\right)^{1/2} = \frac{\sqrt{3}}{6}\tau \doteq 0.288\,675\tau, \qquad (3)$$

where τ is the rounding interval, which equals twice the L^∞ error. For the experiment, $\tau = 0.1$, and thus the expected mean L^2 error is approximately 0.028 867 5. The numerical results for t, u, v are close to this theoretical value. The mean L^2 error for r is a little smaller, because values of relative humidity at the top of the atmosphere are near zero, which does not follow the assumed uniform distribution of variable values within each rounding interval.

The two-domain quantization scheme also produces a smoother reconstructed function compared to the rounding plus lossless scheme. Figure 5 gives a hypothetical example of two reconstructed functions produced by the two-domain quantization scheme (blue dashed curve) and by the rounding plus lossless coding scheme (brown dashed curve), against the original function (black solid curve). The background grid mesh indicates the floors and ceilings for the rounding operation.

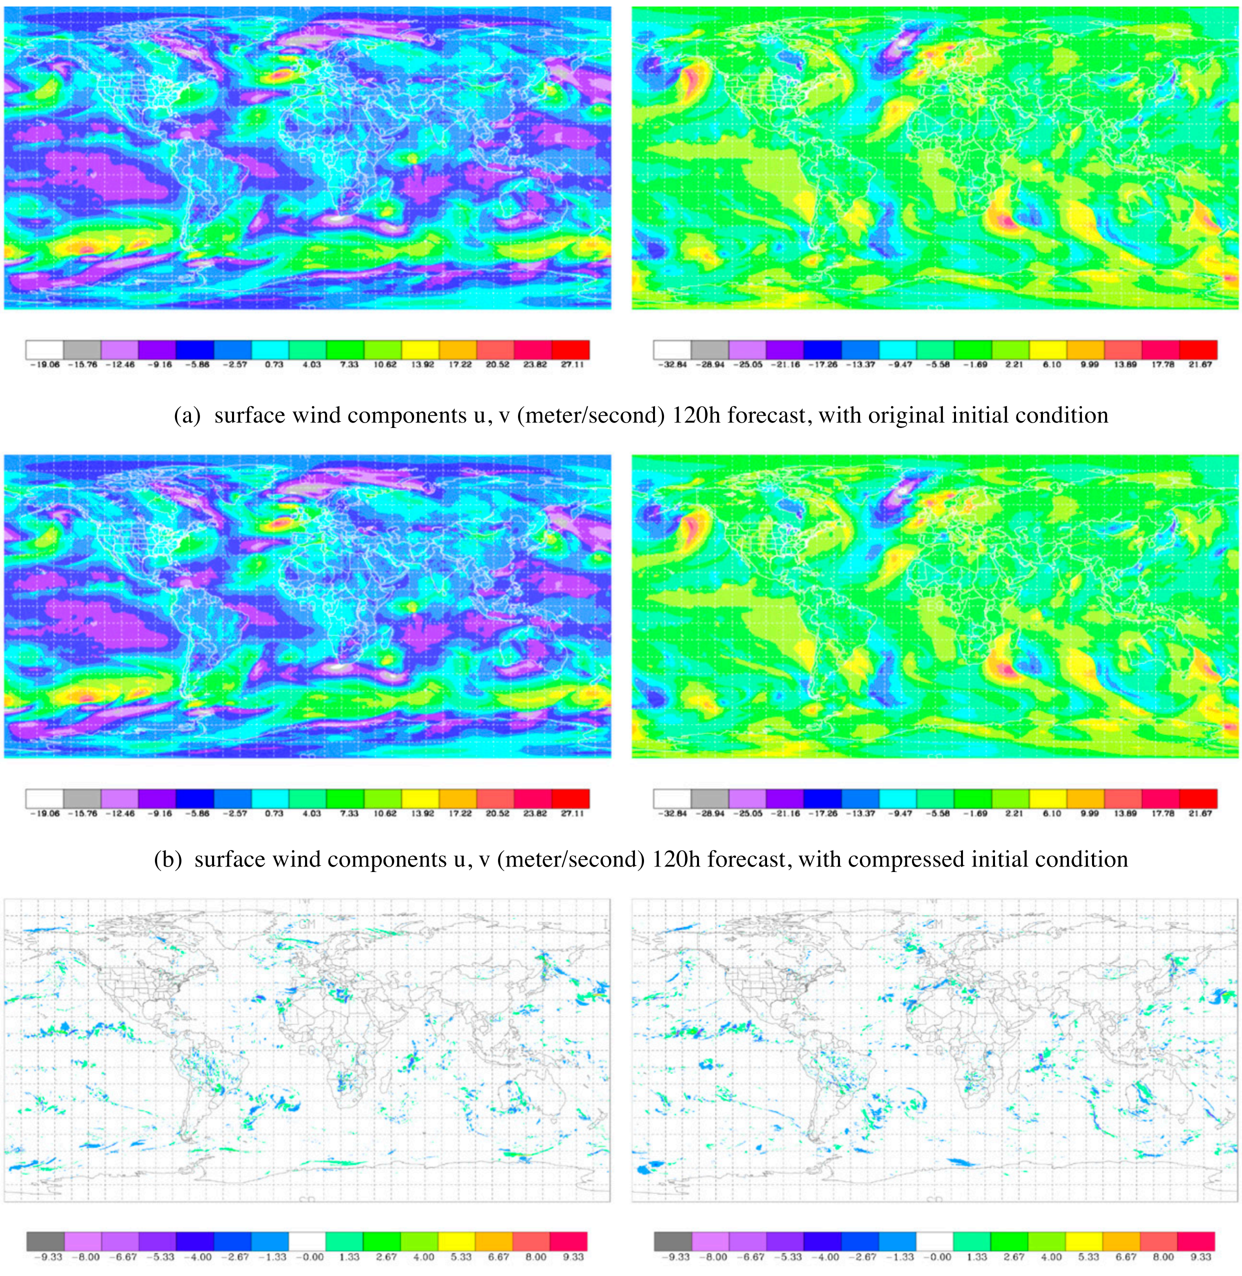

(c) difference fields of u, v (meter/second)

FIG. 6. Comparison of 120-h forecasts initialized with original and compressed kinetic fields. (top) 120-h forecast of u, v with original initial condition; (middle) as in (top), but with compressed initial condition; and (bottom) the difference images of the two forecasts. FIM runs at 30-km resolution, with the initial condition at 0000 UTC 28 Oct 2012 (Hurricane Sandy case).

The example illustrates the difference between the reconstructed functions by the two compression schemes.

4. Model runs with original and compressed initial conditions

Model initial condition datasets are often archived or transmitted to different computer systems to run numerical models for operational or research purposes. To compress these datasets, it is important to understand how data compression impacts the fidelity of their simulation results. In this experiment, we select two cases, one in the fall of 2012, during the life span of Hurricane Sandy, and one in March 2014, during a spring snowstorm.

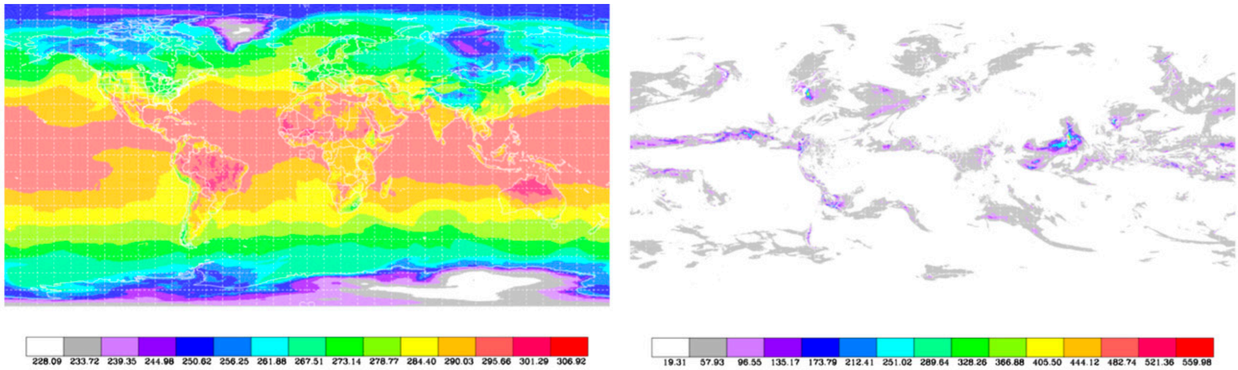

(a) surface temperature t (degree K) and precipitation (millimeter) 120h forecast, with original initial condition

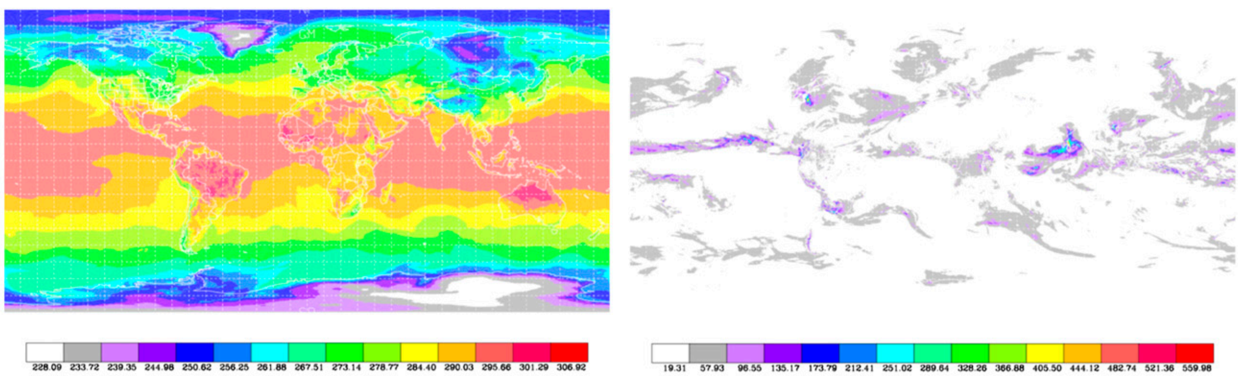

(b) surface temperature t (degree K) and precipitation (millimeter) 120h forecast, with compressed initial condition

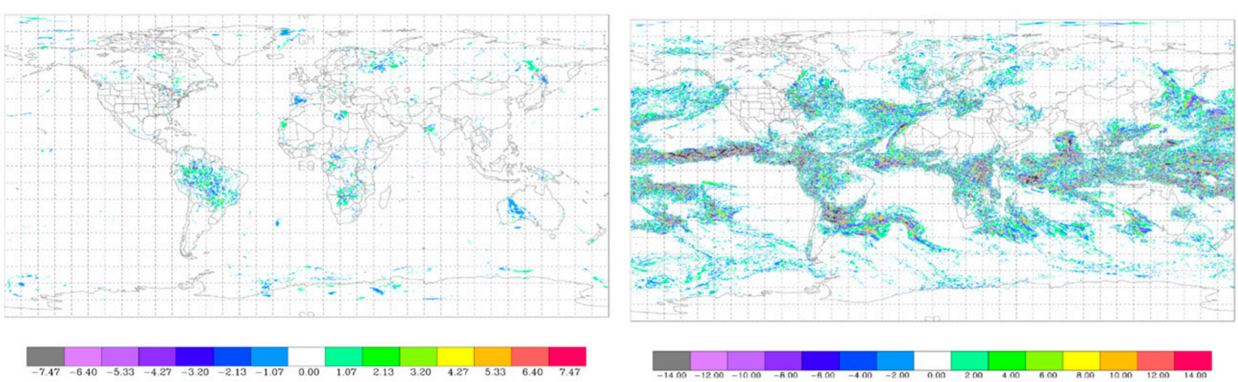

(c) difference fields of t (degree K) and precipitation (millimeter)

FIG. 7. Comparison of 120-h forecasts initialized with original and compressed kinetic fields. (top) The 120-h forecast of t and precipitation with original initial condition; (middle) as in (top), but with compressed initial condition; and (bottom) difference images of the two forecasts. FIM runs at 30-km resolution, with initial condition at 0000 UTC 28 Oct 2012 (Hurricane Sandy case).

Both cases impacted the continental United States with strong sustained surface winds and heavy precipitations. Numerical simulations of the tracks and intensities of those fast-evolving storm systems are critical to the forecasts of these high-impact weather events. They are also good cases to test the impact of compression of initial conditions on the numerical simulations.

The experiment is conducted as follows. First, we save the initial model kinetic state, wind components u, v. We compress the model variables at a specified precision to obtain a new, compressed model state $\hat{u}$, $\hat{v}$. Then, we run FIM with both sets of initial conditions to 120 h. After the model runs finish, we visually and numerically examine and compare the forecast results from both runs.

TABLE 5. Mean L^2 norm differences between the model runs with original and compressed initial kinetic fields (Hurricane Sandy case).

Forecast time (h)	Temp (K)	Wind u (m s^{-1})	Wind v (m s^{-1})	Height (m)
24	0.090 74	0.147 37	0.144 99	0.348 84
48	0.127 98	0.226 82	0.222 53	0.480 97
72	0.167 77	0.294 07	0.299 94	0.619 95
96	0.219 47	0.387 58	0.380 57	0.894 10
120	0.269 69	0.475 03	0.479 88	1.215 37

TABLE 6. Mean L^2 norm differences between the model runs with original and compressed initial kinetic fields (spring snowstorm case).

Forecast time (h)	Temp (K)	Wind u (m s^{-1})	Wind v (m s^{-1})	Height (m)
24	0.092 42	0.152 36	0.150 58	0.357 06
48	0.124 64	0.225 36	0.221 80	0.479 03
72	0.163 02	0.299 62	0.295 80	0.643 66
96	0.207 96	0.381 61	0.377 48	0.854 83
120	0.257 54	0.466 94	0.456 57	1.194 46

a. *Hurricane Sandy case*

Tropical Cyclone Sandy was the 18th named tropical cyclone of the 2012 Atlantic hurricane season (1 June–November 30). The cyclone formed in the central Caribbean on 22 October 2012 and intensified into a hurricane as it tracked north across Jamaica, eastern Cuba, and the Bahamas. Then, it moved northeast of the United States until turning west toward the mid-Atlantic coast on 28 October. Sandy transitioned into a post-tropical cyclone just prior to moving onshore near Atlantic City, New Jersey. Associated with Sandy, heavy rainfall occurred in Maryland, Virginia, Delaware, and New Jersey, and widespread heavy snow in the Appalachian Mountains from western North Carolina northeastward through southwestern Pennsylvania.

The model simulation for this case was initialized at 0000 UTC 28 October 2012. The initial u, v fields at 0000 UTC 28 October 2012 are compressed with L^∞ error <0.0625.

Figures 6 and 7 show a comparison of 120-h forecasts of u, v, t, and precipitation, and the differences of forecasts when initialized with original and compressed initial data.

There are small differences between the numerical forecasts with the original and compressed initial kinetic state, and the dynamic structures of the atmospheric circulation (temperature and wind fields), after 120 h of numerical simulation, are very similar.

To further examine the error evolution with time, Table 5 lists the mean L^2 errors for t, u, v, and geopotential height h, at different forecast times.

b. *A spring storm case*

The spring storm case was associated with a low pressure system that developed off the Florida coast on 24 March 2013. This system moved quickly as it made its way to the Northeast on 25 March and rapidly intensified as it neared the New England offshore waters. Precipitation associated with the storm occurred in the mid-Atlantic and New England coastal areas during the morning of 25 March and continued through the overnight hours, ending by the early morning on 26 March. An area of 4–7 in. of snow was reported across southern Delaware and parts of extreme southern New Jersey. Winds associated with the system were blustery with gusts generally in the 35–45 mph range.

The model simulation for the case was initialized at 0000 UTC 24 March 2014. The same experiment procedure is repeated for this case.

Results are shown in Table 6, and Figs. 8 and 9. Quite similar results to the Hurricane Sandy case are obtained for this case.

c. *Comparison of 500-hPa height anomaly correlation coefficient scores*

To further assess the impact of the compression of initial datasets on the model forecast skills, we compute the anomaly correlation coefficients of the 500-hPa height fields for FIM forecasts with original and compressed initial datasets. The anomaly correlation coefficient (ACC), a coefficient of correlation between the forecast and analysis relative to the climate mean, is defined as

$$\mathrm{ACC} = \frac{\overline{(f-c)(a-c)}}{\left[\overline{(f-c)^2}\,\overline{(a-c)^2}\right]^{1/2}}, \tag{4}$$

where f and a are forecast and analysis, respectively; and c is the corresponding climate mean.

For the experiment, we use the analysis and climate data produced by Global Forecast System (GFS) of the National Centers for Environmental Prediction (NCEP). We compute the ACC scores for both Hurricane Sandy and the 2014 spring snowstorm cases, at five forecast times, and they are listed in Tables 7 and 8.

It appears that at least for the two experiment cases, there are negligible differences of ACC scores between the forecasts with original and compressed initial conditions.

d. *Remarks on the experiment results for two test cases*

From the above-mentioned experiments, we have the following observations:

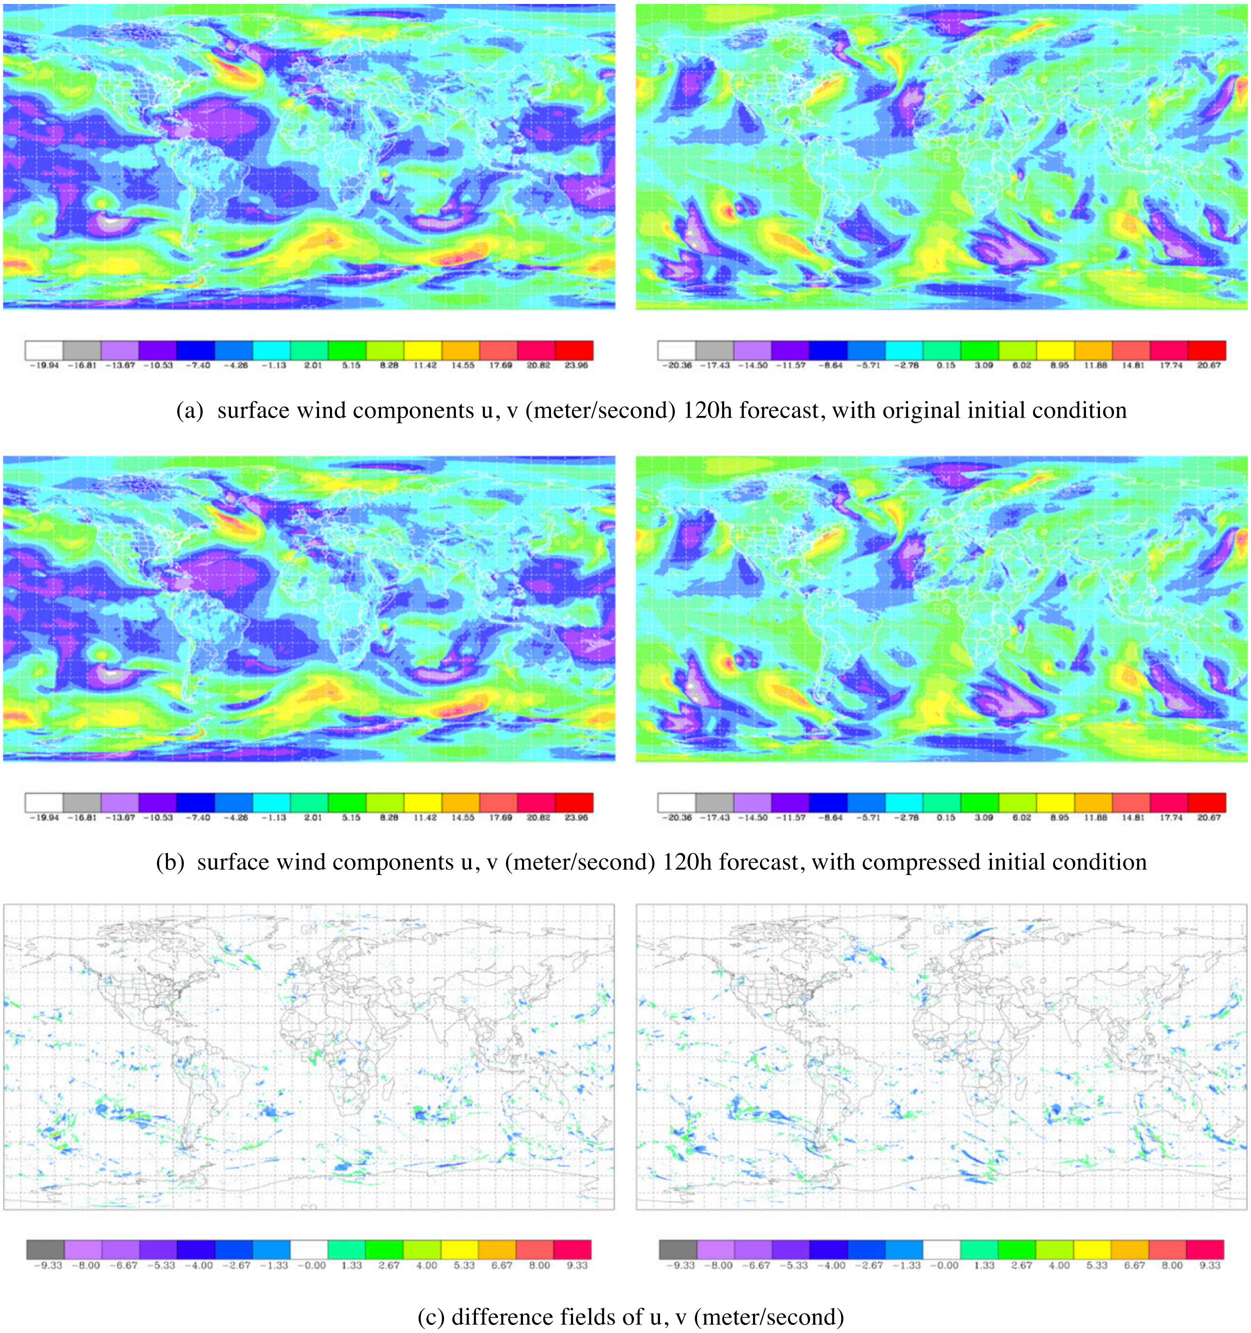

FIG. 8. Comparison of 120-h forecasts initialized with original and compressed kinetic fields. (top) The 120-h forecast of u, v with original initial condition; (middle) as in (top), but with compressed initial condition; (bottom) difference images of the two forecasts. FIM runs at 30-km resolution, with initial condition at 0000 UTC 24 Mar 2014 (2014 spring snowstorm case).

(i) Numerical simulations reach virtually the same state in 5-day forecasts with original and compressed initial kinetic conditions.

(ii) The differences between the two simulations are rather small, and the relative differences (the differences normalized to the magnitude of the variable) are similar for different model variables.

(iii) The differences grow with simulation time. The growth rates are mostly linear, and occasionally between linear and quadratic.

(iv) For the forecasts up to 120 h with original and compressed initial conditions, there are negligible differences in anomaly correlation coefficient scores. It indicates that the impact of compression

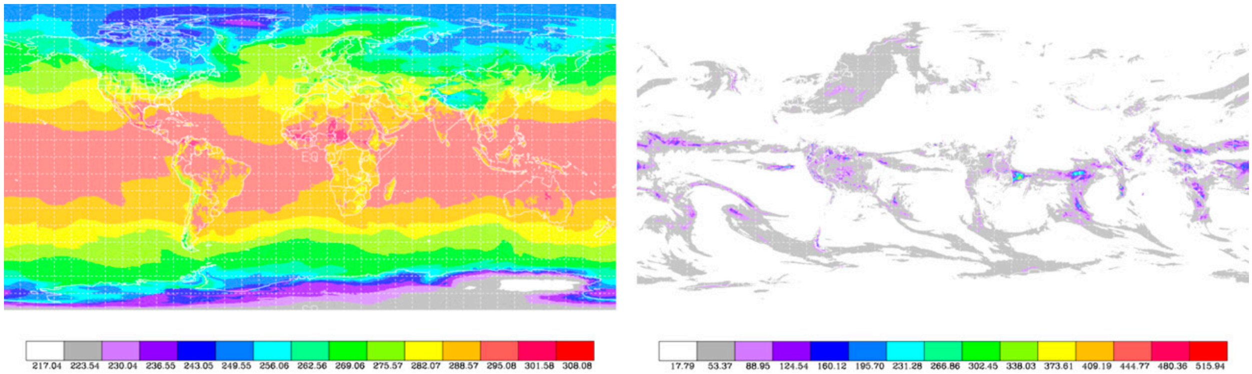

(a) surface temperature t (degree K) and precipitation (millimeter) 120h forecast, with original initial condition

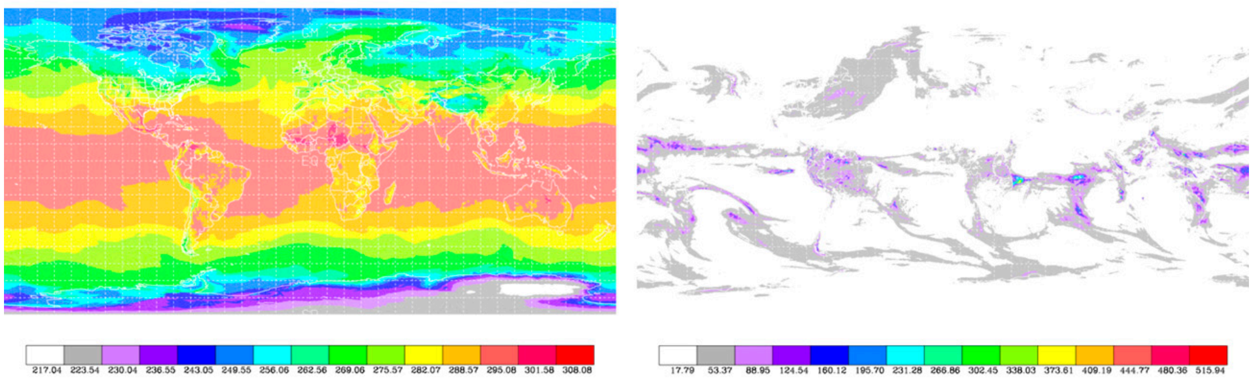

(b) surface temperature t (degree K) and precipitation (millimeter) 120h forecast, with compressed initial condition

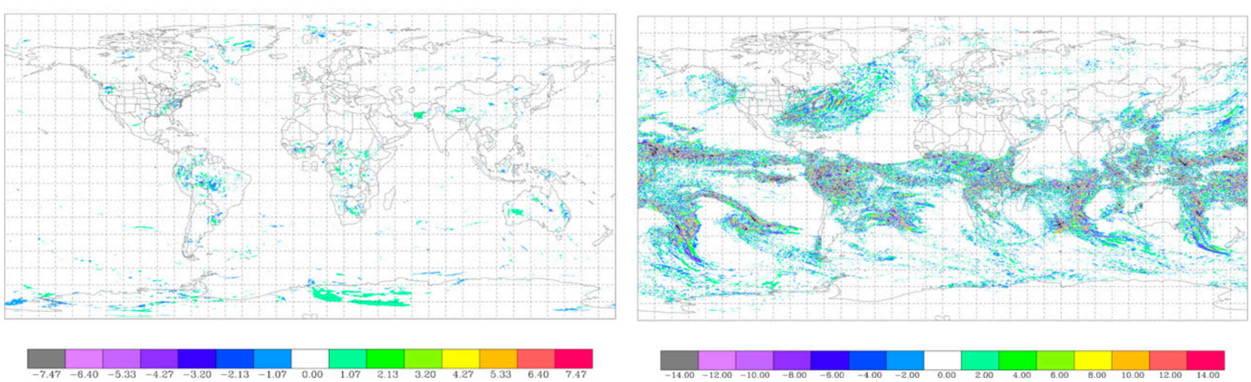

(c) difference fields of t (degree K) and precipitation (millimeter)

FIG. 9. Comparison of 120-h forecasts initialized with original and compressed kinetic fields. (top) The120 h forecast of temperate t and precipitation with original initial condition; (middle) same forecast with compressed initial condition; and (bottom) the difference images of the two forecasts. FIM runs at 30-km resolution, with initial condition at 0000 UTC 24 Mar 2014 (2014 spring snowstorm case).

of initial conditions on model forecast skill is minimal.

(v) In a sense, this experiment is more about the model sensitivity to the perturbation of initial condition than the compression performance of model initial condition. From this perspective, FIM is stable with a small perturbation to its initial condition.

It is apparent that these observations are model dependent, closely related to the model numerics and physics parameterization schemes. Nonetheless, experiment results give us some guidelines for what we could expect from other global weather models, in terms of their numerical sensitivities to small perturbations in initial conditions caused by compression.

TABLE 7. Comparison of 500-hPa height ACCs: Hurricane Sandy case.

Initial dataset	24 h	48 h	72 h	96 h	120 h
Original	99.70	99.01	97.44	94.32	89.74
Compressed	99.70	99.01	97.44	94.34	89.70

TABLE 8. Comparison of 500-hPa height ACCs: 2014 spring snowstorm case.

Initial dataset	24 h	48 h	72 h	96 h	120 h
Original	99.54	98.50	96.15	93.13	88.51
Compressed	99.54	98.49	96.14	93.13	88.59

5. Discussions and summary

We proposed a wavelet compression technique for native datasets on icosahedral grids. The experiments have shown that the technique achieves good compression performance for both classic lossy data compression and precision-controlled lossy data compression.

The major differences between the proposed data compression technique and existing ones are as follows: 1) the native icosahedral grid datasets are compressed directly, as opposed to being interpolated to the Cartesian grid and then compressed; and 2) a two-domain quantization scheme is applied, in which a transformed domain quantization efficiently reduces L^2 and L^∞ errors and a physical domain quantization satisfies the L^∞ error requirement.

For many visualization applications, using classic lossy wavelet compression, without a strict precision specification, the proposed compression technique can achieve dramatic data compressions. It has been demonstrated that a 50:1 or higher compressions can be achieved for many model variables and 30:1 or higher compressions for other variables. At this compression level, no noticeable visual difference appears in rendered images.

For most meteorological applications, precision requirements are specified for compressed model variables. Significant compressions are achieved for commonly used precision requirements. For example, at 0.0625 precision threshold, compression ratios of 15.1:1, 14.5:1, and 7.9:1 are achieved for temperature, zonal wind component, and relative humidity, respectively.

Compared to the GRIB2 with JPEG2000 packing option, which is currently only available for Cartesian grid data, the proposed technique achieves a better compression with the same L^∞ norm error specification. In addition, as we have indicated earlier, at the same level of L^∞ norm error, the proposed compression technique has a smaller mean L^2 norm error, close to one-half of that of the GRIB2 compressed data.

The numerical experiment conducted in this research also includes a preliminary assessment of the impact of compressed initial conditions on numerical simulations. This is the first attempt of such an experiment that is designed to answer the question of whether a precision-controlled lossy data compression can be used to compress data archives that may potentially be used as initial model states. The preliminary results indicate that the answer is promising.

Following the same grid arrangement strategy and data compression scheme, the proposed technique can be extended to other unstructured grids, such as staggered icosahedral grids and cubed-sphere grids. For example, for a staggered icosahedral C grid, the grid points on the cell edges within a rhombus can be arranged into three separate rhombus lattices, according to their triangular grid lines. After this grid rearrangement, each rhombus grid can be treated as a Cartesian lattice and the rest of the proposed compression technique can be applied in the same way.

For future research, a more concrete and detailed assessment of the proposed compression technique is desired to further improve its performance. Specifically, we want to investigate how to preprocess model data for better compression, especially for vector fields; we also want to understand how and at what level of compression meaningful loss of information content occurs, which could impair the numerical simulation of weather activities. More model initial condition variables will be compressed and tested to assess their impact on numerical simulations. Computationally, it is desirable to develop a more sophisticated but still practical bit allocation algorithm for quantizations in the transformed and physical domains. It is also desirable to parallelize the compression software to speed up the encoding and decoding process.

Acknowledgments. The authors thank Dr. Chou Shih-Yun for his insightful comments and James Rosinski for his helpful discussions and detailed internal review of the manuscript. The authors thank the anonymous reviewers for their comments and suggestions. The authors also thank Drs. Alexander E. MacDonald and Stan Benjamin, and Sher Schranz for their support of this research.

REFERENCES

Ansari, R., N. Memon, and E. Ceran, 1998: Near-lossless image compression techniques. *J. Electron. Imaging,* **7,** 486–494, doi:10.1117/1.482591.

Antonini, M., M. Barlaud, P. Mathieu, and I. Daubechies, 1992: Image coding using the wavelet transform. *IEEE Trans. Image Process.,* **1,** 205–220, doi:10.1109/83.136597.

Baumgardner, J. R., and P. O. Frederickson, 1985: Icosahedral discretization of the two-sphere. *SIAM J. Numer. Anal.,* **22,** 1107–1115, doi:10.1137/0722066.

Bleck, R., S. Benjamin, J. Lee, and A. E. MacDonald, 2010: On the use of an adaptive, hybrid-isentropic vertical coordinate in global atmospheric modeling. *Mon. Wea. Rev.,* **138,** 2188–2210, doi:10.1175/2009MWR3103.1.

——, and Coauthors, 2015: A vertically flow-following icosahedral grid model for medium-range and seasonal prediction. Part I: Model description. *Mon. Wea. Rev.,* **143,** 2386–2403, doi:10.1175/MWR-D-14-00300.1.

Cohen, A., I. Daubechies, and J.-C. Feauveau, 1992: Biorthogonal bases of compactly supported wavelets. *Commun. Pure Appl. Math.,* **45,** 485–560, doi:10.1002/cpa.3160450502.

Daubechies, I., 1992: *Ten Lectures on Wavelets.* SIAM, 357 pp.

de Saint-Martin, F. M., P. Siohan, and A. Cohen, 1999: Biorthogonal filterbanks and energy preservation property in image compression. *IEEE Trans. Image Process.,* **8,** 168–178, doi:10.1109/83.743852.

DeVore, R., B. Jawerth, and B. J. Lucier, 1992: Image compression through wavelet transform coding. *IEEE Trans. Inf. Theory,* **38,** 719–746, doi:10.1109/18.119733.

Dey, C. H., 2007: Guide to the WMO table driven code form used for the representation and exchange of regularly spaced data in binary form: FM 92 GRIB edition 2. WMO Tech. Rep., 98 pp. [Available online at http://www.wmo.int/pages/prog/www/WMOCodes/Guides/GRIB/GRIB2_062006.pdf.]

Donoho, D. L., 1993: Unconditional bases are optimal bases for data compression and for statistical estimation. *Appl. Comput. Harmonic Anal.,* **1,** 100–115, doi:10.1006/acha.1993.1008.

Du, Q., and L. Ju, 2005: Finite volume methods on spheres and spherical centroidal Voronoi meshes. *SIAM J. Numer. Anal.,* **43,** 1673–1692, doi:10.1137/S0036142903425410.

Finkel, R. A., and J. L. Bentley, 1974: Quad trees: A data structure for retrieval on composite keys. *Acta Inf.,* **4,** 1–9, doi:10.1007/BF00288933.

Heikes, R., and D. Randall, 1995: Numerical integration of shallow-water equations on twisted icosahedral grid. Part II: A detailed description of grid and an analysis of numerical accuracy. *Mon. Wea. Rev.,* **123,** 1881–1887, doi:10.1175/1520-0493(1995)123<1881:NIOTSW>2.0.CO;2.

Jung, T., and Coauthors, 2012: High-resolution global climate simulations with the ECMWF model in Project Athena: Experimental design, model climate, and seasonal forecast skill. *J. Climate,* **25,** 3155–3172, doi:10.1175/JCLI-D-11-00265.1.

Karray, L., P. Duhamel, and O. Rioul, 1998: Image coding with an L^∞ norm and confidence interval criteria. *IEEE Trans. Image Process.,* **7,** 621–631, doi:10.1109/83.668016.

Lee, J.-L., and A. E. MacDonald, 2009: A finite-volume icosahedral shallow-water model on a local coordinate. *Mon. Wea. Rev.,* **137,** 1422–1437, doi:10.1175/2008MWR2639.1.

——, R. Bleck, and A. E. MacDonald, 2010: A multistep flux-corrected transport scheme. *J. Comput. Phys.,* **229,** 9284–9298, doi:10.1016/j.jcp.2010.08.039.

Lessig, C., and E. Fiume, 2008: SOHO: Orthogonal and symmetric Haar wavelets on the sphere. *ACM Trans. Graphics,* **27** (1), Article 4, doi:10.1145/1330511.1330515.

Lucero, A., S. Cabrera, A. Aguirre, M. Argaez, and E. Vidal, 2004: Comparison of bit allocation methods for compressing three-dimensional meteorological data after applying KLT. *2004 IEEE Int. Geoscience and Remote Sensing Symp.,* Vol. 4, IEEE, 2514–2517.

Majewski, D., D. Liermann, P. Prohl, B. Ritter, M. Buchhold, T. Hanisch, G. Paul, and W. Wergen, 2002: The operational global icosahedral–hexagonal gridpoint model GME: Description and high-resolution tests. *Mon. Wea. Rev.,* **130,** 319–338, doi:10.1175/1520-0493(2002)130<0319:Togihg>2.0.CO;2.

Mallat, S. G., 1989: A theory for multiresolution signal decomposition: The wavelet representation. *IEEE Trans. Pattern Anal. Mach. Intell.,* **11,** 674–693, doi:10.1109/34.192463.

Marpe, D., G. Blattermann, J. Ricke, and P. Maab, 2000: A two-layered wavelet-based algorithm for efficient lossless and lossy image compression. *IEEE Trans. Circuits Syst. Video Technol.,* **10,** 1094–1102, doi:10.1109/76.875514.

Putman, W. M., and S.-J. Lin, 2009: A finite-volume dynamical core on the cubed-sphere grid. *Numerical Modeling of Space Plasma Flows: ASTRONUM-2008,* N. Pogorelov et al., Eds., Astronomical Society of the Pacific Conference Series, Vol. 406, ASP, 268–276.

Said, A., and W. A. Pearlman, 1996: A new, fast, and efficient image codec based on set partitioning in hierarchical trees. *IEEE Trans. Circuits Syst. Video Technol.,* **6,** 243–250, doi:10.1109/76.499834.

Satoh, M., T. Matsuno, H. Tomita, H. Miura, T. Nasuno, and S. Iga, 2008: Nonhydrostatic icosahedral atmospheric model (NICAM) for global cloud resolving simulations. *J. Comput. Phys.,* **227,** 3486–3514, doi:10.1016/j.jcp.2007.02.006.

Schröder, P., and W. Sweldens, 1995: Spherical wavelets: Efficiently representing functions on the sphere. *SIGGRAPH '95: Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques,* S. G. Mair and R. Cook, Eds., ACM, 161–172.

Shapiro, J. M., 1993: Embedded image coding using zerotrees of wavelet coefficients. *IEEE Trans. Signal Process.,* **41,** 3445–3462, doi:10.1109/78.258085.

Skamarock, W. C., J. B. Klemp, M. G. Duda, L. D. Fowler, S.-H. Park, and T. D. Ringler, 2012: A multiscale nonhydrostatic atmospheric model using centroidal Voronoi tesselations and C-grid staggering. *Mon. Wea. Rev.,* **140,** 3090–3105, doi:10.1175/MWR-D-11-00215.1.

Strang, G., and T. Nguyen, 1998: *Wavelets and Filter Banks.* Wellesley-Cambridge Press, 520 pp.

Sweldens, W., 1997: The lifting scheme: A construction of second generation wavelets. *SIAM J. Math. Anal.,* **29,** 511–546, doi:10.1137/S0036141095289051.

Tian, J., and R. O Wells Jr., 1996: A lossy image codec based on index coding. *Proceedings of the IEEE Data Compression Conference (DCC '96),* J. A. Storer, Ed., IEEE, 456.

Tomita, H., M. Tsugawa, M. Satoh, and K. Goto, 2001: Shallow water model on a modified icosahedral geodesic grid by using spring dynamics. *J. Comput. Phys.,* **174,** 579–613, doi:10.1006/jcph.2001.6897.

——, M. Satoh, and K. Goto, 2004: A new dynamical framework of global nonhydrostatic model using the icosahedral grid. *Fluid Dyn. Res.,* **34,** 357–400, doi:10.1016/j.fluiddyn.2004.03.003.

Topiwala, P. N., Ed., 1998: *Wavelet Image and Video Compression.* International Series in Engineering and Computer Science, Vol. 450, Springer, 438 pp.

Wang, N., and S. Madine, 1998: FX-NET: A Java-based Internet client interface to the WFO-Advanced workstation. Preprints, *14th Int. Conf. on Interactive Information and Processing Systems (IIPS) for Meteorology, Oceanography, and Hydrology,* Phoenix, AZ, Amer. Meteor. Soc., 427–429.

——, and R. Brummer, 2003: Experiment of a wavelet-based compression technique with precision control. *19th Int. Conf. on Interactive Information and Processing Systems (IIPS) for Meteorology, Oceanography, and Hydrology,* Long Beach, CA, Amer. Meteor. Soc., 4.8. [Available online at https://ams.confex.com/ams/annual2003/techprogram/paper_56240.htm.]

——, and J. Lee, 2011: Geometric properties of the icosahedral-hexagonal grid on the two-sphere. *SIAM J. Sci. Comput.,* **33,** 2536–2559, doi:10.1137/090761355.

Yea, S., and W. A. Pearlman, 2006: A wavelet-based two-stage near-lossless coder. *Trans. Image Process.,* **15,** 3488–3500, doi:10.1109/TIP.2006.877525.

Ziv, J., and A. Lempel, 1977: A universal algorithm for sequential data compression. *IEEE Trans. Inf. Theory,* **23,** 337–343, doi:10.1109/TIT.1977.1055714.