Dmitry Kolomenskiy · Ryo Onishi · Hitoshi Uehara

WaveRange: wavelet-based data compression for three-dimensional numerical simulations on regular grids

Received: 11 May 2021 / Revised: 27 August 2021 / Accepted: 13 October 2021 / Published online: 1 January 2022
© The Author(s) 2021

Abstract A wavelet-based method for compression of three-dimensional simulation data is presented and its software framework is described. It uses wavelet decomposition and subsequent range coding with quantization suitable for floating-point data. The effectiveness of this method is demonstrated by applying it to example numerical tests, ranging from idealized configurations to realistic global-scale simulations. The novelty of this study is in its focus on assessing the impact of compression on post-processing and restart of numerical simulations.

Keywords High-performance computing · Data compression · Wavelets

1 Introduction

Partial differential equations often arise in physical sciences from three-dimensional (3D) continuum models, yielding boundary value problems for continuous field variables defined over 3D spatial domains. Numerical solution of these problems involves discretization, and among all available methods, many employ regular grids such that the discretized field variables can be stored as three-dimensional arrays. Regular grids prevail in the high-performance computing (HPC) for enabling fast and efficient implementation of high-order numerical methods with good parallel scalability. HPC simulations based on Cartesian grids are extremely diverse and include, to name a few, particle-laden fluid flows (Fornari et al. 2016), solar flares (Bakke et al. 2018) and neutron transport inside the core of a nuclear reactor (Baudron and Lautard 2007). In Earth science, curvilinear grids such as yin–yang (Kageyama and Sato 2004) or cubed sphere (Ronchi et al. 1996) are used for global weather simulations (Nakano et al. 2017), geodynamo simulations (Kageyama et al. 2004), calculation of seismograms (Komatitsch et al. 2013), etc. Toroidal grids are used in tokamak plasma simulations (Korpilo et al. 2016).

HPC produces large volumes of output data. A numerical simulation using several hundred or thousands of processor cores would allocate three-dimensional arrays totaling to several gigabytes or terabytes. If the

D. Kolomenskiy (✉) · R. Onishi
Global Scientific Information and Computing Center (GSIC), Tokyo Institute of Technology, Tokyo, Japan
E-mail: dkolom@gmail.com; d.kolomenskiy@skoltech.ru

R. Onishi
E-mail: onishi.ryo@gsic.titech.ac.jp

Present Address:

D. Kolomenskiy
Center for Design, Manufacturing and Materials, Skolkovo Institute of Science and Technology, Moscow, Russia

R. Onishi · H. Uehara
Japan Agency for Marine–Earth Science and Technology (JAMSTEC), Yokohama, Japan

solution evolves in time, a new three-dimensional dataset is produced at every time step. This massive data flow is characteristic of big data applications (Hilbert 2016), and it is not surprising that high-performance numerical simulations using regular grids hit the limitations of the contemporary data handling technologies. In particular, data storage capacity is finite. To alleviate this constraint in practice, *in situ* data reduction is routinely performed during the simulation and only the selected integral physical quantities or time-resolved sequences are stored. Nevertheless, it is often required to store the three-dimensional fields for purposes such as scientific visualization, restart of the simulation or additional post-processing. These large datasets quickly saturate the available disk space if stored as floating-point arrays without compression.

Lossless data compression tools, such as LZMA compression, reduce the typical floating-point binary data file size by less than 20%. Accepting some data loss, it is a common practice to store the simulation output fields in single precision and down-sample the data, e.g., save every second point value in each direction. Such reduced datasets, while being of insufficient information capacity for the simulation, are often suitable for post-processing. Given that the differences between neighboring point values can be interpreted as wavelet coefficients, a more refined version of the mentioned approach is to apply wavelet transform to the field and encode the significant portion of the wavelet coefficients using a common data compression method such as entropy coding. This technique is currently widely in use for image compression, being part of the JPEG2000 standard (Taubman and Marcellin 2002). Its suitability for the computational fluid dynamic (CFD) data compression has been evaluated by Schmalzl (2003), alongside other image compression algorithms, and later revisited by Woodring et al. (2011). Application of 3D orthogonal wavelet transforms to volume data has been presented by Muraki (1993). A related method has been recently implemented for multi-resolution rendering and storage of geoscience models with discontinuities (Peyrot et al. 2019). In the work of Sakai et al. (2011, 2013a, 2013b), wavelet compression has been used in the context of numerical simulation of industrial fluid flows using the building-cube method, with focus on aeroacoustics. Overall good performance has been reported in terms of the compression ratio, accuracy and parallel performance for large datasets. However, error control was not explicitly handled.

Of all types of output, the restart data may pose the most stringent accuracy constraint on the lossy compression, because it is commonly expected that restart should not influence the final result of the simulation. One can expect that the required restart data accuracy depend on the physical model. Indeed, there exist models that are insensitive to ample reduction in the width of the floating-point significand (Hatfield et al. 2018). The use of lossy data compression techniques has also been advocated by showing that compression effects are often unimportant or disappear in post-processing analyses (Baker et al. 2016), and substantial gain in the compression ratio can be achieved while keeping the error at acceptable level in terms of physically motivated metrics (Laney et al. 2013). Consequently, it appears reasonable to adjust the restart data storage to the precision justified by the level of model error. We further investigate into this issue by considering two atmospheric dynamic simulations in the present paper.

In the context of fluid dynamics and atmospheric science, since the ability of wavelets to provide compressed representation of turbulent flows was recognized (Farge 1992), a significant body of research focused on the development of wavelet-based adaptive numerical methods allowing to lower the computational complexity and memory requirements of high-Reynolds number flow simulations (Schneider and Vasilyev 2010). Studies taking the perspective of CFD data storage remain relatively sparse. Besides the aforementioned work, a wavelet transform-vector quantization compression method for ocean models was proposed by Bradley and Brislawn (1993), the effect of lossy wavelet-based compression on barotropic turbulence simulation data has been studied by Wilson (2002), a hybrid method with supercompact multi-wavelets was suggested by Kang et al. (2003), and trade-offs in accuracy, storage cost and execution times using different wavelet transforms were considered by Li et al. (2015).

It should be mentioned that wavelet bases are not the only that yield sparse representation of turbulent flow fields. Decompositions such as proper orthogonal decomposition (POD, Berkooz et al. 1993; Bala-jewicz et al. 2013) or dynamic mode decomposition (DMD, Schmid 2010) are also used for this purpose and employed in CFD output data compression methods (Lorente et al. 2010; Bi et al. 2014). Each method has its own advantages, but in this paper we only consider the wavelet-based approach that may be more suitable for large datasets for its lower computational complexity, compared with the POD or DMD. A comparison of dimensionality reduction using POD and wavelet coherent structure identification can be found in a study by Schlegel et al. (2009). The wavelet-based method presented in this work does not require any time history, i.e., it can be applied to compress a single time snapshot.

Sub-band coding (SBC) (Røsten et al. 2004) and the use of more general filter banks than the discrete wavelet (Duval and Røsten 2000) have been considered in the context of seismic data compression. Another,

conceptually different family of methods can be described as prediction-based compression algorithms (see Najmabadi et al. 2017; Lakshminarasimhan et al. 2011; Liang et al. 2018; and references therein) that exploit spatiotemporal patterns in the data. A comparative discussion of different existing approaches to scientific data reduction, including lossy data compression, can be found in a recent review paper by Li et al. (2018), also see surveys by Balsa Rodríguez et al. (2014); Beyer et al. (2014) for GPU-based volume rendering applications.

Despite the fact that there exist a variety of compression methods that can compress 3D data, they are still rarely used for simulation restart and long-term simulation output data storage. One of the reasons is scarcity of published information regarding the impact of compression on post-processing and restart. Quantification of such impact constitutes the main novelty in the present study. The first objective of this work is to implement a data compression method suitable for files that contain three-dimensional floating-point arrays output from numerical simulation. We mainly target applications in Earth science such as atmospheric dynamic simulation, but the method and the software are designed to fit broader use. Our approach is similar to the method of Sakai et al. (2013a) conceptually, but differs in many aspects such as error control and wavelet transform depth. Therefore, a self-contained description of the method is provided in Sect. 2. The computer code is implemented in C/C++; it is open-source and accessible via <https://github.com/pseudospectators/WaveRange>. It is described in Appendix F. Our second objective is to assess the impact of compression on post-processing and restart. We evaluate the performance of WaveRange and devise practical recommendations for its users. This constitutes Sect. 3 of the paper. We particularly focus on the relationship between reconstruction error and compression ratio, as well as its effect on the accuracy of post-processing and simulation restart. The compression and decompression performance with consideration of computational cost is examined in Sect. 4. A comparison with other compression methods is provided. Section 5 contains concluding remarks.

2 Problem definition and description of the method

On the highest level, the compression method consists of the following steps: (1) wavelet transform; (2) quantization; and (3) entropy coding. Reconstruction is the inverse of the above operations in the reverse order. These steps are schematically shown in the data flow diagram in Fig. 1. After an introductory comment on the input data in Sect. 2.1, the three algorithmic steps are described in Sects. 2.2, 2.3 and 2.4, respectively.

The motivation of this three-step design is the following. The wavelet transform can provide compressed representation of a signal if it is locally correlated, and it becomes particularly efficient if the fine-scale activity is sparse. For example, turbulent flow fields satisfy these criteria, and it is known that only few nonzero wavelet coefficients are sufficient to represent the dynamically active part of the flow (Schneider and Vasilyev 2010). Small detail coefficients use as much memory as large ones when stored in the floating-point format, but quantization can reduce the small-detail storage to a few (less than 8) bytes per coefficient. For instance, detail coefficients that are smaller than the accuracy threshold are approximated as integer zeros. As long as the fine-scale activity is sparse, the quantized wavelet transform contains constant sequences of one-byte integers. In particular, it contains sequences of zeros. Such sequences can be

Fig. 1 Data flow diagram of the compression and reconstruction method. Note that ‘dequantization’ is not the exact inverse of the quantization, i.e., $F \neq \tilde{F}$ in general

efficiently compressed by entropy coding. Thus, the present compression method is intended to minimize the compressed file size at a given accuracy.

2.1 Data grids

We restrict our attention to data sampled on single- or multi-block grids with each block using three-dimensional Cartesian indexing, as shown in Fig. 2. Numerical methods that involve such kinds of topology are common in HPC for the ease and efficiency of data management. Atmospheric flow simulations of the global scale can be performed using a yin–yang grid that consists of two overlapping blocks. Regional and urban simulations can incorporate geometrical representation of landscape features and buildings by using immersed boundary approaches (Lundquist et al. 2010; Matsuda et al. 2018) that effectively reduce the computational domain to a rectangular box.

Let $\{f_{i_x, i_y, i_z}\}$ be a three-dimensional scalar field sampled on a grid $\{x_{i_x, i_y, i_z}, y_{i_x, i_y, i_z}, z_{i_x, i_y, i_z}\}$, where $i_x = 1, \dots, n_x$; $i_y = 1, \dots, n_y$; $i_z = 1, \dots, n_z$. This grid may be curvilinear. However, the positions of the grid nodes in the physical space are not required by the data compression algorithm. Only the array of values f_{i_x, i_y, i_z} and the number of grid points in each direction n_x , n_y and n_z constitute the input data. If the grid is multi-block, either each block can be treated independently by the compression algorithm, or the blocks can be merged in one array if their dimensions match. Therefore, in the following discussion we will refer to the rectilinear grid schematically shown in Fig. 2, without loss of generality.

2.2 Wavelet transform

The wavelet transform produces an array of real values $\{F_{i_x, i_y, i_z}\}$ with the same number of elements as in the original array $\{f_{i_x, i_y, i_z}\}$, i.e., containing $n_x \times n_y \times n_z$ elements in total. We use a three-dimensional multi-resolution transform based on the biorthogonal Cohen–Daubechies–Feauveau 9/7 (CDF9/7) wavelet, expressed in terms of lifting steps (Daubechies and Sweldens 1998; Getreuer 2013). These wavelets have proved efficient for image compression; therefore, they are likely to be efficient for scientific data having similar locally correlated spatial structure. Numerical experiments with sample CFD velocity data confirmed that this transform yields higher compression ratio than using lower-order biorthogonal wavelets or orthogonal wavelets such as Haar, Daubechies (D4...D20) and Symlets (Sym4...Sym10). Previous work by Li et al. (2015) also showed that the CDF9/7 wavelet performed better than the Haar wavelet for scientific data compression, albeit using thresholding of the wavelet coefficients rather than entropy coding.

Basically, the transform consists of a finite sequence of filtering steps, called lifting steps, applied to one-dimensional (1D) signal. Let g_i , $i = 1, \dots, m$ be a 1D array of real values. One level of the forward transform consists in calculating the values of approximation coefficients s_j , $j = 1, \dots, \lceil m/2 \rceil$ and detail coefficients d_j , $j = 1, \dots, \lfloor m/2 \rfloor$ ($\lceil \cdot \rceil$ and $\lfloor \cdot \rfloor$ are the ceiling and the floor functions, respectively) using lifting steps

Fig. 2 Schematic drawing of a yin–yang grid (Kageyama and Sato 2004) and a cubed sphere grid (Ronchi et al. 1996) used in Earth science computations, and a Cartesian grid block with point-value data stored at the grid nodes

$$\begin{aligned}
s_j^{(0)} &= g_{2j-1}, \\
d_j^{(0)} &= g_{2j}, \\
d_j^{(1)} &= d_j^{(0)} + \alpha(s_j^{(0)} + s_{j+1}^{(0)}), \\
s_j^{(1)} &= s_j^{(0)} + \beta(d_j^{(1)} + d_{j-1}^{(1)}), \\
d_j^{(2)} &= d_j^{(1)} + \gamma(s_j^{(1)} + s_{j+1}^{(1)}), \\
s_j^{(2)} &= s_j^{(1)} + \delta(d_j^{(2)} + d_{j-1}^{(2)}), \\
s_j &= \zeta s_j^{(2)}, \\
d_j &= d_j^{(2)} / \zeta,
\end{aligned} \tag{1}$$

where $\alpha = -1.5861343420693648$, $\beta = -0.0529801185718856$, $\gamma = 0.8829110755411875$, $\delta = 0.4435068520511142$ and $\zeta = 1.1496043988602418$ (Daubechies and Sweldens 1998; Getreuer 2013). For boundary handling, coefficients $s_{[m/2]+1}^{(0)}$, $s_{[m/2]+1}^{(1)}$, $d_0^{(1)}$ and $d_0^{(2)}$ are defined and set identical to zero. If m is odd, then the missing element $d_{[m/2]}^{(0)}$, which is necessary for calculating $d_{[m/2]}^{(1)}$ and $s_{[m/2]}^{(1)}$ subsequently, is determined using extrapolation

$$\begin{aligned}
d_{[m/2]}^{(0)} &= -\frac{2}{1 + 2\beta\gamma} \left(\alpha\beta\gamma s_{[m/2]-1}^{(0)} + \beta\gamma d_{[m/2]-1}^{(0)} \right. \\
&\quad \left. + (\alpha + \gamma + 3\alpha\beta\gamma) s_{[m/2]}^{(0)} \right).
\end{aligned} \tag{2}$$

The output of (1) is an array of size m in which the first $[m/2]$ elements contain the approximation coefficients s_j and the last $[m/2]$ elements contain the detail coefficients d_j .

The inverse transform admits coefficients s_j and d_j at input, and resolves the lifting steps (1) in the reverse order to produce g_j at the output. More specifically, $s_l^{(2)}$ and $d_l^{(2)}$ are determined from the last two lines, then the third to last equation is solved with respect to $s_l^{(1)}$ and so on.

The three-dimensional transform is constructed by applying the above one-dimensional transform sequentially in the three directions of the three-dimensional data array; see Appendix A for further explanation. We found that repeating the process 4 times is practically sufficient to reach the maximum compression ratio, because, after 4 levels of the transform, the number of approximation coefficients becomes as small as 1/4096 of the total number of elements in the dataset. We therefore always set $L = 4$ in WaveRange. This means that, as long as $n_x \geq 15$, $n_y \geq 15$ and $n_z \geq 15$, the same number (four) of 1D transform levels is realized in all three directions, even in those cases when there is a dramatic difference between n_x , n_y and n_z . Note that the transform is in place, i.e., $\{f_{i_x, i_y, i_z}\}$ and $\{F_{i_x, i_y, i_z}\}$ physically share the same memory by being stored in the same array.

The inverse transform has similar algorithmic structure. It starts with the approximation coefficients and the details at the largest scale and repeats L iterations adding one extra level of details in each direction on each iteration.

An example field and its transform are displayed in Fig. 3. On the left, the original field in physical space, $\{f_{i_x, i_y, i_z}\}$, is displayed. In this example, it contains point values of the velocity component u_x in a turbulent wake flow, which is described in Section C. The magnitude of the point data values is shown on a logarithmic color scale. On the right, the output of Algorithm 2 is shown, in which the approximation coefficients and the detail coefficients on all levels are packed in one three-dimensional array $\{F_{i_x, i_y, i_z}\}$. Seven-eighth of its elements correspond to the smallest scale detail coefficients. They are all small in magnitude: Most of them are below the visibility threshold of the selected color scale, and only few large ones appear in the boundary layer. On the next and subsequent levels, details in the turbulent wake become increasingly larger in magnitude. The approximation coefficients occupy the upper left corner of the domain. They are small in number and large in magnitude. Note that the transform is near lossless (in the terminology of Li et al. 2018) such that all values of $\{f_{i_x, i_y, i_z}\}$ can be calculated from $\{F_{i_x, i_y, i_z}\}$ with floating-point round-off accuracy.

Fig. 3 Example visualization of a 3D dataset and its wavelet transform

2.3 Quantization

From this point on, $\{F_{i_x, i_y, i_z}\}$ is treated as a one-dimensional array of real values F_i , $i = 1, \dots, N$, where $N = n_x \times n_y \times n_z$. An example graphical illustration of quantization is shown in Fig. 4. Quantization represents each element F_i as a set of one-byte numbers (i.e., integer numbers ranging from 0 to 255), as required for the subsequent entropy coding. In general, entropy coders are not limited to 256 size alphabet, but this size is convenient as it corresponds to the *char* type, native in C language. A double-precision floating-point variable F_i can be losslessly quantized using eight one-byte numbers. However, in applications related to the numerical simulation of turbulent flows, quantization with some data loss has to be accepted in order to achieve the desired high compression ratio.

Lossy compression is achieved by $J < 8$ one-byte integer numbers $Q_i^{(j)}$ per variable F_i , indexed with a superscript $\langle j \rangle$, with common offsets $F_{\min}^{(j)}$ independent of i , using the following approximation:

$$F_i \approx \sum_{j=1}^J (Q_i^{(j)} \delta^{(j)} + F_{\min}^{(j)}), \quad (3)$$

with the approximation error no greater than $\delta^{(J)}$. The latter is controlled by the relative tolerance ε , which is a user-specified parameter of the compression routine. Note that ε can be regarded as the desired L^∞ error in f in physical space, normalized with $\max |f|$, but $\delta^{(J)}$ in (3) sets the absolute error in F_i in wavelet space. To relate $\delta^{(J)}$ with ε , we define

$$\varepsilon_F = \varepsilon \max |f_{i_x, i_y, i_z}| / \eta, \quad (4)$$

where the maximum is taken over all elements, and η is a constant coefficient. By trial and error, we have found that the value $\eta = 1.75$ guarantees that $\|f - \tilde{f}\|_\infty / \max |f| \approx \varepsilon$ as long as we use four levels of the wavelet transform. Quantization adds random noise with amplitude ε_F to the wavelet coefficients F , i.e., $\max |F - \tilde{F}| = \varepsilon_F$. The pointwise error in the physical space, $f - \tilde{f}$, is a weighted sum of $F - \tilde{F}$, with the weight determined by the lifting coefficients and by the length of the filter. From this consideration, it may be possible to evaluate η analytically, but the empirical value of 1.75 proves acceptable in all test cases considered in this paper, see D. We then assign $\delta^{(J)} = \varepsilon_F$. The values of $Q_i^{(j)}$, $\delta^{(j)}$, $F_{\min}^{(j)}$ and J are determined

Fig. 4 Illustrative example of the data quantization procedure

as follows from Algorithm 1, where we use square brackets to denote the nearest integer. In the example, $J = 3$ bit planes are required to represent the floating-point data with the desired accuracy ε . The amount of memory required to store one bit plane is 8 times less than for the original dataset. Thus, after quantization, the data size is $J/8$ times the original size.

Data: tolerance ε ; point values F_i , $i = 1, \dots, N$.

Result: bit depth parameter J ; offsets $F_{\min}^{(j)}$; quantization steps $\delta^{(j)}$; bit planes $Q_i^{(j)}$,

$$i = 1, \dots, N, j = 1, \dots, J.$$

$$F_i^{(1)} = F_i, i = 1, \dots, N;$$

$$\varepsilon_F = \varepsilon \max |f_{i_x, i_y, i_z}| / \eta, i_x = 1, \dots, n_x, i_y = 1, \dots, n_y, i_z = 1, \dots, n_z;$$

$$q = 256;$$

$$j \leftarrow 1;$$

repeat

$$F_{\min}^{(j)} = \min_{k=1, \dots, N} F_k^{(j)};$$

$$F_{\max}^{(j)} = \max_{k=1, \dots, N} F_k^{(j)};$$

if $(F_{\max}^{(j)} - F_{\min}^{(j)}) / (q - 1) > \varepsilon_F$ **then**

$$\delta^{(j)} = (F_{\max}^{(j)} - F_{\min}^{(j)}) / (q - 1);$$

else

$$\delta^{(j)} = \varepsilon_F;$$

end

for $i = 1$ to N **do**

$$Q_i^{(j)} = \lceil (F_i^{(j)} - F_{\min}^{(j)}) / \delta^{(j)} \rceil;$$

$$F_i^{(j+1)} = F_i^{(j)} - (Q_i^{(j)} \delta^{(j)} + F_{\min}^{(j)});$$

end

$$j \leftarrow j + 1;$$

until $\delta^{(j)} < \varepsilon_F$;

Algorithm 1: Quantization algorithm.

2.4 Entropy coding

Entropy coding is a technique allowing to reduce the quantity of storage required to hold a message without any loss of information. The message can be any sequence of characters drawn from some selected alphabet. The implementation that we use in the present work is based on an alphabet of 256-bit symbols, i.e., integers from 0 to 255. Conceptually, entropy coding consists in representing frequently occurring subsequences with few bits and rarely occurring ones with many bits. Shannon's coding theorem serves the entropy of a message as a theoretical bound to possible lossless compression (Shannon 1948).

The entropy coding method that we employ in our present work is the range coding (Martin 1979). Our current implementation uses an open-source coder *rngcod13* developed by Schindler (1999). The size of the encoded files produced by this coder is within a fraction of per cent from the theoretical bound, and the operation speed is faster compared to other similar methods such as arithmetic coding.

The input message consists of the bit planes $Q_i^{(j)}$ obtained during the quantization step. These data arrays are divided in blocks of 60,000 elements. The frequencies of each value in a block are counted and copied to the output stream of the coder. The data block is then encoded using the calculated frequencies. Given a

stream of 256-bit symbols $Q_i^{(j)}$ and their frequencies, the coder produces a shorter stream of bits to represent these symbols. The output stream is written in a file, appended with metadata necessary for decoding.

The output file size depends on the contents of the input message $Q_i^{(j)}$. It can be less than 1% if only a few elements of $Q_i^{(j)}$ are nonzero. It can be above 80% if the input message resembles white noise. Practical situations in between these two extremes are considered in Sect. 3.

3 Illustrative examples

In the subsequent sections of this paper, we discuss numerical examples that serve to evaluate the performance of the method and to gain better understanding of different properties of the compressed data. A homogeneous isotropic turbulence dataset, which can be regarded as a highly idealized representation of atmospheric flow, is examined in Sect. 3.1. Quantitative measures such as the error norm, the relative compressed file size and the compression ratio are introduced, and their scaling with respect to the tolerance and the data size is analyzed. A similar analysis with application to seismology is discussed in Sect. 3.2. After gaining insight into the basic typical features of the compressed data, the discussion proceeds to realistic atmospheric flow simulations. In Sect. 3.3, a global weather simulation for typhoon prediction is considered, with special attention to restart of the simulation from a compressed restart data file. Restarts using different levels of compression tolerance are tested. Restarts of an urban-scale weather simulation are discussed in Sect. 3.4.

3.1 Fluid turbulence simulation

The homogeneous isotropic turbulence (HIT) dataset considered here is similar to the velocity fields analyzed by Onishi et al. (2013). It was obtained by integrating the incompressible Navier–Stokes equations in a 2π -periodic cube box, in dimensionless units, using a fourth-order finite-difference scheme, second-order Runge–Kutta time marching and HSMAC velocity–pressure coupling (Onishi et al. 2011). The flow domain was discretized using a uniform Cartesian grid consisting of $n^3 = 512^3$ staggered grid points. Power input was supplied using external isotropic forcing to achieve a statistically stationary state. A snapshot velocity field (u, v, w) was stored in double precision, each component occupying 1 GB of disk space.

The dataset used in the present study is visualized in Fig. 5a using an iso-surface of the vorticity magnitude, and the energy spectrum is shown in Fig. 5b. Blue color corresponds to the original data. The three velocity components are in the range $[-4.08, 4.16]$, $[-5.20, 4.82]$ and $[-4.31, 4.74]$, respectively. The turbulence kinetic energy is equal to $K = 1.44$. With the dissipation rate $\epsilon = 0.262$, the Kolmogorov length scale is equal to $\eta = (v^3/\epsilon)^{1/4} = 0.00845$, such that $k_{\max}\eta = 2.16$, where $k_{\max} = n/2$. The Taylor microscale $\lambda = \sqrt{10\nu K/\epsilon} = 0.246$ yields the Reynolds number $Re_\lambda = \sqrt{2K/3}\lambda/\nu = 218$. Homogeneous isotropic turbulence presents a challenge for the data compression. Turbulent flow fills the entire domain with fluctuations at all scales, the smallest scale being of the same order as the discretization grid step (see Fig. 5b).

Before analyzing the compression performance in this case, let us discuss the error control properties. Among the variety of physically motivated error metrics we choose L^∞ , which is intuitive and perhaps the most stringent. The L^∞ norm is the most sensitive norm to large local errors. Scientific computing commonly relies on numerical evaluation of derivatives (using, e.g., finite-difference schemes). Therefore, even if only one data point value contains a large error, this error is likely to propagate in space and in time in the subsequent numerical simulation or post-processing. Other, problem specific metrics will be discussed later on. Figure 6a shows the relative L^∞ error of the velocity field reconstructed from compressed data. The compression algorithm takes the tolerance ϵ as a control parameter. This allows to plot the reconstruction error as a function of ϵ . Each velocity component is scaled by its maximum absolute value, yielding

$$e_\infty = \frac{\max |\tilde{f} - f|}{\max |f|}, \quad (5)$$

where f stands for one of the components (u , v or w) of the original velocity field, and $\tilde{f}$ is its reconstruction from the compressed data. The maxima and the minima are calculated over all grid points. In Fig. 6a, the dash-dot diagonal line visualizes the identity relationship between ϵ and e_∞ . The actual data for all three velocity components closely follows this trend, with the discrepancy only becoming noticeable when ϵ is smaller than the accumulated round-off error, and for large ϵ when the discrete nature of quantization becomes apparent as there remain only few nonzero detail coefficients. For all $\epsilon \in [10^{-14}, 10^{-3}]$, the difference between ϵ and e_∞ is less than 15%, i.e., the desired error control is successfully achieved.

Fig. 5 HIT dataset. **a** Flow visualization using a vorticity magnitude iso-surface corresponding to 15% of the maximum value. Black lines show the extent of the computational domain. Iso-surfaces of the original field (cyan) and reconstructed after compression with $\varepsilon = 10^{-2}$ (magenta) are superimposed and overlap almost perfectly. **b** Energy spectra of the original and the reconstructed velocity fields

Figure 6b shows a plot of the relative compressed file size

$$\Sigma = \frac{s(\varphi)}{s(f)} \times 100\% \quad (6)$$

versus e_∞ , where $s(f) = 1$ GB is the disk space required to store a 512^3 array of double-precision values, and $s(\varphi)$ is the respective compressed file size in GB. The dash-dot diagonal line in this plot represents the storage requirement for a 512^3 array of hypothetical custom-precision values, which is 100% in the case of double precision, 50% for the single precision and 0 for the full loss of information, and all intermediate values are linearly interpolated. This line sets a reference in terms of the compression achievable by simply quantizing the floating-point array elements with tolerance ε and using smaller, but constant number of bits per element. Its slope is equal to 6.39% of storage per one decimal order of magnitude of accuracy.

The point markers connected by solid lines show the amount disk space actually used by the compressed velocity field files and the respective reconstruction error. When the accuracy is in the range between 10^{-14} and 10^{-3} , it is well approximated with a fitting

$$\Sigma_{HIT} = (-0.12 - 0.052 \log_{10} e_\infty) \times 100\%, \quad (7)$$

which is shown with a dotted magenta line. The negative intercept can be explained as follows. When ε (and, consequently, e_∞) is large, only a few significant bits are sufficient to represent the field (u, v, w) with the desired accuracy. Hence, the wavelet transform after quantization becomes sparse, and it is very efficiently compressed by entropy coding. In the intermediate range of ε , increasing accuracy by one order of magnitude comes at a cost of 5.2% increased storage. Least significant bit planes of the wavelet coefficients are not sparse, they are noise-like, and their lossless compression ratio by entropy coding asymptotically tends to a fairly low value typical of noise, $0.0639/0.052 \approx 1.23$, in the hypothetical limit of $\varepsilon \rightarrow 0$. However, for $\varepsilon < 10^{-14}$, accumulation of round-off errors becomes significant, Σ sharply increases to 76%, while e_∞ saturates at 5×10^{-15} . Round-off error could be avoided by switching to integer wavelet transform ensuring perfect reconstruction. Extrapolation of the linear trend suggests that this approach may be superior than, for example, direct application of LZMA encoding, as shown by crosses in Fig. 6b. It is also noteworthy that compression with $\varepsilon = 10^{-8}$ allows to reduce the data storage by a factor of 3, which is significantly better than using the native single-precision floating-point format. Compression with $\varepsilon = 10^{-6}$ increases this ratio to 5. On the other hand, from the figure it is apparent that multi-fold reduction of the file size, which is our objective, entails data loss.

3.2 Seismology simulation

Let us proceed with an example from seismology. Similarly to the analysis in the previous section, we compress an example dataset with a given ε , measure the compressed file size, then reconstruct the fields

Fig. 6 Compression of the HIT velocity components. **a** L^∞ error norm as a function of the tolerance ε ; **b** compressed file size in per cent of the original file size, as a function of the L^∞ error norm. The fit (7) is shown with a dotted line. The dashed vertical line shows the accuracy of single-precision storage and the dash-dot lines correspond to compression using quantization only

from the compressed format and measure the L^∞ error e_∞ . The example dataset is a restart file for a synthetic seismogram computation of a large earthquake using SPECfem3D, a spectral element code based on a realistic fully three-dimensional Earth model (Komatitsch et al. 2013). In the simulation, the domain is split in as many slices as the number of processors used. Data that correspond to different slices are stored in separate files, and in this example, we only process one of them. The file contains 31 single-precision arrays, but only the first 9 are large: the displacement, velocity and acceleration of the crust and mantle (3×19904633 elements each), of the inner core (3×1270129 elements each) and of the outer core (2577633 elements each). Only these arrays are compressed, while the remaining 22 small arrays containing in total 5250 single-precision numbers are directly copied in the end of the compressed file. Note that, although the grid is three-dimensional, the spectral element solver packs the variables in one-dimensional arrays. This packing preserves spatial localization; therefore, wavelet compression remains efficient. Figure 7 shows the relative compressed file size $\Sigma_1 = s(\varphi)/s(f_1)$ versus e_∞ , where $s(f_1) = 757$ MB is the original single-precision file size. The result is close to the trend obtained in the previous sections for the turbulent fluid flow, the latter shown with a dotted line.

3.3 Global-scale simulation of tropical cyclones

In this section, we evaluate the compression of restart files produced in a global weather simulation using the Multi-Scale Simulator for the Geoenvironment (MSSG), which is a coupled nonhydrostatic atmosphere–ocean–land model developed at the Center for Earth Information Science and Technology, Japan Agency for Marine–Earth Science and Technology (JAMSTEC) (Takahashi et al. 2013). Its atmospheric component

includes a large eddy simulation (LES) model for the turbulent atmospheric boundary layer and a six-category bulk cloud micro-physics model (Onishi and Takahashi 2012). Longwave and shortwave radiation transfer is taken into account using the Model simulation radiation transfer code version 10 (MstranX) (Sekiguchi and Nakajima 2008). In the global weather simulation mode, MSSG uses a yin–yang grid system in order to relax Courant–Friedrichs–Lewy condition on the sphere. Higher-order space and time discretization schemes are employed.

In a study by Nakano et al. (2017), nine tropical cyclones were simulated within the framework of the Global 7km mesh nonhydrostatic Model Intercomparison Project for improving TYphoon forecast (TYMIP-G7). The main objective of that project was to understand and statistically quantify the advantages of high-resolution global atmospheric models toward the improvement of TC track and intensity forecasts. Thus, in MSSG, each horizontal computational domain covered 4056×1352 grids in the yin–yang latitude–longitude grid system. The average horizontal grid spacing was 7 km. The vertical level comprised 55 vertical layers with a top height of 40 km and the lowermost vertical layer at 75 m. It was recognized that, when requiring the output data for every 1 or 3 h over 5-day periods be stored for analyses, the total volume of storage summed up to a huge amount.

Here, we use the initial data of July 29, 2014 at 12:00 UTC. The dataset is stored in multiple files. They include a text header file that contains descriptive parameters of the dataset such as its size, hydrodynamic field identifiers and domain decomposition parameters. The hydrodynamic field variables are stored in 1024 binary files, each containing all $n_f = 12$ fields within the same sub-domain. Thus, each sub-domain contains $n_f \times 254 \times 43 \times 55$ grid point data in double precision, totaling to 55 MB of data in one file. First 512 of these files belong to the Yin grid and the remaining 512 to the Yang grid.

Either the compression can be performed on each data subset file independently or, alternatively, the subsets can be merged before applying the wavelet transform. It is noticed in Sect. 3.1 that the compression ratio has a tendency to increase with the data size. Hence, all subsets of each hydrodynamic field are concatenated as parts of a three-dimensional array of size $4064 \times 2752 \times 55$. The fields are processed sequentially requiring 4.6 GB of RAM for the input array plus up to 6.4 GB for the encoded output data and for temporary arrays.

First, let us quantify the compression performance of this restart dataset. Figure 8a shows that the relative error e_∞ varies almost identically to the threshold ε , where e_∞ is calculated as the maximum relative error over all data points of all fields. It saturates at the level of 10^{-14} due to round-off. The compressed file size, shown in Fig. 8b, is significantly smaller than in the previously considered cases of turbulent incompressible velocity data. The linear fit

$$\Sigma_T = (-9.5 - 2.7 \log_{10} e_\infty) \times 100\%, \quad (8)$$

shown with a green dotted line, has a greater offset from the dash–dot diagonal and a less steep slope compared with the HIT fit (7), which is shown with a magenta dotted line. The global weather simulation is more complex than the incompressible Navier–Stokes, the fields contained in the restart files are heterogeneous and have sparser wavelet transform that compresses more efficiently.

Fig. 7 Compression of the seismology simulation restart file. Compressed file size in per cent of the original file size, as a function of the L^∞ error norm; The fit reference (7) is shown with a dotted line. The dash–dot line corresponds to compression using quantization only

To measure the effect of lossy restart data compression on the simulation accuracy, the following protocol was implemented.

- Compress the original restart data with some given tolerance ε ;
- Using the compressed file, reconstruct the full-size restart data;
- Restart the weather simulation and let it continue for 120 hours of physical time;
- Compare the time evolution of selected physical quantities with the original simulation not using data compression.

We focus on typhoon Halong. Figure 9a–c shows the typhoon core trajectory, minimum pressure in the core and the wind speed, respectively. The best track from observation is shown using the black color, and the result of the original simulation is shown using the red color. The typhoon trajectory is predicted well during the first three days, after that it deviates more to the north during the simulation. The predicted pressure drop and the increase in the wind speed are slightly advanced in time compared with the observation data, but the maximum wind speed is evaluated accurately.

The results of the restarts with $\varepsilon = 10^{-16}$, 10^{-6} and 10^{-4} are shown with different colors. All of them visually coincide with the original result during the first 24 h, after which the discrepancy grows large enough to be visible, but it remains in all cases significantly smaller than the difference with respect to the observation.

For $\varepsilon = 10^{-3}$ or larger, it was found impossible to restart the simulation because of numerical instability. In fact, the onset of such instability is already noticeable in the case of $\varepsilon = 10^{-4}$, as the wind speed becomes

Fig. 8 Compression of the typhoon simulation restart file. **a** L^∞ error norm as a function of the tolerance ε ; **b** compressed file size in per cent of the original file size, as a function of the L^∞ error norm. The fits (7) and (8) that correspond to Σ_{HT} and Σ_T , respectively, are shown with the magenta and green dotted lines. The dashed vertical line shows the accuracy of single-precision storage and the dash-dot diagonal line corresponds to compression using quantization only

slightly different in the very beginning of that simulation. We investigate further on this effect by plotting the difference between the restart and the original results on a logarithmic scale. We consider the L^2 error norm of the wind speed, obtained by summation over all grid points in latitude–longitude square window Ω of size $10^\circ \times 10^\circ$ centered on the typhoon as predicted in the original simulation,

$$\|U_{\text{restart}} - U_{\text{original}}\|_2 = \left(\frac{1}{\#\Omega} \sum_{p \in \Omega} (U_{\text{restart}} - U_{\text{original}})^2 \right)^{1/2}, \quad (9)$$

where U_{restart} and U_{original} is the velocity magnitude in the restart simulation and in the original simulation, respectively. For all simulations with $\varepsilon \leq 10^{-5}$, the error increases polynomially as the power ≈ 2 of the physical time after the restart point, until saturation after about 72 h. This trend arises from the nonlinear dynamics of the system and it shows no distinguishable deterministic relation with ε as long as the latter is sufficiently small. For $\varepsilon = 10^{-4}$, the error increases rapidly during the first time iterations, but then it decreases and ultimately follows the same trend as described previously. It is apparent that larger values of ε entail faster initial error growth and, ultimately, numerical divergence that cannot be accommodated by the physical model (Fig. 10).

Figure 8 shows that successful restarts belong to the intermediate regime (where Σ evolves according to equation (8)) and to the high-accuracy tail of the compression diagram, such that the compressed file size is greater than 2.5, i.e., the compression ratio is less than 40. The low-accuracy end of the diagram showing the file size of less than 1% can be practical for the data archiving only if it is not intended as input for restarting the simulation. Taking into consideration these opposing requirements of simulation reliability and data storage efficiency, it is advisable to compress the restart data of global weather simulations with the relative tolerance of $\varepsilon = 10^{-6}$.

3.4 Urban-scale simulation

In a study by Matsuda et al. (2018), a tree-crown-resolving large eddy simulation coupled with a three-dimensional radiative transfer model was applied to an urban area around the Tokyo Bay. Source terms that represent contributions of the ground surface, buildings, tree crowns and anthropogenic heat were integrated within the MSSG model in order to perform urban-scale simulations. In particular, tree crowns were taken into account using the volumetric radiosity method. The landscape was set based on geographic information system (GIS) data from the Tokyo Metropolitan Government. The initial and side-boundary atmospheric conditions were imposed by the linear interpolation of the mesoscale data provided by the Japan

Fig. 9 Results of the typhoon forecast depending on the restart file accuracy ε : **a** trajectory of the center; **b** time evolution of the minimum pressure; **c** time evolution of the maximum wind speed

Fig. 10 Time evolution of the L^2 average error in the wind speed integrated over a $10^\circ \times 10^\circ$ window, for different levels of the restart file accuracy ε

Meteorological Agency. The computational domain was a rectangular box discretized with uniform grid step (5 m) in two perpendicular horizontal directions and slightly stretched in the vertical direction (from 5 m near the sea level to 15 m in the upper layers). We consider restart data for a simulation using $N = 2500 \times 2800 \times 151$ grid points. Only the hydrodynamic fields are compressed since all other restart data use much less disk space. The domain is decomposed in equal Cartesian blocks of $50 \times 25 \times 151$ points. These datasets are stored in 5600 files, each containing 21 hydrodynamic fields. The files occupy 166 GB of disk space in total.

As shown in Fig. 11, data compression dramatically reduces the storage requirement. We have compared two approaches. The first ('united file') is to read the data from all sub-domains and concatenate in one array per field, then transform and encode each field and write all encoded data in one binary file. This is the same procedure as used in Sect. 3.3. The second approach ('divided file') consists in processing each sub-domain independently, producing 5600 compressed binary files. Since it does not require any communication between sub-domains, processing can be executed in parallel with ideal speedup. The parallel speedup comes at a price of larger compressed file size, by $\approx 2\%$ of the original data size. The sub-domain files are smaller than the united file; therefore, their compression ratio is overall lower, as explained in Sect. 3.1. In addition, part of the difference is due to the sub-domain data being normalized with the respective local maximum absolute value instead of using the global maximum. The difference may be insignificant when considering the high-accuracy end of the plot, e.g., for $\varepsilon = 10^{-14}$, Σ increases from 21 to 23%. However, when the tolerance is set to $\varepsilon = 10^{-6}$, the variation of the compressed file size from 6 to 8% of the original restart file size can be considered as relatively large. Using MPI communication for parallel wavelet transform and range coding, it may be possible to achieve better trade-off between parallel speedup and compression ratio.

We follow similar procedure as in Sect. 3.3 to evaluate the effect of lossy compression upon restart. Taking restart files from a previous simulation as initial data, three new simulations have been performed for the period of 16:00-16:10 JST (Japan Standard Time) on August 11, 2007. With the wind speed of about 2 m s^{-1} , the ten-minute time span of these simulations yields the characteristic advection length of 1.2 km, which is less than the domain size but an order of magnitude greater the size of a typical building. The first of the simulations resumes from the original files, while the second and the third resume from compressed data with $\varepsilon = 10^{-6}$ and 10^{-12} , respectively. We use the united compressed file format that introduces no artifacts at the boundaries between sub-domains. In the following discussion, we analyze the output data written on disk in the end of these simulations.

Table 1 shows the residual relative error in the L^∞ and L^2 norms, respectively, calculated as

$$e_\infty(\varepsilon) = \frac{\max_{i=1,\dots,N} |\tilde{f}_i(\varepsilon) - f_i|}{\max_{i=1,\dots,N} |f_i|} \quad (10)$$

and

$$e_2(\varepsilon) = \frac{\left[\frac{1}{N} \sum_{i=1}^N (\tilde{f}_i(\varepsilon) - f_i)^2 \right]^{1/2}}{\max_{i=1,\dots,N} |f_i|}, \quad (11)$$

Fig. 11 Compression of the urban weather simulation restart file. Compressed file size in per cent of the original file size, as a function of the L^∞ error norm. For reference, the dotted magenta line shows Σ_{HIT} (7), the dashed vertical line shows the accuracy of single-precision storage and the dash-dot diagonal line corresponds to compression using quantization only

where f denotes a hydrodynamic field obtained in the reference simulation that resumed from the original restart data, and $\tilde{f}(\varepsilon)$ stands for the respective field computed starting from the compressed data. To simplify the notation, f and $\tilde{f}(\varepsilon)$ are treated as one-dimensional arrays.

The relative L^∞ error is of order 100% in both cases for most of the field variables, except for pressure fluctuation that reaches 14% and for the base density and pressure, both of which are constant in time and therefore remain of the same order as ε or less. The relative L^2 error is, in general, two orders of magnitude smaller than the respective L^∞ error, which means that, in most part of the domain, the pointwise residual error is much smaller than the respective peak values. Interestingly, the error tends to be smaller for $\varepsilon = 10^{-6}$ than for $\varepsilon = 10^{-12}$, this may be explained by greater artificial dissipation introduced at $\varepsilon = 10^{-6}$ that slows down the convection of coherent structures.

Comparing the present results with the global numerical simulation described in Sect. 3.3, one of the key differences is in the spatial resolution. In this building-resolving simulation, the eddy turnover time of the smallest wake vortices is of order of seconds. Therefore, after 10 min (i.e., by the end of the simulation), the small structures de-correlate, producing large pointwise error.

To gain better insight, let us consider a horizontal plane at 20 m altitude above the sea level. Figure 12 shows the velocity magnitude distribution U in different cases. The top row panels (a), (b) and (c) correspond to the result at 16:10 JST of the simulation resumed from the original restart data. Panel (a) displays the entire slice while (b) and (c) zoom on selected sub-domains. The result of a restart with $\varepsilon = 10^{-6}$ is shown in Fig. 12d–f. Large-scale structures are essentially the same as in Fig. 12a–b. However, a careful comparison between panels (f) and (c) reveals significant differences on a smaller scale. To focus on such small-scale discrepancy, we calculate the difference between the velocity fields obtained with and without compression, $|\Delta U| = |U(\varepsilon = 10^{-6}) - U(\varepsilon = 0)|$, and display it in Fig. 12g–i on an exaggerated color scale. The darker tone of panels (g) and (h) suggests that $|\Delta U|$ is generally much smaller than U . There are, however, many bright spots around the buildings that mark the small-scale differences in the wake. A zoom

Table 1 Relative error after restart

Field	$e_{\infty}(\varepsilon=10^{-6})$	$e_2(\varepsilon=10^{-6})$	$e_{\infty}(\varepsilon=10^{-12})$	$e_2(\varepsilon=10^{-12})$
fl	1.1901	2.1675×10^{-2}	0.7986	5.0943×10^{-2}
fp	1.0919	2.6380×10^{-2}	1.0227	6.6958×10^{-2}
fr	1.0423	2.4495×10^{-2}	1.3071	7.1998×10^{-2}
ro	0.2440	1.6876×10^{-3}	0.5434	3.1061×10^{-3}
ps	0.1189	9.8013×10^{-4}	0.1432	6.0759×10^{-3}
rqv	0.0849	2.3355×10^{-4}	0.5417	8.5168×10^{-4}
rqq	0.4826	8.8831×10^{-4}	0.6498	1.4386×10^{-3}
surf	0.6702	2.8779×10^{-3}	0.7086	5.1263×10^{-3}

fl: longitudinal momentum; fp: latitudinal momentum; fr: altitudinal momentum; ro: density fluctuation; ps: pressure fluctuation; rqv: water vapor density; rqq: sub-grid-scale turbulence kinetic energy; surf: surface flux variable

on one of these spots displayed in Fig. 12i reveals that, locally, $|\Delta U|$ is of the same order of magnitude as U , in agreement with the global L^∞ error evaluations shown in Table 1. In addition, the error is spatially organized in a pattern characteristic of mixing layers. The vorticity plots in Fig. 13a–c show that ω_z is small in the bulk of the fast current (lower bottom corner of the sub-domain), but many strong small-scale vortices are present in the mixing layer as well as in the slow current around the buildings. Although these small vortices show qualitatively similar arrangement in Fig. 13a as in 13b, the exact position differs by as much as the core size. Consequently, the error $|\Delta\omega_z| = |\omega_z(\varepsilon = 10^{-6}) - \omega_z(\varepsilon = 0)|$ is a superposition of strong well-localized peaks. It is worth mentioning that a restart with $\varepsilon = 10^{-12}$ has led to very similar results. This is expected as small-scale vortices shed from the buildings evolve rapidly and chaotically. Exact deterministic repetition of such simulation requires that the initial data be exact. On the other hand, the initial error has negligible effect on a kilometer scale.

4 Compression performance with consideration of computational cost

Our intention is to compress data from numerical simulations. Let us first consider a scenario when the same computer is used for the simulation and for the data compression. Writing the output data in a divided file, as explained in Sect. 3.4, enables parallel execution of the compression program, albeit at a cost of slight increase in the compressed file size. Simulations often use time-marching schemes, each time iteration incorporates differential operators of linear computational complexity (i.e., proportional to the number of grid points N) in the case of, e.g., evaluating derivatives using finite differences or having an even higher complexity (e.g., $N \log N$ or N^2) if spectral methods are used or linear systems are to be solved. Although wavelet transform and range coding are known to be relatively expensive operations, their computational complexity is linear. This means that the number of arithmetic operations necessary to compress one snapshot of the output data is, at worst, proportional to the number of arithmetic operations required for one time iteration in the simulation. The proportionality constant may be greater than one if the simulation only uses explicit low-order schemes and the right-hand side of the evolution equation is simple enough. However, typical practical problems in scientific computing are computationally more intense than compression using one pass of wavelet transform followed by quantization and range coding. Besides that, it is rare to write on disk the output after every time step. Temporal sampling is commonly used (Li et al. 2018), which dramatically reduces the data compression cost in comparison with the simulation cost.

A different scenario would be to perform a simulation on a supercomputer and compress/decompress the result files on a desktop computer. The elapsed time of data compression can be long in such situation, and that is the case we focus on in this section. For the performance analysis, we use the HIT dataset described in Sect. 3.1. The simulation was performed on the Earth Simulator supercomputer system (NEC SX-ACE), while the compression/decompression analyzed on an HP Z640 workstation with two Intel Xeon E5-2620v4 8-core CPUs at 2.10 GHz clock rate, 8×8 GB DDR4-2400 RAM and a RAID 1 pair of Seagate 2 TB 7200 RPM SATA HDDs. A succinct discussion of the WaveRange software performance and optimization aspects will be followed by a comparison with other open-source libraries.

4.1 Performance assessment using a roofline model

We used Intel Advisor 2019 Update 4 (build 594873) for the performance analysis. WaveRange was built using gcc 7.4.0 with $-O2$ level optimization and AVX2 support, under the Ubuntu 18.04.1 operating system. The x -velocity component of the HIT dataset was read from a file in the FluSI HDF5 format, compressed with $\varepsilon = 10^{-6}$ and written on disk also using HDF5. Then it was decompressed. The compression and the decompression programs have been profiled separately. Therefore, the results are presented in two columns in Table 2 and in two panels in Fig. 14.

In the compression procedure, the wavelet transform is the least time-consuming step. This can be explained by the relatively large fraction of vector operations in it: there are 12 vector loops and 12 scalar kernels shown in Fig. 14 with the red and the blue circles, respectively. Two of them achieve the L1 bandwidth bound. The need for strided memory access is the main factor that limits further optimization of the transform (these loops are marked with blue circles situated below the DRAM bandwidth line). In contrast, only 1 of the 8 range coder kernels has been vectorized, and the most time-consuming one is the function that encodes a symbol using frequencies. This explains why range coding is the most time-consuming step. Quantization and other parts of the program count 2 vector loops and 4 scalar kernels. In

Fig. 12 Wind speed distribution in the urban weather simulation, on a plane at 20 m altitude above the sea level. Left column shows the entire horizontal span of the computational domain, middle column shows a zoom on a selected $4\text{ km} \times 4\text{ km}$ area, right column shows a deeper zoom on a $400\text{ m} \times 400\text{ m}$ area. **a, b, c** Restart from the original initial data, results shown for 16:10 JST. **d, e, f** Restart from the compressed input data using $\varepsilon = 10^{-6}$, results shown for 16:10 JST. **g, h, i** Difference between the results using the original and the compressed initial data. Note that the color scale of the bottom row figures has been adjusted to accentuate on the regions of large $|\Delta U|$

Fig. 13 Vertical vorticity component at 20 m altitude above the sea level, in a $400\text{ m} \times 400\text{ m}$ sub-domain that corresponds to the right column of Fig. 12. All data are for 16:10 JST. **a** Restart from the original initial data, **b** restart from the compressed input data using $\varepsilon = 10^{-6}$, **c** difference between the results using the original and the compressed initial data

the quantization process, type conversion from double to char presents difficulties for the automatic optimization. The overall time in vectorized loops amounts to 11% of the total CPU time.

Decompression with WaveRange takes longer time than compression. Although the inverse wavelet transform has 15 vector and 6 scalar kernels, its execution takes longer than the forward transform used in the compression. Range decoding is by far the most expensive part, as it amounts to 61% of the decompression CPU time. Only 1 of the 10 range coder kernels has been vectorized. The most time-consuming function is the one that calculates cumulative frequency of the next symbol. Dequantization is relatively cheap: it only takes 13% of the total CPU time. The overall vector time ratio is 12%, which is nearly the same as in the compression program.

4.2 Comparison with other methods

Baselines in terms of lossless compression achievable with general-purpose utilities are provided in Appendix E, showing the smallest relative compressed file size of 81.2%. This means that many-fold data reduction cannot be achieved without loss of information in this example.

The performance of different lossy compression methods, including WaveRange, is displayed in Fig. 15, which shows Σ , θ_c and θ_d as functions of the L^∞ error norm. Three other lossy methods have been evaluated. Scale–offset is an HDF5 filter that performs a scale and offset operation, then truncates the result to a minimum number of bits. Scaling is performed by multiplication of each data value by 10 to the power of a given scale factor, which can be adjusted in order to reach the desired compressed file size or accuracy. SZ-1.4 (Tao et al. 2017) and ZFP (Diffenderfer et al. 2019) are two state-of-the-art scientific data compressors. The former is based on the Lorenzo predictor and the latter uses a custom transform that operates multi-dimensional blocks of small size. Both are implemented as HDF5 filters and both can operate in a fixed error mode. To obtain the plots in Fig. 15, we followed the same procedure as described in the previous sections: the HIT x -velocity data file was compressed, the file size was measured, then it was decompressed and the L^∞ error was measured. In addition, the execution time was measured for the compression and for the decompression separately.

Let us first discuss the compression performance in terms of Σ . The scale–offset compression generally produces larger files than WaveRange. Note that this method is equivalent to quantization (3) in the limit of large L^∞ error. ZFP produces slightly larger files than WaveRange, except when the set tolerance is smaller than 10^{-14} and the L^∞ error of WaveRange increases significantly due to round-off. The performance of SZ-1.4 varies largely depending on the control parameter configuration. In our test, we optimized it for the maximum compression ratio at a given error magnitude. Thus, the maximum quantization interval number was increased as the error decreased. With this setting, SZ-1.4 could produce smaller files than WaveRange, but only for the relative error greater than 10^{-8} . SZ-1.4 switched to lossless mode when the relative error smaller than 10^{-10} was requested.

Considering the throughput θ_c and θ_d , ZFP was generally the fastest in our tests, although it was outperformed by SZ-1.4 in certain cases by a small amount. WaveRange was 3 times slower than ZFP for the compression and 5 times slower for the decompression. This is not surprising, considering that the transform used in ZFP was optimized for the maximum throughput at the cost of a certain decrease in the compression ratio. Interestingly, despite the relatively high algorithmic complexity, WaveRange compression was faster than the scale–offset compression.

5 Conclusions

A wavelet-based method for compression of data output from numerical simulation of fluid flows using block Cartesian grids has been presented. The method consists of a discrete wavelet transform, quantization

Table 2 Performance characteristics of WaveRange as applied to the HIT dataset

	Compression	Decompression
Total CPU time	20.63 s	24.06
Wavelet transform	23%	26%
Range coding	41%	61%
Other (incl. quantization)	36%	13%
Time in vectorized loops	11%	12%
Time in scalar code	89%	88%

Fig. 14 Roofline model. **a** Compression performance; **b** decompression performance. Circles show loops that belong to the wavelet transform routine, squares correspond to range encoding/decoding and diamonds belong to quantization/dequantization or the rest of the program. Blue markers show scalar loops and red markers show vectorized loops. The size of the marker and its line width signifies its relative contribution to the program total elapsed time: less than 1% for the smallest markers, between 1% and 15% for the medium-size markers and more than 15% for the large markers. Markers may overlap

adapted for floating-point data and entropy coding. It is designed such as to guarantee the desired pointwise reconstruction accuracy. An open-source software implementation has been provided, see <https://github.com/pseudospectators/WaveRange>.

The data compression properties have been analyzed using example numerical tests from different kinds of problems, from idealized fluid flows to realistic seismology and weather simulations. In particular, it is found that, in the most challenging (from the compression point of view) case of homogeneous isotropic turbulence, compression allows to reduce the data storage by a factor of 3 using $\varepsilon = 10^{-8}$, which is significantly better than storing the same data in single-precision floating-point format. The method shows favorable scaling with the data size, i.e., greater compression ratios are achieved for larger datasets. Compression of the wake turbulence is also slightly better, compared to the reference homogeneous isotropic turbulence case. This is explained by greater inhomogeneity of the wake turbulence, which means that there are fewer large wavelet coefficients.

The compression performance depends on the flow type. For the realistic data generated from global and urban weather simulations, the file size can be reduced by a factor of about 15 using the threshold value $\varepsilon = 10^{-6}$. Both simulations can be successfully restarted from the reconstructed data. The reconstruction error has shown no significant effect on the dynamics of large-scale structures, which are typically the main

Fig. 15 Comparative performance of different scientific data compression methods. **a** Compressed file size in per cent of the original file size. The inset shows a zoom on the interval where WaveRange shows improvement over the other methods; **b** compression throughput; **c** decompression throughput. The dashed vertical line shows the accuracy of single-precision storage and the dash-dot lines correspond to compression by quantization only. The HIT x -velocity file is compressed and decompressed

objects of interest. It should be noted, however, that the small-scale structures may randomize very quickly if any small initial error is introduced in the simulation.

Appendix

Three-dimensional wavelet transform

WaveRange applies one level of the 1D wavelet transform in the x direction, then in y and finally in z . After that, it moves to the next level: The approximation coefficients on the first level are taken as input data for 1D transforms. Then, the same procedure is repeated on the second level, etc., as explained in Algorithm 2.

Data: 3D array of point values in physical space f_{i_x, i_y, i_z} , $i_x = 1, \dots, n_x$,

$$i_y = 1, \dots, n_y, i_z = 1, \dots, n_z.$$

Result: approximation coefficients and detail coefficients placed in the same array

$$F_{i_x, i_y, i_z}.$$

$$m_x \leftarrow n_x; m_y \leftarrow n_y; m_z \leftarrow n_z;$$

$$\{F_{i_x, i_y, i_z}\} \equiv \{f_{i_x, i_y, i_z}\};$$

for $l = 1$ to L **do**

if $m_x > 1$ **then**

for $i_y = 1$ to m_y ; $i_z = 1$ to m_z **do**

 copy data elements F_{j_x, i_y, i_z} in a contiguous array g_{j_x} , where $j_x = 1, \dots, m_x$;

 apply the 1D wavelet transform (1) and copy the result in F_{j_x, i_y, i_z} ;

end

end

if $m_y > 1$ **then**

for $i_x = 1$ to m_x ; $i_z = 1$ to m_z **do**

 copy data elements F_{i_x, j_y, i_z} in a contiguous array g_{j_y} , where $j_y = 1, \dots, m_y$;

 apply the 1D wavelet transform (1) and copy the result in F_{i_x, j_y, i_z} ;

end

end

if $m_z > 1$ **then**

for $i_x = 1$ to m_x ; $i_y = 1$ to m_y **do**

 copy data elements F_{i_x, i_y, j_z} in a contiguous array g_{j_z} , where $j_z = 1, \dots, m_z$;

 apply the 1D wavelet transform (1) and copy the result in F_{i_x, i_y, j_z} ;

end

end

$$m_x \leftarrow \lceil m_x/2 \rceil; m_y \leftarrow \lceil m_y/2 \rceil; m_z \leftarrow \lceil m_z/2 \rceil;$$

end

Algorithm 2: Three-dimensional wavelet transform.

Compression ratio and its dependence on the size of the dataset

Compression ratio

$$r = \frac{s(f)}{s(\varphi)} \quad (12)$$

is a commonly used performance metrics, particularly suitable in those situations when the compressed file size is many times smaller than the original size. Thus, Fig. 16 reveals that hundredfold compression is attainable by setting $\varepsilon = 10^{-2}$. This may be a good setting for the purpose of qualitative flow visualization, for example, as illustrated in Fig. 5a: vortex filaments consist of many small cyan and magenta patches, which means that two iso-surfaces coincide almost perfectly, namely, the cyan iso-surface that visualizes the vorticity magnitude calculated from the original velocity, and the magenta iso-surface that is obtained using the velocity field reconstructed from the compressed data with $\varepsilon = 10^{-2}$. Figure 5b confirms that the error mainly contaminates the smallest scales, while the energy-containing large scales are much less affected. This observation supports the interpretation of quantization as adding thermal noise to the original data. Note that we chose $\varepsilon = 10^{-2}$ to magnify the numerical error. For $\varepsilon = 10^{-6}$ or less, the reconstructed spectrum would be visually identical with the original. Finally, Fig. 16 confirms that asymptotic scaling (7) is realized as soon as $e_\infty < 10^{-3}$.

Let us now consider how the compression ratio r scales with the size of the dataset. There are at least two obvious ways to obtain a smaller HIT dataset from the original 512^3 arrays: down-sampling and sub-domain extraction. The former means that only every 2nd (or 4th, etc.) grid point in each direction is retained, and all other data points are discarded. The latter means that only the first 256 (or 128, etc) points in each direction are retained. We have thus constructed two sequences of datasets of size $n^3 = 32^3, 64^3, 128^3, 256^3$ and 512^3 , where the largest dataset is the original one. The compression ratio using tolerance $\varepsilon = 10^{-6}$ is displayed in Fig. 17a, b.

Figure 17a shows r of the down-sampled data. Starting from the rightmost point, there is a major decrease in r between $n = 512$ and 256 , i.e., after discarding every second point. It is followed by a slower decay that accompanies subsequent down-sampling. The original dataset is a result of direct numerical simulation of turbulence, which implies that the inertial range is fully resolved, i.e., the distance between the neighboring grid points is smaller than the Kolmogorov scale. Consequently, the turbulent velocity fluctuation at the smallest scale contains disproportionately less energy than any larger scale in the inertial range. This means that the numerical values of the smallest scale wavelet coefficients contain much less nonzero significant digits. Therefore, they are compressed much more efficiently. In the inertial range, the number of nonzero significant digits becomes larger as the wavenumber decreases, which explains further gradual decrease of r with decreasing n . The evolution of r with n is thus related to the decay of the Fourier energy spectrum.

Figure 17b shows r of the sub-domain data. In this case, the energy per unit volume of the smallest scale velocity fluctuation field does not depend on n . Hence, smallest scale wavelet coefficients are of the same order of magnitude, regardless of n . As a result, in this case, the compression ratio r varies less with n than in the previous case. It is practically constant, $r = 5.5$, between $n = 128$ and 512 . There is, nevertheless, some moderate decrease down to $r = 4.5$ at $n = 32$. It can be explained by the efficiency of the entropy coding becoming lower as the dataset becomes smaller.

These scalings confirm that the chosen method of compression is particularly suitable for datasets produced by large-scale, high-resolution numerical simulations. In addition, the largest compression ratio will be achieved if the data are stored in one single file rather than divided in multiple sub-domain files treated independently. Further, Appendix C investigates into the effects of spatial inhomogeneity by considering fluid flow past a solid cylinder.

Wake turbulence

The fluid velocity field in this case is obtained from a numerical simulation of viscous incompressible flow past a periodic array of circular cylinders. The flow configuration is similar to the numerical wind tunnel with a cylindrical obstacle considered by (Ravi et al. 2016). Here it is described in dimensionless units. The fluid domain is a rectangular box with length of 10, width of 8 and height of 4. The boundary conditions on the exterior faces of the domain are periodic in all the three directions. In addition to that, a vorticity sponge condition is applied over the 48 grid slabs adjacent to the outflow boundary, to avoid the wake reentering the domain. The cylinder immersed in the fluid has the diameter of 1.894. Its axis is oriented vertically and it is located at 1.5 length units downstream from the upstream boundary of the periodic domain. The fluid has the kinematic viscosity of 0.001 and density of 1. Mean inflow velocity of 1.246 is imposed.

The computational domain is discretized using a uniform Cartesian grid consisting of $960 \times 768 \times 384$ points. The no-slip boundary condition at the surface of the cylinder is modeled using the volume

penalization method with the penalization parameter $C_\eta = 5 \cdot 10^{-4}$. For more information about the numerical method, see Engels et al. (2016).

A small cylindrical detail attached on the surface of the cylinder at the angular distance of 135° from the front stagnation point served to quickly break the bilateral symmetry of the flow. Subsequently, random noise introduced during the startup phase provoked three-dimensional instabilities. The three components of the velocity field (u, v, w) at time $t = 330$ were saved, respectively, in three separate files in double precision. Each file occupied 2.2 GB of hard disk space.

The flow configuration and the result of the simulation are visualized in Fig. 18. The wake is apparently turbulent with a variety of scales and heterogeneity reminiscent of industrial flows. Grey color shows the cylinder, cyan shows an iso-surface of the vorticity magnitude calculated using the original velocity field, and magenta shows a similar iso-surface for $\varepsilon = 10^{-2}$. The two iso-surfaces overlap.

Similarly to the previous HIT test case, the procedure of compression with prescribed tolerance ε and subsequent reconstruction has been applied to the velocity components of the turbulent wake. Again, the relative error measured in L^∞ norm has appeared almost identical to ε , except for the smallest and for the largest ε . Figure 19a shows the compressed file size as a function of the error norm, componentwise. The compressed file size is again normalized with the original file size. The error norm is defined by (5). The compression method is equally effective for all velocity components, despite the anisotropy of the velocity field. Perfect reconstruction cannot be reached because of round-off errors, but reconstruction with 10^{-14} accuracy is achieved from a compressed file slightly larger than one half of the original size. It can also be noticed that, if the velocity field were stored in a single-precision file, the error would be 100,000 times larger than when storing the same field using the compressed format in an equally large file. Similarly to Fig. 6b for the HIT test case, the diagonal dash-dot line in Fig. 19a shows the gain in compression achieved by discarding the least significant digits, i.e., by reducing the precision of each point value. The difference between this upper bound and the actually measured file size is the combined effect of wavelet transform and entropy coding.

The compression ratio as a function of the error norm is shown in Fig. 19b. Greater compression ratios can be achieved when precision requirements are less stringent. For instance, in this example one can achieve 8 times reduction in the volume of data if stored with 10^{-6} accuracy. For very low accuracy, the compression ratio saturates at $r \approx 400$, which is close to the maximum compression ratio obtained for the HIT data.

In the intermediate range of ε , the compressed file size as a function of the relative L^∞ error can be approximated as

$$\Sigma_{Cyl} = (-0.19 - 0.053 \log_{10} e_\infty) \times 100\%, \quad (13)$$

which is shown in Fig. 19a using green dots. It can be compared with Σ_{HIT} given by (7), which is superposed on the same figure using magenta dots. The values of Σ_{Cyl} are smaller than those of Σ_{HIT} . This can be explained by comparing the histograms displayed in Fig. 20. Only the first component, u , is shown for clarity. The results for v and w are similar.

To guarantee fair comparison between two histograms for the two different flow fields, u is normalized by its half-span before applying the wavelet transform yielding $\tilde{u}$. The same normalization is used in the data

Fig. 16 Compression of the HIT velocity components. Compression ratio versus relative L^∞ error. The dashed vertical line shows the accuracy of single-precision storage and the dash-dot lines correspond to compression using quantization only

Fig. 17 Compression ratio r of (a) uniformly down-sampled velocity datasets containing n^3 out of 512^3 grid points; (b) velocity data in sub-domains of size n^3

compression algorithm. As the number of bits required to represent a point value of $\tilde{u}$, after quantization, is proportional to $\log_{10} |\tilde{u}|$, the latter quantity is used to produce the histogram. The interval between its minimum and maximum is divided in a finite number of bins and the number of point values falling in each bin is counted. Note that the maximum values are almost identical for both datasets. The result is normalized such that the area under the curve integrates to 1.

By comparing the histograms for the HIT and the wake velocity datasets, one can see, for example, that the HIT field has relatively many coefficients of order of magnitude 10^{-2} , but less at 10^{-6} . The expected value for the HIT case is -4 , whereas in the cylinder wake case it is equal to -5.1 . The standard deviation is similar in both cases: 1.3 and 1.1, respectively. It follows that the HIT wavelet coefficients are, on average, almost one order of magnitude larger than the cylinder wake wavelet coefficients. Consequently, for equal compression ratio, the L^∞ error is expected to be one order of magnitude larger for the HIT data than for the cylinder wake. This is in agreement with the observed difference between the linear fits in Fig. 19a. In addition, the slightly larger skewness of the cylinder wake PDF explains why the difference becomes slightly smaller when the tolerance is decreased - also compare the slopes of (7) and (13).

WaveRange treats individual components of a vector field independently. Although it must be possible to exploit correlation between multiple scalar fields, this procedure is not straightforward. We have tested two approaches. The results are compared in Table 3. The ‘velocity polar’ method consists in transforming the velocity components u , v and w to a magnitude ρ and two angles ϑ and ϕ such that

$$u = \rho \sin \vartheta \cos \phi, \quad v = \rho \sin \vartheta \sin \phi, \quad w = \rho \cos \vartheta. \quad (14)$$

The scalar fields ρ , ϑ and ϕ are then compressed with a prescribed tolerance ϵ . The ‘vorticity Cartesian’ method calculates the vorticity $\boldsymbol{\Omega} = \nabla \times \mathbf{U}$, which is the curl of the velocity vector $\mathbf{U} = (u, v, w)$ and the spatial average of the velocity $\mathbf{U}_0$. Then, the three components of the vorticity are compressed with tolerance ϵ . The velocity is reconstructed from $\boldsymbol{\Omega}$ and $\mathbf{U}_0$ using the Biot–Savart formula. All differential and integral operators are approximated using a Fourier spectral method.

Both methods introduce error due to additional computations. For the ‘velocity polar’ method, computation entails a larger round-off error than in the original ‘velocity Cartesian’ case. The ‘vorticity

Fig. 18 Flow visualization of the turbulent wake simulation. Black lines show the extent of the computational domain. The vorticity magnitude is visualized using an iso-surface at 10% of the maximum value. Iso-surfaces of the original field (cyan) and reconstructed after compression with $\varepsilon = 10^{-2}$ (magenta) are superimposed and overlap almost perfectly. The surface of the cylinder is colored in grey

Cartesian' method involves numerical differentiation which has a truncation error. For these reasons, we select relatively large compression tolerance ε in order to achieve the reconstruction error of approximately 10^{-4} . It is much larger than the truncation and the round-off errors.

To quantify the reconstruction accuracy of a vector field using a single scalar-valued metric, the maximum relative L_∞ error norm is selected among the three velocity components,

$$e_{vec} = \max \left(\frac{\|\tilde{u} - u\|_\infty}{\|u\|_\infty}, \frac{\|\tilde{v} - v\|_\infty}{\|v\|_\infty}, \frac{\|\tilde{w} - w\|_\infty}{\|w\|_\infty} \right) \quad (15)$$

The compressed file size Σ_{vec} is calculated as the sum of the compressed file sizes divided by the original storage size of the three velocity components in double precision, and the compression ratio is equal to $r_{vec} = 1/\Sigma_{vec}$.

Note that, in all methods, e_{vec} is calculated on the reconstructed Cartesian velocity components $\tilde{u}$, $\tilde{v}$, $\tilde{w}$, but the tolerance ε is set on the transformed field in the 'velocity polar' and 'vorticity Cartesian' methods. Therefore, in the two latter cases, ε is iteratively varied until e_{vec} becomes close to 10^{-4} . The final values of e_{vec} are shown in the last column of Table 3, and the corresponding ε is in the second last column. From the values of Σ_{vec} and r_{vec} in, respectively, the second and the third columns in Table 3 we conclude that the original 'velocity Cartesian' method provides the best compression for the desired 10^{-4} accuracy.

Such inefficiency of the 'velocity polar' and 'vorticity Cartesian' representations can be explained by the loss of uniformity in the spatial distribution of pointwise errors, and unequal errors of $\tilde{u}$, $\tilde{v}$ and $\tilde{w}$. This implies that some point values are stored with higher precision than necessary. In addition, the 'velocity polar' representation suffers from the discontinuity in ϑ and ϕ artificially introducing small scales in the field, which require more wavelet coefficients to be stored. To find a suitable vector field transform may be a promising direction for the future work.

Reliability of the error control

The wavelet transform offers control over the L^2 norm, but we aim for the L^∞ error control. The pointwise error in the physical space is a linear combination of the quantization errors of the wavelet coefficients within the stencil. The latter are all smaller than ε_F by construction. The weights are constants specific to the CFD9/7 wavelet. The number of terms in the sum is bounded by the number of operations in the lifting steps times the number of spatial dimensions ($= 3$) times the number of levels of the transform ($= 4$). From these considerations, one may derive a theoretical upper bound on the ratio between the L^∞ error in physical space and the filtering threshold ε_F , and thus obtain a theoretical estimate for the parameter η in (4). However, we

Fig. 19 Compression of the wake velocity components. **a** Compressed file size in per cent of the original file size, as a function of the L^∞ error norm; **b** compression ratio versus the relative L^∞ error. The fits (7) and (13) that correspond to Σ_{HIT} and Σ_{Cyl} , respectively, are shown with the magenta and green dotted lines. The dashed vertical line shows the accuracy of single-precision storage and the dash-dot lines correspond to compression using quantization only

Fig. 20 Histogram of the order of magnitude of the wavelet coefficients of u velocity component in two different test cases

do not attempt such rigorous analysis. The value $\eta = 1.75$ is an empirical constant. Moreover, it does not exactly guarantee that $e_\infty = \varepsilon$.

In order to gain quantitative information about reliability of the error control in WaveRange across different application scenarios, we compiled the L^∞ error data e_∞ from all examples (HIT, wake turbulence, typhoon, urban-scale and seismology simulations) with different tolerance values ε into one dataset. The

Table 3 Evaluation of the effect of derived representations of the velocity vector field

	Σ_{vec} (%)	r_{vec}	ϵ	e_{vec}
Velocity Cartesian (base)	3.3	30.6	10^{-4}	9.5×10^{-5}
Velocity polar	4.3	23.2	3×10^{-4}	9.9×10^{-5}
Vorticity Cartesian	5.2	19.1	3.5×10^{-5}	10.1×10^{-5}

error-to-tolerance ratio e_∞/ϵ was then calculated for each sample in the dataset. For the HIT and the wake turbulence, e_∞ was taken separately for each velocity component. For the typhoon, urban-scale and seismology simulations, the maximum e_∞ over all fields was taken. For the urban-scale simulation, the united and the divided storage schemes were both included in the analysis. This yielded a set of 186 values of e_∞/ϵ . We then discarded those samples that corresponded to $\epsilon \leq 10^{-14}$, for which the error control failed because of the round-off errors. The remaining set contained 138 samples. We then calculated a histogram of e_∞/ϵ . The result is shown in Fig. 21.

The peak of the histogram is at $e_\infty/\epsilon = 1$. The expected value of e_∞/ϵ is estimated as 0.93, and the standard deviation is equal to 0.20. By numerical integration of $PDF(e_\infty/\epsilon)$ we found that, in 83% of the cases included in the analysis, e_∞/ϵ was less or equal 1. In 94% cases it was less or equal 1.2 and in no case it exceeded 1.7.

Lossless compression of the HIT dataset

In addition to comparing WaveRange with other methods, it is important to provide baselines in terms of lossless compression achievable with general-purpose utilities. This is summarized in Table 4. Here we use the x -velocity component of the HIT dataset.

DEFLATE (Deutsch 1996) is an algorithm widely used for general purposes. The best compression, i.e., the smallest compressed file size, is achieved when the level of compression is set to 9. The execution time can be minimized by setting the level of compression to 1. It is customary to apply shuffling as a preconditioner to facilitate floating-point data compression with DEFLATE. The idea of shuffling is to break apart floating-point elements of an array into their mantissa and exponent components, then change the byte order in the data stream such as to place the first byte of every element in the first chunk, then the second byte of every element in the second chunk, etc. In the scientific datasets, values at neighboring points are usually close to each other. Therefore, shuffling produces a data stream that includes many continuous subsequences of identical entries, which DEFLATE can compress well. From this point of view, shuffling is similar to the quantization described in Sect. 2.3.

The last two lines in Table 4 show the performance of Szzip and 7-Zip. Szzip (Yeh et al. 2002) is an implementation of the extended-Rice lossless compression algorithm designed for use with scientific data. In our test, we activated the optional nearest neighbor coding method. 7-Zip applies the LZMA method (Pavlov 2019), and we set the level of compression to 9, which is the maximum level. The crosses in Fig. 6(b) correspond to the compressed file size obtained with this method. All algorithms have been applied as HDF5 filters (HDF5 version 1.10.0-patch1), with the exception of 7-Zip, which was used as a standalone application (version 16.02).

Fig. 21 Histogram of the error-to-tolerance ratio

Table 4 Lossless compression of the HIT data: relative compressed file size Σ , compression ratio r , compression throughput θ_c and decompression throughput θ_d

	Σ	r	θ_c	θ_d
DEFLATE, fastest	96.3%	1.038	23.2 MB/s	144.1 MB/s
DEFLATE, best	95.6%	1.046	16.9 MB/s	153.0 MB/s
Shuffling and DEFLATE, fastest	82.6%	1.211	31.0 MB/s	282.6 MB/s
Shuffling and DEFLATE, best	81.2%	1.231	12.7 MB/s	290.2 MB/s
Szip	84.4%	1.184	61.3 MB/s	124.8 MB/s
7-Zip	87.4%	1.144	1.6 MB/s	17.0 MB/s

The performance metrics are the relative compressed file size Σ as defined by (6), compression ratio r as defined by (12), compression throughput $\theta_c = s(f)/t_c$ and decompression throughput $\theta_d = s(f)/t_d$, where $s(f) = 1$ GB is the storage space required for a 512^3 array of double-precision values, t_c is the CPU time for compression and t_d is the CPU time for decompression. DEFLATE with the lowest level of compression reduced the file size by only a very small amount, down to 96.3% of the original size. Switching to the maximum level of compression helps to gain additional 0.7%, but shuffling the data brings a dramatic improvement, allowing to reach 81.2%, which is the best result among all lossless methods considered here. Szip is slightly less efficient from the file size reduction viewpoint, but its compression throughput is higher. However, the decompression throughput is lower for Szip than for DEFLATE. 7-Zip turns out to be less efficient by all metrics, which is not surprising, since it is not designed for use with scientific data. Overall, the compression ratio is 1.231 at most.

Software framework

This appendix contains an overview of the software. Installation instructions and user's notes can be found in <https://github.com/pseudospectators/WaveRange/blob/master/README.md>

Software functionalities

WaveRange application includes an encoder that compresses the floating-point data and a decoder that reconstructs the data from the compressed format. Two different executables are built for these two purposes, respectively, when WaveRange is used as a standalone application. WaveRange's 'generic' interface can read and write plain Fortran and C/C++ files containing one or several multi-dimensional arrays. In addition, the current version of WaveRange includes 'specialized' input and output interfaces compatible with a spectral incompressible Navier–Stokes solver FluSI (Engels et al. 2016) and with MSSG, the Multi-Scale Simulator for the Geoenvironment (Takahashi et al. 2013). Each of these two solvers generate two types of large files: regular output for, e.g., flow visualization, and restart files. WaveRange can compress both. The input floating-point data for compression may be stored in one file or divided in multiple files, in a format specific to the computation software employed.

WaveRange can also be used as a library. In that case, the encoder and the decoder functions are called from the user's program. The original, the encoded and the reconstructed datasets are passed as parameters. Disk input/output is to be implemented by the user.

Software architecture

WaveRange is written in C and C++. The lower-level functions related with the discrete wavelet transform and range coding are in C. Their source codes are stored in the directories `waveletcdf97_3d/` and `rangecod/`, respectively. On a higher level, the C++ code in `core/` implements the functions `void encoding_wrap` and `void decoding_wrap` that execute all operations necessary for compression and reconstruction, respectively, in accordance with the data flow diagram in Fig. 1. These functions are called after the input data files are read and before the output files are written. Two similar functions, `void encoding_wrap_f` and `void decoding_wrap_f`, offer compatibility with Fortran. The `int main` functions are specific to each interface. The respective C++ source files are contained in the directories `generic/`, `flusi/` and `mssg/`, the executables are generated in `bin/generic/`, `bin/flusi/` and `bin/mssg/`. Note that the FluSI interface requires the hierarchical data format (HDF5) library (Folk et al.

1999), but the generic and MSSG interfaces have no external dependencies. The WaveRange library files appear in bin/lib/.

Acknowledgements This work is supported by the FLAGSHIP2020, MEXT within the Post-K Priority Issue 4 (Advancement of meteorological and global environmental predictions utilizing observational ‘Big Data’). The authors thank Mr. Koji Goto and Dr. Keigo Matsuda for their help with handling the global- and urban-scale simulations, and Prof. Seiji Tsuboi for providing the seismology simulation data.

Data availability The HIT and the wake turbulence datasets, in the compressed format, are available at <https://osf.io/pz4n8/>. Access to the weather simulation data can be granted upon request, under a collaborative framework between JAMSTEC and related institutes or universities.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Baker AH, Hammerling DM, Mickelson SA, Xu H, Stolpe MB, Naveau P, Sanderson B, Ebert-Uphoff I, Samarasinghe S, De Simone F, Carbone F, Gencarelli CN, Dennis JM, Kay JE, Lindstrom P (2016) Evaluating lossy data compression on climate simulation data within a large ensemble. *Geosci Model Dev* 9(12):4381–4403. <https://doi.org/10.5194/gmd-9-4381-2016>
- Bakke H, Frogner L, Gudiksen BV (2018) Non-thermal electrons from solar nanoflares - in a 3D radiative MHD simulation. *Astron Astrophys* 620:L5. <https://doi.org/10.1051/0004-6361/201834129>
- Balajewicz MJ, Dowell EH, Noack BR (2013) Low-dimensional modelling of high-reynolds-number shear flows incorporating constraints from the navier-stokes equation. *J Fluid Mech* 729:285–308
- Balsa Rodríguez M, Gobetti E, Iglesias Guitián JA, Makhinya M, Marton F, Pajarola R, Suter SK (2014) State-of-the-art in compressed GPU-based direct volume rendering. *Comput Graph Forum* 33(6):77–100. <https://doi.org/10.1111/cgf.12280>
- Baudron AM, Lautard JJ (2007) MINOS: a simplified Pn solver for core calculation. *Nucl Sci Eng* 155(2):250–263. <https://doi.org/10.13182/NSE07-A2660>
- Berkooz G, Holmes P, Lumley JL (1993) The proper orthogonal decomposition in the analysis of turbulent flows. *Ann Rev Fluid Mech* 25(1):539–575
- Beyer J, Hadwiger M, Pfister H (2014) A survey of GPU-based large-scale volume visualization. In: Borgo R, Maciejewski R, Viola I (eds) *EuroVis—STARs*, <https://doi.org/10.2312/eurovisstar.20141175>
- Bi C, Ono K, Yang L (2014) Parallel POD compression of time-varying big datasets using m-swap on the K computer. In: 2014 IEEE International congress on big data, pp 438–445
- Bradley JN, Brislawn CM (1993) Wavelet transform-vector quantization compression of supercomputer ocean models. In: [Proceedings] DCC ’93: Data compression conference, pp 224–233. <https://doi.org/10.1109/DCC.1993.253127>
- Daubechies I, Sweldens W (1998) Factoring wavelet transforms into lifting steps. *J Fourier Anal Appl* 4(3):247–269
- Deutsch P (1996) DEFLATE compressed data format specification version 1.3. Tech. rep
- Diffenderfer J, Fox A, Hittinger J, Sanders G, Lindstrom P (2019) Error analysis of ZFP compression for floating-point data. *SIAM J Sci Comput* 41(3):A1867–A1898. <https://doi.org/10.1137/18M1168832>
- Duval L, Røsten T (2000) Filter bank decomposition of seismic data with application to compression and denoising. In: SEG Annual International Meeting, Soc. Expl. Geophysicists, pp 2055–2058. <https://doi.org/10.1190/1.1815847>
- Engels T, Kolomenskiy D, Schneider K, Sesterhenn J (2016) FluSI: a novel parallel simulation tool for flapping insect flight using a Fourier method with volume penalization. *SIAM J Sci Comput* 38:S3–S24
- Farge M (1992) Wavelet transforms and their applications to turbulence. *Ann Rev Fluid Mech* 24(1):395–458
- Folk M, Cheng A, Yates K (1999) HDF5: a file format and I/O library for high performance computing applications. *Proc Supercomput* 99:5–33
- Fornari W, Formenti A, Picano F, Brandt L (2016) The effect of particle density in turbulent channel flow laden with finite size particles in semi-dilute conditions. *Phys Fluids* 28(3):033301. <https://doi.org/10.1063/1.4942518>
- Getreuer P (2013) Wavelet CDF 9/7 implementation. Retrieved from <http://www.getreuer.info/home/waveletcdf97/>, November 11, 2021
- Hatfield S, Subramanian A, Palmer T, Düben P (2018) Improving weather forecast skill through reduced-precision data assimilation. *Mon Weather Rev* 146(1):49–62
- Hilbert M (2016) Big Data for development: a review of promises and challenges. *Dev Pol Rev* 34:135–174
- Kageyama A, Sato T (2004) Yin-Yang grid: an overset grid in spherical geometry. *Geochem Geophys.* <https://doi.org/10.1029/2004GC000734>

- Kageyama A, Kameyama M, Fujihara S, Yoshida M, Hyodo M, Tsuda Y (2004) A 15.2 TFlops simulation of geodynamo on the Earth Simulator. In: SC '04: Proceedings of the 2004 ACM/IEEE conference on supercomputing, pp 35, <https://doi.org/10.1109/SC.2004.1>
- Kang H, Lee D, Lee D (2003) A study on CFD data compression using hybrid supercompact wavelets. *KSME Int J* 17(11):1784–1792. <https://doi.org/10.1007/BF02983609>
- Komatitsch D, Tsuboi S, Tromp J (2013) The spectral-element method in seismology. American Geophysical Union (AGU), pp 205–227. <https://doi.org/10.1029/157GM13>
- Korpilo T, Gurchenko AD, Gusakov EZ, Heikkinen JA, Janhunen SJ, Kiviniemi TP, Leerink S, Niskala P, Perevalov AA (2016) Gyrokinetic full-torus simulations of ohmic tokamak plasmas in circular limiter configuration. *Comput Phys Commun* 203:128–137. <https://doi.org/10.1016/j.cpc.2016.02.021>
- Lakshminarasimhan S, Shah N, Ethier S, Klasky S, Latham R, Ross R, Samatova NF (2011) Compressing the incompressible with ISABELA: In-situ reduction of spatio-temporal data. Euro-Par 2011 Parallel Processing. Lecture Notes in Computer Science, Springer, Berlin Heidelberg, pp 366–379
- Laney D, Langer S, Weber C, Lindstrom P, Wegener A (2013) Assessing the effects of data compression in simulations using physically motivated metrics. In: SC '13: Proceedings of the international conference on high performance computing, networking, storage and analysis, pp 1–12, <https://doi.org/10.1145/2503210.2503283>
- Li S, Gruchalla K, Potter K, Clyne J, Childs H (2015) Evaluating the efficacy of wavelet configurations on turbulent-flow data. In: 2015 IEEE 5th Symposium on large data analysis and visualization (LDAV), pp 81–89, <https://doi.org/10.1109/LDAV.2015.7348075>
- Li S, Marsaglia N, Garth C, Woodring J, Clyne J, Childs H (2018) Data reduction techniques for simulation, visualization and data analysis. *Comput Graph Forum* 37(6):422–447. <https://doi.org/10.1111/cgf.13336>
- Liang X, Di S, Tao D, Li S, Li S, Guo H, Chen Z, Cappello F (2018) Error-controlled lossy compression optimized for high compression ratios of scientific datasets. 2018 IEEE international conference on big data (Big Data) pp 438–447
- Lorente LS, Vega JM, Velazquez A (2010) Compression of aerodynamic databases using high-order singular value decomposition. *Aerosp Sci Technol* 14(3):168–177
- Lundquist KA, Chow FK, Lundquist JK (2010) An immersed boundary method for the weather research and forecasting model. *Mon Weather Rev* 138(3):796–817
- Martin GNN (1979) Range encoding: an algorithm for removing redundancy from a digitized message. In: Video & data recording conference, Southampton, UK
- Matsuda K, Onishi R, Takahashi K (2018) Tree-crown-resolving large-eddy simulation coupled with three-dimensional radiative transfer model. *J Wind Eng Ind Aerod* 173:53–66
- Muraki S (1993) Volume data and wavelet transforms. *IEEE Comput Graph Appl* 13(4):50–56. <https://doi.org/10.1109/38.219451>
- Najmabadi SM, Offenhäuser P, Hamann M, Jainbalkya G, Hempert F, Glass CW, Simon S (2017) Analyzing the effect and performance of lossy compression on aeroacoustic simulation of gas injector. *Computation*. <https://doi.org/10.3390/computation5020024>
- Nakano M, Wada A, Sawada M, Yoshimura H, Onishi R, Kawahara S, Sasaki W, Nasuno T, Yamaguchi M, Iriguchi T, Sugi M, Takeuchi Y (2017) Global 7 km mesh nonhydrostatic model intercomparison project for improving typhoon forecast (TYMIP-G7): experimental design and preliminary results. *Geosci Model Dev* 10(3):1363–1381. <https://doi.org/10.5194/gmd-10-1363-2017>
- Onishi R, Takahashi K (2012) A warm-bin-cold-bulk hybrid cloud microphysical model. *J Atmos Sci* 69(5):1474–1497. <https://doi.org/10.1175/JAS-D-11-0166.1>
- Onishi R, Baba Y, Takahashi K (2011) Large-scale forcing with less communication in finite-difference simulations of stationary isotropic turbulence. *J Comput Phys* 230(10):4088–4099
- Onishi R, Takahashi K, Vassilicos JC (2013) An efficient parallel simulation of interacting inertial particles in homogeneous isotropic turbulence. *J Comput Phys* 242:809–827
- Pavlov I (2019) 7z format. Retrieved from <http://www.7-zip.org/7z.html>, November 11, 2021
- Peyrot JL, Duval L, Payan F, Bouard L, Chizat L, Schneider S, Antonini M (2019) HexaShrink, an exact scalable framework for hexahedral meshes with attributes and discontinuities: multiresolution rendering and storage of geoscience models. *Comput Geosci* 23(4):723–743. <https://doi.org/10.1007/s10596-019-9816-2>
- Ravi S, Kolomenskiy D, Engels T, Schneider K, Wang C, Sesterhenn J, Liu H (2016) Bumblebees minimize control challenges by combining active and passive modes in unsteady winds. *Sci Rep* 6:35043
- Ronchi C, Iacono R, Paolucci PS (1996) The Cubed Sphere: a new method for the solution of partial differential equations in spherical geometry. *J Comput Phys* 124(1):93–114. <https://doi.org/10.1006/jcph.1996.0047>
- Røsten T, Ramstad TA, Amundsen L (2004) Optimization of sub-band coding method for seismic data compression. *Geophys Prospect* 52(5):359–378. <https://doi.org/10.1111/j.1365-2478.2004.00422.x>
- Sakai R, Sasaki D, Nakahashi K (2011) Large-scale CFD data compression for building-cube method using wavelet transform. In: Kuzmin A (ed) *Computational fluid dynamics 2010*. Springer, Berlin Heidelberg, Berlin, Heidelberg, pp 465–470
- Sakai R, Sasaki D, Nakahashi K (2013a) Parallel implementation of large-scale CFD data compression toward aeroacoustic analysis. *Comput Fluids* 80:116–127
- Sakai R, Sasaki D, Obayashi S, Nakahashi K (2013b) Wavelet-based data compression for flow simulation on block-structured cartesian mesh. *Int J Numer Meth Fl* 73(5):462–476. <https://doi.org/10.1002/fld.3808>
- Schindler M (1999) Range encoder homepage. Retrieved from <http://www.compressconsult.com/rangecoder/>, November 11, 2021
- Schlegel M, Noack BR, Comte P, Kolomenskiy D, Schneider K, M F, Luchtenburg DM, Scouten JE, Tadmor G, (2009) Reduced-order modelling of turbulent jets for noise control. In: Juvé D, Manhart M, Munz CD (eds) *Numerical Simulation of Turbulent Flows and Noise Generation, Notes on Numerical Fluid Mechanics and Multidisciplinary Design*, vol 104. Springer, Berlin Heidelberg, Berlin, Heidelberg, pp 3–27

- Schmalzl J (2003) Using standard image compression algorithms to store data from computational fluid dynamics. *Comput Geosci* 29(8):1021–1031. [https://doi.org/10.1016/S0098-3004\(03\)00098-0](https://doi.org/10.1016/S0098-3004(03)00098-0)
- Schmid PJ (2010) Dynamic mode decomposition of numerical and experimental data. *J Fluid Mech* 656:5–28
- Schneider K, Vasilyev OV (2010) Wavelet methods in computational fluid dynamics. *Ann Rev Fluid Mech* 42(1):473–503
- Sekiguchi M, Nakajima T (2008) A k-distribution-based radiation code and its computational optimization for an atmospheric general circulation model. *J Quant Spectrosc Ra* 109(17):2779–2793. <https://doi.org/10.1016/j.jqsrt.2008.07.013>
- Shannon CE (1948) A mathematical theory of communication. *Bell Syst Tech J* 27(379–423):623–656
- Takahashi K, Onishi R, Baba Y, Kida S, Matsuda K, Goto K, Fuchigami H (2013) Challenge toward the prediction of typhoon behaviour and down pour. *J Phys Conf Ser* 454(1):012072
- Tao D, Di S, Chen Z, Cappello F (2017) Significantly improving lossy compression for scientific data sets based on multidimensional prediction and error-controlled quantization. In: *Proceedings of the 31st IEEE International parallel and distributed processing symposium (ipdps)*, Orlando, Florida, pp 1129–1139
- Taubman D, Marcellin M (2002) *JPEG2000 image compression fundamentals, standards and practice*, The Springer International Series in Engineering and Computer Science, vol 642, 1st edn. Springer, US
- Wilson JP (2002) Wavelet-based lossy compression of barotropic turbulence simulation data. In: *Proceedings DCC 2002. Data compression conference*, pp 479. <https://doi.org/10.1109/DCC.2002.1000022>
- Woodring J, Mniszewski S, Brislawn C, DeMarle D, Ahrens J (2011) Revisiting wavelet compression for large-scale climate data using JPEG 2000 and ensuring data precision. In: *2011 IEEE symposium on large data analysis and visualization*, pp 31–38, 10.1109/LDAV.2011.6092314
- Yeh PS, Xia-Serafino W, Miles L, Kobler B, Menasce D (2002) Implementation of CCSDS lossless data compression in HDF. In: *Earth Science Technology Conference—2002*, Pasadena, California, p A3P2

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.