

Computer Programs in Physics



Unlocking massively parallel spectral proper orthogonal decompositions in the PySPOD package ☆,☆☆

Marcin Rogowski^{a,*}, Brandon C.Y. Yeung^b, Oliver T. Schmidt^b, Romit Maulik^f,
Lisandro Dalcin^{a,**}, Matteo Parsani^{a,c}, Gianmarco Mengaldo^{d,e,***}

^a King Abdullah University of Science and Technology (KAUST), Computer Electrical and Mathematical Science and Engineering Division (CEMSE), Extreme Computing Research Center (ECRC), Thuwal, Saudi Arabia

^b Department of Mechanical and Aerospace Engineering, University of California San Diego, La Jolla, CA, USA

^c King Abdullah University of Science and Technology (KAUST), Physical Science and Engineering Division (PSE), Thuwal, Saudi Arabia

^d Department of Mechanical Engineering, National University of Singapore, Singapore, SG

^e Honorary Research Fellow, Department of Aeronautics, Imperial College London, London, UK

^f Information Science and Technology Department, Pennsylvania State University, State College, PA, USA

ARTICLE INFO

Keywords:

Spectral proper orthogonal decomposition
SPOD
Parallel
Distributed
MPI
Modal decomposition
Dynamical systems

ABSTRACT

We propose a parallel (distributed) version of the spectral proper orthogonal decomposition (SPOD) technique. The parallel SPOD algorithm distributes the spatial dimension of the dataset preserving time. This approach is adopted to preserve the non-distributed fast Fourier transform of the data in time, thereby avoiding the associated bottlenecks. The parallel SPOD algorithm is implemented in the PySPOD library and makes use of the standard message passing interface (MPI) library, implemented in Python via `mpi4py`. An extensive performance evaluation of the parallel package is provided, including strong and weak scalability analyses. The open-source library allows the analysis of large datasets of interest across the scientific community. Here, we present applications in fluid dynamics and geophysics, that are extremely difficult (if not impossible) to achieve without a parallel algorithm. This work opens the path toward modal analyses of big quasi-stationary data, helping to uncover new unexplored spatio-temporal patterns.

Program summary

Program Title: PySPOD

CPC Library link to program files: <https://doi.org/10.17632/jf5bf26jci.1>

Developer's repository link: <https://github.com/MathEXLab/PySPOD>

Licensing provisions: MIT License

Programming language: Python

Nature of problem: Large spatio-temporal datasets may contain coherent patterns that can be leveraged to better understand, model, and possibly predict the behavior of complex dynamical systems. To this end, modal decomposition methods, such as the proper orthogonal decomposition (POD) and its spectral counterpart (SPOD), constitute powerful tools. The SPOD algorithm allows the systematic identification of space-time coherent patterns. This can be used to understand better the physics of the process of interest, and provide a path for mathematical modeling, including reduced order modeling. The SPOD algorithm has been successfully applied to fluid dynamics, geophysics and other domains. However, the existing open-source implementations are serial, and they prevent running on the increasingly large datasets that are becoming available, especially in computational physics. The inability to analyze via SPOD large dataset in turn prevents unlocking novel mechanisms and dynamical behaviors in complex systems.

☆ The review of this paper was arranged by Prof. W. Jong.

☆☆ This paper and its associated computer program are available via the Computer Physics Communications homepage on ScienceDirect (<http://www.sciencedirect.com/science/journal/00104655>).

* Corresponding author.

** Corresponding author.

*** Corresponding author at: Department of Mechanical Engineering, National University of Singapore, Singapore, SG.

E-mail addresses: marcin.rogowski@kaust.edu.sa (M. Rogowski), dalcin@gmail.com (L. Dalcin), mpegim@nus.edu.sg (G. Mengaldo).

<https://doi.org/10.1016/j.cpc.2024.109246>

Received 8 August 2023; Received in revised form 7 May 2024; Accepted 10 May 2024

Available online 16 May 2024

0010-4655/© 2024 Elsevier B.V. All rights reserved, including those for text and data mining, AI training, and similar technologies.

Solution method: We provide an open-source parallel (MPI distributed) code, namely PySPOD, that is able to run on large datasets (the ones considered in the present paper reach about 200 Terabytes). The code is built on the previous serial open-source code PySPOD that was published in <https://joss.theoj.org/papers/10.21105/joss.02862.pdf>. The new parallel implementation is able to scale on several nodes (we show both weak and strong scalability) and solve some of the bottlenecks that are commonly found at the I/O stage. The current parallel code allows running on datasets that was not easy or possible to analyze with serial SPOD algorithms, hence providing a path towards unlocking novel findings in computational physics.

Additional comments including restrictions and unusual features: The code comes with a set of built-in postprocessing tools, for visualizing the results. It also comes with extensive continuous integration, documentation, and tutorials, as well as a dedicated website in addition to the associated GitHub repository. Within the package we also provide a parallel implementation of the proper orthogonal decomposition (POD), that leverages the I/O parallel capabilities of the SPOD algorithm.

1. Introduction

Data that depends on both space and time, also referred to as spatio-temporal data, is ubiquitous. It can represent the Earth's atmosphere, the flow past an aircraft, and the ocean's dynamics, among many other phenomena and processes. Usually, spatio-temporal data is high-dimensional and its interpretation non obvious. Yet, spatio-temporal data, especially the one related to physical processes, contain coherent patterns. These, if uncovered, may provide a better understanding of critical aspects of the underlying physical processes of interest. Hence, tools to analyze and make sense of this type of data are of paramount importance.

Over the last several years, many tools have been proposed to mining information from spatio-temporal data, such as the proper orthogonal decomposition (POD) [1,2], the dynamic mode decomposition (DMD) [3], and the spectral proper orthogonal decomposition (SPOD) [4–8]. Yet, the majority of these tools comes with limited parallel capabilities for handling large datasets.

In this paper, we propose a first parallel (distributed) SPOD algorithm, that allows large-scale modal decomposition analyses. The algorithm has been tested on both geophysical and fluid mechanics data, reaching 199 terabytes (TB) of size. Weak and strong scalability have been thoroughly assessed for the entire algorithm as well as for its main constituting steps, including input/output handling, discrete Fourier transform, and eigenvalue computations. The efficiency of the novel parallel algorithm proposed in this paper allows analysis of data at unprecedented scales, opening the path towards uncovering new physics in large datasets, that existing SPOD packages were unable to tackle [9,10].

This paper is organized as follows. In section 2, we introduce the SPOD method. In section 3, we detail our parallelization strategy. In section 4, we present the results obtained on two large datasets, one related to fluid mechanics and the other to geophysics (in particular atmospheric physics). In section 5, we present the strong and weak scalability analysis of the parallel SPOD algorithm. In section 6, we provide some concluding remarks.

2. The spectral proper orthogonal decomposition

2.1. A note on suitable data

The spectral proper orthogonal decomposition, also referred to as spectral POD or simply SPOD, extracts coherent structures from statistically stationary data. The data considered have spatial and temporal dependence, and describe a stochastic (also referred to as random) process denoted by $q(\mathbf{x}, t)$, where \mathbf{x} represents the spatial coordinate, and t the time. Usually, $\mathbf{x} = (x, y, z) \in \mathbb{R}^3$ for the processes of interest, where x, y , and z are Cartesian coordinates. In practice, both two- and three-dimensional data may be used. We now define in more detail the concept of stationarity and the assumptions that allow us to use ensemble and time averages interchangeably.

The *stationarity* assumption under which the SPOD operates is typically intended in the *weak sense*, also referred to as *wide-sense* or *covariance stationarity*. This implies that $q(\mathbf{x}, t)$, where we dropped the probability parametrization ξ , has first- and second-order moments (i.e., average and autocovariance) that do not vary with time. The mathematical formalization of wide-sense stationarity is as follows. Let $q(\mathbf{x}, t)$ be a continuous-time stochastic process, $E[\cdot]$ be the expectation operator, and C be the covariance. If

- the expectation operator is independent of t :

$$E[q(\mathbf{x}, t)] = \mu(\mathbf{x}), \quad (1)$$

- the covariance depends only on the difference between two times, $t - t'$:

$$C(\mathbf{x}, \mathbf{x}', t, t') = C(\mathbf{x}, \mathbf{x}', \tau), \quad \text{where } \tau = t - t', \quad (2)$$

- the average ‘power’ is bounded, and does not go to infinity:

$$E[|q(\mathbf{x}, t)|^2] < \infty, \quad (3)$$

then $q(\mathbf{x}, t)$ is said to be *wide-sense stationary*.

Many fields in computational physics, including fluid dynamics and geophysics, give rise to wide-sense stationary problems. In section 4 we will apply SPOD to examples of wide-sense stationary problems in these two disciplines.

As a final yet important note in this section, we add that the stochastic processes considered are also ergodic in addition to being wide-sense stationary. This means that the expectation, $E[q]$, coincides with the ensemble average of different realizations of $q(\mathbf{x}, t)$, that in turn is equal to the long-time average, $q(\mathbf{x}, t)$. Therefore, when we say zero-average stochastic process, we refer to a process where $\bar{q} = E[q]$ has been removed. Removal of the mean facilitates the interpretation of the SPOD eigenvalues as perturbation energy or variance. In the following, without loss of generality we will always assume that $\bar{q} = 0$. In sections 2.2 and 2.3, we will use some of the notions introduced here to derive the continuous and discrete SPOD approaches.

2.2. Theory

For the sake of reader's convenience, we report here the theory behind SPOD, closely following [5].

The task of the SPOD is to identify a deterministic function $\phi(\mathbf{x}, t)$ (or a set of functions) that best approximates the weak-sense stationary and zero-average process $q(\mathbf{x}, t)$. In mathematical terms, this translates into finding the function $\phi(\mathbf{x}, t)$ that maximizes the expected value of the normalized projection of the stochastic function $q(\mathbf{x}, t)$, that is,

$$\lambda = \frac{E[|\langle q(\mathbf{x}, t), \phi(\mathbf{x}, t) \rangle_{\mathbf{x}, t}|^2]}{\langle \phi(\mathbf{x}, t), \phi(\mathbf{x}, t) \rangle_{\mathbf{x}, t}}. \quad (4)$$

In equation (4), we assume that any realization of $q(\mathbf{x}, t)$ belongs to a Hilbert space, H , with a space-time inner product, $\langle \cdot, \cdot \rangle_{\mathbf{x}, t}$, and expect-

tation operator, $E[\cdot]$, here taken to be the ensemble average. The inner product in equation (4), $\langle \cdot, \cdot \rangle_{\mathbf{x},t}$, between two generic variables, \mathbf{u} and \mathbf{v} , is defined as

$$\langle \mathbf{u}, \mathbf{v} \rangle_{\mathbf{x},t} = \int_{-\infty}^{\infty} \int_{\Omega} \mathbf{u}^*(\mathbf{x}, t) \mathbf{W}(\mathbf{x}) \mathbf{v}(\mathbf{x}, t) d\mathbf{x} dt, \quad (5)$$

where Ω denotes the spatial domain, $\mathbf{W}(\mathbf{x})$ the spatial weighting, and the asterisk superscript represents the conjugate transpose. By invoking the Karhunen-Loève (KL) theorem [11,12], we know that there exists a set of mutually orthogonal deterministic functions that form a complete basis in H . This can be defined as $\hat{q}(\mathbf{x}, f) = \sum_{k=1}^{\infty} a_k(f) \phi_k(\mathbf{x}, f)$, where $\hat{(\cdot)}$ denotes the Fourier transform in time. The eigenfunctions, ϕ_k , and their associated eigenvalues, λ_k , are solutions to the eigenvalue problem in the frequency domain:

$$\int_{\Omega} \hat{C}(\mathbf{x}, \mathbf{x}', f) \mathbf{W}(\mathbf{x}') \phi(\mathbf{x}', f) d\mathbf{x}' = \lambda(f) \phi(\mathbf{x}, f), \quad (6)$$

where $\hat{C}(\mathbf{x}, \mathbf{x}', f)$ is the cross-spectral density tensor, i.e., the Fourier transform of C :

$$\hat{C}(\mathbf{x}, \mathbf{x}', f) = \int_{-\infty}^{\infty} C(\mathbf{x}, \mathbf{x}', \tau) e^{-i2\pi f \tau} d\tau. \quad (7)$$

In equation (7), the two-point space-time correlation tensor $C(\mathbf{x}, \mathbf{x}', \tau)$ is defined as

$$C(\mathbf{x}, \mathbf{x}', \tau) = E[q(\mathbf{x}, t) q^*(\mathbf{x}', t')], \quad (8)$$

where we used equation (2), under the assumption of wide-sense stationarity of the stochastic process $q(\mathbf{x}, t)$ and $\tau = t - t'$ is the difference between the two times t and t' . This implies that the covariance depends only on the difference between two times, t and t' and, therefore, we can write, $C(\mathbf{x}, \mathbf{x}', t, t') \rightarrow C(\mathbf{x}, \mathbf{x}', \tau)$. The last step allows reformulating the problem in the spectral (also referred to as frequency) domain, as per equation (6). This significant result was first presented in [4,13] and revisited in [5]. It provides eigenmodes at each frequency that inherit the same properties of the more traditional (non-spectral) POD.

The SPOD formulation just summarized leads to monochromatic SPOD modes that optimally characterize the second-order space-time moments of the continuous-time stochastic process considered [5].

2.3. Practical implementation

In practice, the continuous-time stochastic process $q(\mathbf{x}, t)$ introduced in section 2.1 and used in section 2.2 is provided as discrete data. The discrete data consist of snapshots of the wide-sense stationary time series, $\mathbf{q}(\mathbf{x}, t_i)$, $t_i = 1, \dots, N_f$, from which we subtract the temporal mean, $\bar{\mathbf{q}}$. Each snapshot $\mathbf{q}(\mathbf{x}, t_i)$ is sampled at M_{space} spatial points with coordinates $\mathbf{x} \in \mathbb{R}^{M_{\text{space}} \times d}$ (usually two- or three-dimensional, that is, $d = 2$ or 3, respectively), and records M_{vars} variables.

To derive the SPOD algorithm, we recast each snapshot of the discrete multidimensional data into a vector of dimension $\mathbf{q}(\mathbf{x}, t_i) = \mathbf{q}_i \in \mathbb{R}^M$, where $M = M_{\text{space}} M_{\text{vars}}$. We can then assemble the data matrix (also referred to as snapshot matrix),

$$\mathbf{Q} = [\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_{N_f}] \in \mathbb{R}^{M \times N_f}. \quad (9)$$

The data described by equation (9) can arise from simulations and observations of a wide-sense stationary stochastic system. We assumed the data to be composed exclusively of real numbers. However, this assumption is not strictly necessary, as the data can, in principle, be complex.

The first step to obtaining the discrete analog of the frequency-domain eigenvalue problem in equation (6) consists of segmenting the data along the time direction into L (possibly overlapping) blocks

$$\mathbf{Q}^{(\ell)} = [\mathbf{q}_1^{(\ell)}, \dots, \mathbf{q}_{N_f}^{(\ell)}] \in \mathbb{R}^{M \times N_f},$$

$$\ell = 1, \dots, L \implies \begin{cases} \mathbf{Q}^{(1)} = [\mathbf{q}_1^{(1)}, \dots, \mathbf{q}_{N_f}^{(1)}] \in \mathbb{R}^{M \times N_f}, \\ \mathbf{Q}^{(2)} = [\mathbf{q}_1^{(2)}, \dots, \mathbf{q}_{N_f}^{(2)}] \in \mathbb{R}^{M \times N_f}, \\ \dots \\ \mathbf{Q}^{(L)} = [\mathbf{q}_1^{(L)}, \dots, \mathbf{q}_{N_f}^{(L)}] \in \mathbb{R}^{M \times N_f}. \end{cases} \quad (10)$$

Each data block ℓ in equation (10) contains N_f time snapshots, overlaps by N_{overlap} time snapshots with the adjacent block, and is regarded as equally representative of the whole data by the ergodic assumption. Indeed, it is a realization of the stochastic process described by the discrete data \mathbf{Q} in equation (9). This approach of partitioning the time series into overlapping data blocks is the well-known Welch periodogram method [14,7]. A 50% overlap, i.e., $N_{\text{overlap}} = N_f/2$, is accepted best practice. In the following, we will use the terms *realization* and *data block* interchangeably.

The second step consists of applying the discrete Fourier transform (DFT) in time to each data block or realization in equation (10):

$$\mathbf{Q}^{(\ell)} \xrightarrow{\text{DFT}} \hat{\mathbf{Q}}^{(\ell)} = [\hat{\mathbf{q}}_1^{(\ell)}, \hat{\mathbf{q}}_2^{(\ell)}, \dots, \hat{\mathbf{q}}_{N_f}^{(\ell)}] \in \mathbb{C}^{M \times N_f}, \quad \ell = 1, \dots, L. \quad (11)$$

We note that each Fourier-transformed data block $\hat{\mathbf{Q}}^{(\ell)}$ contains N_f frequencies. Wide-sense stationarity and ergodicity allow us to reorganize the Fourier-transformed data into N_f data matrices, one per frequency, f_k . In particular, we collect all realizations of the DFT at the k -th frequency into

$$\hat{\mathbf{Q}}_k = [\hat{\mathbf{q}}_k^{(1)}, \hat{\mathbf{q}}_k^{(2)}, \dots, \hat{\mathbf{q}}_k^{(L)}] \in \mathbb{C}^{M \times L}, \quad \text{for all frequencies } f_k, k = 1, \dots, N_f. \quad (12)$$

For fluid mechanical and geophysical applications, we usually have $L \ll M$. For the parallelization of the SPOD algorithm we will make use of this notion.

The third step is to construct the cross-spectral density matrix for each frequency, f_k . This step corresponds to the discrete counterpart of equation (7), and can be readily achieved by calculating

$$\hat{C}_k = \frac{1}{L-1} \hat{\mathbf{Q}}_k \hat{\mathbf{Q}}_k^* \in \mathbb{C}^{M \times M}, \quad \text{for all frequencies } f_k, k = 1, \dots, N_f, \quad (13)$$

where $L-1$ is a normalization factor known as Bessel's correction, that is only appropriate if the data is centered about the sample mean rather than the long-time (i.e., true) mean. The construction of the cross-spectral density matrix in equation (13) finally allows us to write the discrete analog of the frequency-domain eigenvalue problem defined in equation (6):

$$\begin{aligned} \hat{C}_k \mathbf{W} \Phi_k &= \Phi_k \Lambda_k, \quad \text{with} \\ \Phi_k &= [\phi_k^{(1)}, \phi_k^{(2)}, \dots, \phi_k^{(L)}] \in \mathbb{C}^{M \times L} \quad (\text{SPOD modes}), \\ \Lambda_k &= \text{diag}(\lambda_k^{(1)}, \lambda_k^{(2)}, \dots, \lambda_k^{(L)}) \in \mathbb{R}^{L \times L} \quad (\text{modal energies}). \end{aligned} \quad (14)$$

The SPOD modes, Φ_k , and associated modal energies (or eigenvalues), Λ_k , can be computed by solving equation (14) for each frequency, f_k . In practice, to alleviate the computational burden of this step, one usually turns to the method of snapshots [15]:

$$\hat{\mathbf{Q}}_k^* \mathbf{W} \hat{\mathbf{Q}}_k \Psi_k = \Psi_k \Lambda_k, \quad \Phi_k = \hat{\mathbf{Q}}_k \Psi_k \Lambda_k^{-1/2}. \quad (15)$$

By construction, for a given frequency f_k , the modes Φ_k are orthonormal, i.e., $\Phi_k^* \mathbf{W} \Phi_k = \mathbf{I}$, where \mathbf{I} is the identity matrix. Modes at different frequencies are instead not orthonormal, that is, $\Phi_{k_1}^* \mathbf{W} \Phi_{k_2} \neq \mathbf{I}$, where $k_1 \neq k_2$, but orthonormal under the full space-time inner product.

Finally, the SPOD modes can be grouped per frequency, as follows:

$$\Phi = [\Phi_1, \dots, \Phi_{N_{f/2}}] = [\underbrace{\phi_1^{(1)}, \dots, \phi_1^{(L)}}_{\Phi_1}, \underbrace{\phi_2^{(1)}, \dots, \phi_2^{(L)}}_{\Phi_2}, \dots, \underbrace{\phi_{N_{f/2}}^{(1)}, \dots, \phi_{N_{f/2}}^{(L)}}_{\Phi_{N_{f/2}}}], \quad (16)$$

where, for real data, $\mathbf{Q} \in \mathbb{R}^{M \times N_t}$, the total number of frequencies is $N_{f/2} = \lceil \frac{N_f}{2} \rceil + 1$. This is because the transformed data at negative frequencies correspond to the conjugates of the positive frequencies, and it is therefore redundant. For additional details on the SPOD method, the interested reader can refer to [5,7].

2.4. SPOD for data compression

As shown in [16], and further explored in [8], it is possible to compute a matrix of expansion coefficients \mathbf{A} . This can be constructed using a weighted oblique projection of the data onto the modal basis

$$\mathbf{A} = (\Phi^* \mathbf{W} \Phi)^{-1} \Phi^* \mathbf{W} \mathbf{Q} \\ = [\underbrace{a_1^{(1)}, a_1^{(2)}, \dots, a_1^{(L)}}_{\mathbf{a}_1}, \underbrace{a_2^{(1)}, a_2^{(2)}, \dots, a_2^{(L)}}_{\mathbf{a}_2}, \dots, \underbrace{a_{N_f}^{(1)}, a_{N_f}^{(2)}, \dots, a_{N_f}^{(L)}}_{\mathbf{a}_{N_f}}]. \quad (17)$$

where $\mathbf{A} \in \mathbb{C}^{(L \times N_f) \times N_t}$ is the matrix containing the expansion coefficients and $\Phi \in \mathbb{C}^{M \times (L \times N_f)}$ is a matrix which gathers all the SPOD modes arranged by frequency as in equation (16). The full matrix of expansion coefficients, constructed using all modes and frequencies, has dimensions $L \times N_f \times N_t$. In practice, it is common to use only a portion L_r of the total number of modes, and (eventually) a portion N_{f_r} of the total number of frequencies, where we denote with Φ_r and \mathbf{A}_r , the reduced number of SPOD modes and expansion coefficients, respectively. This reduction recasts the original high-dimensional data into a smaller SPOD latent space of dimension $L_r \times N_{f_r}$.

Once both SPOD modes and expansion coefficients are available, it is possible to reconstruct the original high-dimensional data as follows

$$\tilde{\mathbf{Q}} = \Phi_r \mathbf{A}_r, \quad (18)$$

where $\tilde{\mathbf{Q}}$ is an approximation of the original data \mathbf{Q} , given the truncation imposed on the number of SPOD modes and frequencies. The storage of the expansion coefficients and SPOD modes required to reconstruct the high-dimensional data in equation (18), can potentially lead to significant savings in terms of memory storage. In addition, the ability to truncate number of SPOD modes and number of frequencies can be beneficial to just store the frequencies of interest (e.g., removing high-frequency noise from the data), and capture low-rank behavior of the process under study [5,6].

For instance, if we consider a problem constituted by 30,000 time snapshots, 37,000,000 spatial points, and 1 variable, we have: $N_t \times M_{\text{space}} \times M_{\text{vars}} = 30,000 \times 37,000,000 \times 1$. To store this dataset in memory, we require approximately 8.88 TB in double floating-point precision, and 4.44 TB in single floating-point precision.

However, if we, for instance, store the first 3 SPOD modes and 100 frequencies of interest, the amount of storage memory required is significantly lowered. To store the SPOD modes, we would need $3 \times 100 \times 37,000,000 \times 2$, that leads to 0.18 TB of memory in double floating-point precision, and 0.089 TB in single floating-point precision. To store the time coefficients, we would need $3 \times 100 \times 30,000 \times 2$, that leads to 0.00014 TB in double floating-point precision and 0.000072 TB in single floating-point precision. Hence, if we store 3 SPOD modes and 100 frequencies, the storage memory required is only 2% of the memory required to store the original dataset, and the data compression is achieved by having full control on what modes and frequencies are left out (if any).

3. Parallelization strategy

The parallelization efforts have been mainly focused on the matrix operations required by the SPOD algorithm, and on the data Input/Output (I/O). We outline the parallelization of the algorithm first (section 3.1), and reserve section 3.2 for I/O, as it is a crucial aspect to achieve competitive scalability results.

3.1. SPOD algorithm

The parallelization strategy for the SPOD algorithm uses a single-program multiple-data (SPMD) approach (the interested reader can also refer to other parallelized modal decompositions, such as [17], that is specialized to boundary-layer flows). It allows maintaining the structure of the code similar to that of the serial one, only introducing parallel communication and synchronization in a limited number of places. For the SPOD algorithm, this consists of decomposing the spatial dimensions M_{space} (conveniently flattened) of the data matrix in equation (9). This is a practical and advantageous choice as it allows preserving all operations in time – in particular, the DFT – without needing expensive all-to-all communication patterns. Additionally, the number of spatial points M_{space} is usually much larger than the number of variables M_{vars} , thus allowing for considerably higher parallelism. The decomposition of the space dimension allows a straightforward parallel implementation that only requires one single MPI collective reduction operation. Most of the MPI operations are implemented in an auxiliary utility module. This allows using MPI routines off-the-shelf, requiring minor modifications to the original (serial) code.

Entering into more detail, and assuming that the data distribution is achieved right after I/O, the parallelization of the algorithm becomes trivial. Indeed, it consists of a simple parallel reduction operation for the inner product derived from the method of snapshots in equation (15). The MPI-based implementation uses the mpi4py package [18,19]. In Fig. 1, we depict the parallel SPOD algorithm. In red, we represent the parallel operations required, while in blue are operations that remain unchanged compared to the serial code. The six steps reported in Fig. 1 are described in the following.

- Step 1) Distribute data: it consists of distributing the spatial dimension of each of the data blocks \mathbf{Q}^ℓ , $\ell = 1, \dots, L$ across the MPI ranks available.
- Step 2) DFT (see equation (11)): it consists of performing the DFT along the time dimension. This operation remains unchanged as time has not been distributed.
- Steps 3, 4) Inner product and reduction (see equation (15)): the inner product operation involves contracting the spatial dimension, hence it requires a parallel reduction operation.
- Step 5) Eigen-decomposition (see equation (15)): this operation does not require any parallel handling as there is no manipulation of the distributed spatial dimension. Since the size of the blocks is typically small, eigen-decompositions are cheap to compute. Therefore, we perform this step redundantly in all MPI ranks. If this task ever becomes a computational bottleneck, it would be straightforward to replace with a distributed-memory implementation exploiting the already available CPU resources, or alternatively, GPU off-loading for acceleration.
- Step 6) Parallel I/O: This operation involves writing the SPOD modes to disk. The I/O handling is described in the next section 3.2, and is key to the scalability and overall performance of our implementation.

3.2. I/O handling

We dedicate a separate section to the I/O aspect of PySPOD as, in most practical scenarios, it plays a critical role in the performance of




Fig. 1. Schematic of the parallel SPOD algorithm. The key aspect is to obtain an appropriate data decomposition layout that allows preserving all time operations as done in serial (i.e., the DFT), and decompose only the spatial dimensions of the data. Once the data is in the required parallel layout, the parallelization of the SPOD algorithm becomes trivial and consists only of a parallel reduction (step 4) of the inner product (step 3). (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

PySPOD overall. This is especially prominent in a strong scaling scenario such as that presented in section 5.2. We divide this section into background and related work (section 3.2.1), and the PySPOD I/O implementation (section 3.2.2).

3.2.1. Background and related work

Since the early 1990s, I/O has been acknowledged as a significant contributor to parallel application performance [20]. However, relatively little attention was given to I/O compared to the compute and communication subsystems during that time [21]. Since then, the performance gap between the I/O throughput and the compute capability of supercomputers has only increased, hence exacerbating data movement challenges [22]. A 2012 report titled “Storage challenges at Los Alamos National Lab” noted that parallel file systems lack large parallel I/O, which would be both high bandwidth and resistant to I/O patterns without additional tuning [23]. Consequently, I/O systems remain relatively slow and complex, making it difficult for non-expert users to utilize them fully.

Middleware libraries were introduced to improve usability and simplify parallel I/O. After all, as [22] noted, users need not be aware of their data’s low-level layout and organization, which opens possibilities for specialized I/O middleware. MPI-IO enables simple N-1 I/O (N processes, 1 file) [24], and it has been an ongoing effort since 1999. It contains optimizations reducing the number of distinct I/O requests using techniques such as data sieving and collective I/O for noncontiguous accesses [25]. MPI-IO, however, is still low-level, as the developer needs to calculate the offsets within the file, while MPI aims to access the data efficiently. Higher-level libraries such as HDF5 [26], Parallel netCDF [27], ADIOS [22], and ADIOS 2 [28] have also been introduced to isolate developers from low-level file system details and alleviate the burden of platform-specific performance optimization. These libraries often build on top of MPI-IO and introduce custom file formats.

Although robust solutions for parallel I/O are available, achieving good performance often requires application-specific and system-specific tuning. Researchers have dedicated significant effort to optimizing applications that utilize the popular Lustre file system. Tools such as Lustre IO Profiler [29] were introduced to monitor and understand the

I/O activities on the file system. The authors of this tool highlight the importance of incorporating system-specific details, such as the number of Object Storage Targets (OSTs), into the MPI distribution used in order to improve the achieved I/O bandwidth. The authors of Y-Lib [30] also point out the low performance of MPI-IO on Lustre file systems and propose a solution that minimizes contention for file system resources by controlling the number of OSTs with which each aggregator process communicates. This solution is shown to outperform MPI-IO in many cases. Similarly, the authors of [31] demonstrate that Parallel netCDF with MPI-IO does improve the performance over netCDF; however, even with supercomputer-specific optimizations, the performance still disappoints once a large number of MPI ranks is used. The authors utilize asynchronous I/O (quilt) servers alongside Parallel netCDF to further improve the effective bandwidth. Other solutions explore tuning the middleware or file system parameters, such as those presented in [32–35]. These approaches often use sophisticated techniques, such as genetic algorithms, to identify and tune parameters of the I/O stack (HDF5, MPI-IO, and Lustre/GPFS parameters) [34,35]. The wealth of work on the topic highlights the difficulty of achieving good performance on current parallel file systems.

Despite the availability of powerful solutions and tools to tune the performance, many developers still choose not to use them, even in leadership scale systems [29]. An analysis conducted in 2015 using Darshan [36] found that, depending on the supercomputing facility, between 50 to 95% of monitored jobs used POSIX I/O exclusively. However, as [36] notes, relying on POSIX I/O does have to result in poor application performance. In fact, in some cases, using a naive N-N (N processes, N files) approach, where each process reads from or writes to a unique file, can achieve better performance and scalability than N-1, as it does not incur overhead to maintain data consistency. This has been observed in PanFS, GPFS, and Lustre file systems. However, the N-N approach presents challenges related to usability, such as when restarting with a different number of processes, and high metadata costs at scale [23,37].

3.2.2. PySPOD I/O

In PySPOD, we take advantage of the N–N model’s good performance, but like authors of [37], slightly modify it to be N–M (generally $M < N$) mapping, where M is not to be confused with ‘italic’ M adopted in section 2.3. Such a setting, where N is the number of processes and M is the number of files, is better suited to our application characteristics. We use a simple two-phase I/O remnant of that proposed in [21] back in 1993.

In particular, we first read the data from disk in a contiguous manner; afterwards, we re-distribute the data according to the parallel decomposition our application needs, as explained in section 3.1. Parallel data re-distribution can be performed either with collective communications, or point-to-point communications. In our case, we use non-blocking point-to-point MPI communications, as it results in higher reliability in the HPC systems that have been used in this work.

In the following, we enumerate the reasons for our decision to adopt this two-phase I/O, and categorize these reasons into *application-specific* and *performance-specific*.

Application-specific

- PySPOD is designed to analyze large datasets, which are often split over multiple files. One such example is the Climate Data Store (CDS), which provides rich climate datasets [38], and that has been used in the scalability analyses presented in sections 5.2 and 5.3. CDS, however, limits the amount of data that can be downloaded in a single request. The resulting datasets’ size can be in the order of tens of terabytes split over several hundred files. Another example arises from computational fluid dynamics (CFD) simulators such as SSDC, Nektar++ and Charles [39–41]. Such software often writes output using one file per requested timestep. Considering the data size and PySPOD’s competitive scalability performance, we decided against converting the data to a common format (such as HDF5), as the conversion could be more expensive than the analysis, and we implemented several native readers for different data formats.
- Using two-phase I/O allows reusing sequential readers. The design of keeping the I/O (first phase) separate from the data distribution (second phase) effectively makes the I/O phase into many concurrent but sequential I/O streams. This design makes it easy for a user or a developer to implement support for additional file formats in PySPOD. The programming burden is greatly reduced, as there are plenty of sequential reader modules in the Python Package Index which can be used with minimal effort. This design choice removes the burden of “thinking parallel” from a user/developer perspective, as the data distribution is performed in the second phase, and it does not depend on the file format.

Performance-specific

- Using a two-phase I/O results in fewer and more contiguous requests to the storage system, which is preferred on a parallel file system. For illustration, most often, the data is split over multiple files, where each file represents a different range of timesteps, and the spatial coordinates follow the first dimension. In our design, each process reads contiguous data from at least one input file, i.e., all spatial dimensions for a subset of timesteps (or, more generally, the first dimension, i.e., time), which is good for performance. An alternative design, such as using MPI-IO to immediately read only a subset of spatial variables for all timesteps, may result in every process accessing every file, depending on the MPI implementation used. As we discussed in section 3.2.1, MPI-IO has been shown to perform poorly on Lustre file systems [30]. In fact, an MPI-IO reader is also implemented in PySPOD; however, we also obtained poor

performance results in early experimentation; hence, we decided to extend and optimize the two-phase reader.

- Separating I/O from MPI removes the risk of performance degradation due to an underoptimized MPI distribution. As discussed in section 3.2.1, MPI-IO’s performance may depend heavily on the MPI distribution used and its knowledge of the underlying I/O system. By separating I/O and data distribution, we read the data from the disk efficiently and then explicitly distribute it with MPI. Since we use point-to-point communication, which is at the very core of MPI, it is much less likely that those functions will perform poorly.

Application-specific and performance-specific

- SPOD output is expected in the form of one output file per frequency and mode. However, because the data for each frequency is spread across all processes, this would lead to N–1 storage access for each file. This approach would be highly inefficient due to relatively small spatial dimensions. To address this issue, we use a two-phase process similar to the one used for reading files, allowing each file to be written by only one process.

To address the memory capacity limitations of PySPOD, we implemented several precautions. One such measure was to read data in chunks, with each process reading approximately 256 MB of data (value determined empirically). This helps reducing memory overhead since processing each chunk requires twice the memory, as time-split data being sent can only be deallocated once it is received by the target process, and space-split data is being received simultaneously. The chunk size should be as small as possible to reduce the memory overhead of using auxiliary memory but large enough to hide the I/O latency. On our system, 256 MB chunks correspond to roughly 6% of total memory per CPU core.

Additionally, we used a dictionary of NumPy arrays instead of one large array. Such design allows us to deallocate processed data as its FFT transforms replace it in memory, therefore reducing peak memory usage. Specifically, we store the data for each chunk in three dictionary keys, with each NumPy array occupying approximately 85 MB of memory.

Moreover, we implemented reading in such a way that one process first seeks through the metadata to identify which range of timesteps is available in which file. This information, broadcast with MPI, significantly reduces the overhead associated with N processes opening all M files.

4. Datasets

The datasets adopted to test the parallel SPOD algorithm consist of fluid mechanics, and geophysical data. The former uses jet data produced by high-fidelity large-eddy simulation (LES), and is described in section 4.1. The latter uses fifth-generation reanalysis data (ERA5) produced by ECMWF [42], and is described in section 4.2, along with associated SPOD results.

4.1. SPOD analysis of fluid mechanics datasets

In our first example, we analyze the data generated by Yeung et al. [43] using the solver Charles [41], from LES of a supersonic twin-rectangular jet at a Reynolds number of $Re \approx 10^6$ based on the jet exit conditions and the equivalent nozzle diameter. The simulation was previously validated by Brès et al. [44] against the companion experiments of Samimy et al. [45]. The time-resolved data consist of 20,000 snapshots of the 3D flow field, interpolated onto a Cartesian grid and saved at a time interval of $\Delta t = 0.2$. Each snapshot records five primitive variables: density (ρ), velocities (u, v, w), and temperature (T), in single precision, for a total of $N_x \times N_y \times N_z \times N_{\text{var}} = 625 \times 270 \times 344 \times 5$ data

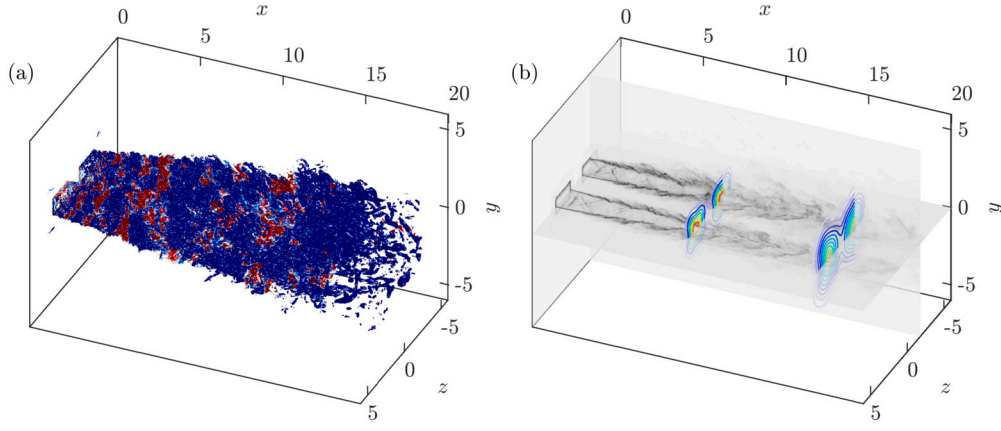


Fig. 2. Instantaneous flow field of the twin-rectangular jet: (a) Q-criterion isosurface, colored by pressure; (b) numerical schlieren on the $y=0$ and $z=-1.8$ planes, with contours of mean streamwise velocity on the $x \in \{8, 16\}$ planes.

points per snapshot. Storage of the database requires 18,392 GB on disk in HDF5 format, or 23,220 GB once loaded into memory. Fig. 2(a) visualizes an instantaneous Q-criterion isosurface, showing the highly turbulent flow field of the high-Reynolds number jet. The colors represent pressure, $p - p_\infty$. Alternating bands of red and blue in the region $x \lesssim 5$ correspond to near-field coherent pressure fluctuations, which radiate sound outward and contribute to far-field noise. Fig. 2(b) shows planar slices of the instantaneous density gradient magnitude, $|\nabla \rho|$, i.e., an artificial schlieren. Shock cells can be observed inside the potential cores. Superimposed on the schlieren are contours of the mean streamwise velocity, \bar{u} . Despite the chaotic instantaneous flow field, \bar{u} displays reflectional symmetries about the major and minor axes, $y=0$ and $z=0$, respectively. The mean flow thus recovers the geometrical symmetries of the twin-rectangular nozzles.

The modal decomposition of axisymmetric jets is typically preceded by an azimuthal Fourier transform, which exploits the rotational invariance of the turbulent statistics. Without loss of generality, the transform reduces the analysis from a single 3D SPOD to one 2D SPOD per azimuthal wavenumber. The absence of azimuthal homogeneity in the twin-rectangular jet precludes such a simplification, thus necessitating a costly 3D analysis, to which PySPOD is ideally suited. To perform SPOD, we assemble the primitive variables into the state vector $\mathbf{q} = [\rho, u, v, w, T]^T$. Since the flow is compressible, we choose the weight matrix

$$\mathbf{W} = \int \int \int \text{diag} \left(\frac{\bar{T}}{\gamma \bar{\rho} M_j^2}, \bar{\rho}, \bar{\rho}, \bar{\rho}, \frac{\bar{\rho}}{\gamma(\gamma-1)\bar{T}M_j^2} \right) dx dy dz, \quad (19)$$

such that the inner product $\langle \mathbf{q}_1, \mathbf{q}_2 \rangle = \mathbf{q}_1^H \mathbf{W} \mathbf{q}_2$ induces the compressible energy norm [46]. Here, $\gamma = 1.4$ is the adiabatic index; $M_j = 1.5$ is the jet Mach number. We select a block size of $N_f = 256$, with 50% overlap, giving $L = 155$ blocks.

The premultiplied SPOD eigenvalue spectra are reported in Fig. 3. The leading eigenvalues show a prominent peak at $f \approx 0.2$, where f is nondimensionalized by the nozzle height and the ambient speed of sound. This peak corresponds to the fundamental screech tone. Screeching of the supersonic twin-rectangular jet stems from acoustic resonance between the nozzle and the shock cells, and has been observed both experimentally [47,45] and numerically [44,48,43] to occur at a similar frequency. The peak at $f \approx 0.2$ persists at least to the second eigenvalue. Large separations between the first, second, and third eigenvalues—termed low-rank behavior [49]—in the frequency range $0.2 \lesssim f \lesssim 0.3$ signal the presence of coherent structures arising from underlying physical instabilities, in this case the screech mechanism.

In Fig. 4, we study these structures by visualizing the pressure component, ϕ_p , of the SPOD modes corresponding to the first four eigenvalues at $f = 0.21$. We recover ϕ_p from the density and temperature




Fig. 3. Premultiplied SPOD eigenvalue spectra of the twin-rectangular jet. The spectra fade from black to white with increasing mode number. Modes corresponding to the highlighted (\bullet) eigenvalues at $f = 0.21$ are reported in Fig. 4.

components of each mode, ϕ_p and ϕ_T , respectively, using the linearized ideal gas equation,

$$\phi_p = \frac{1}{\gamma} \left(\phi_p \bar{T} + \bar{\rho} \phi_T \right). \quad (20)$$

Isovalues of $\text{Re}\{\phi_p\} = \pm 0.0005$ are chosen to highlight the 3D structure of the far-field acoustic waves. Each row in Fig. 4 shows one mode. The left and right columns, 4(a,c,e,g) and 4(b,d,f,h), also display the cross-sectional views of the planes $z = 1.3$ and $y = -0.25$, respectively. These provide insights into the wavepackets within the jet plume, which are well-known to be efficient sources of noise [50]. In Fig. 4(a,b), mode 1 recovers near-perfect antisymmetry about the major-axis plane, $y=0$, and symmetry about the minor-axis plane, $z=0$. In contrast, mode 2 in 4(c,d) is clearly antisymmetric about both planes. Mode 4 in 4(g,h), on the other hand, is symmetric about both planes. The symmetry of mode 3 is difficult to ascertain visually from 4(e,f). In particular, the strong in-phase (4(b,h)) and out-of-phase (4(d)) coupling between the twin jets observed in modes 1, 2, and 4 appear to be lost in mode 3 (4(f)). This is due to insufficient statistical convergence.

A comprehensive discussion of twin-rectangular jet dynamics exceeds the scope of this work. However, our findings illustrate some of the physical insights that can be gleaned from a 3D modal analysis, which PySPOD vastly accelerates.

4.2. SPOD analysis of geophysical datasets

In our second example, we use data from ERA5, a fifth-generation reanalysis dataset produced by the ECMWF that combines model data

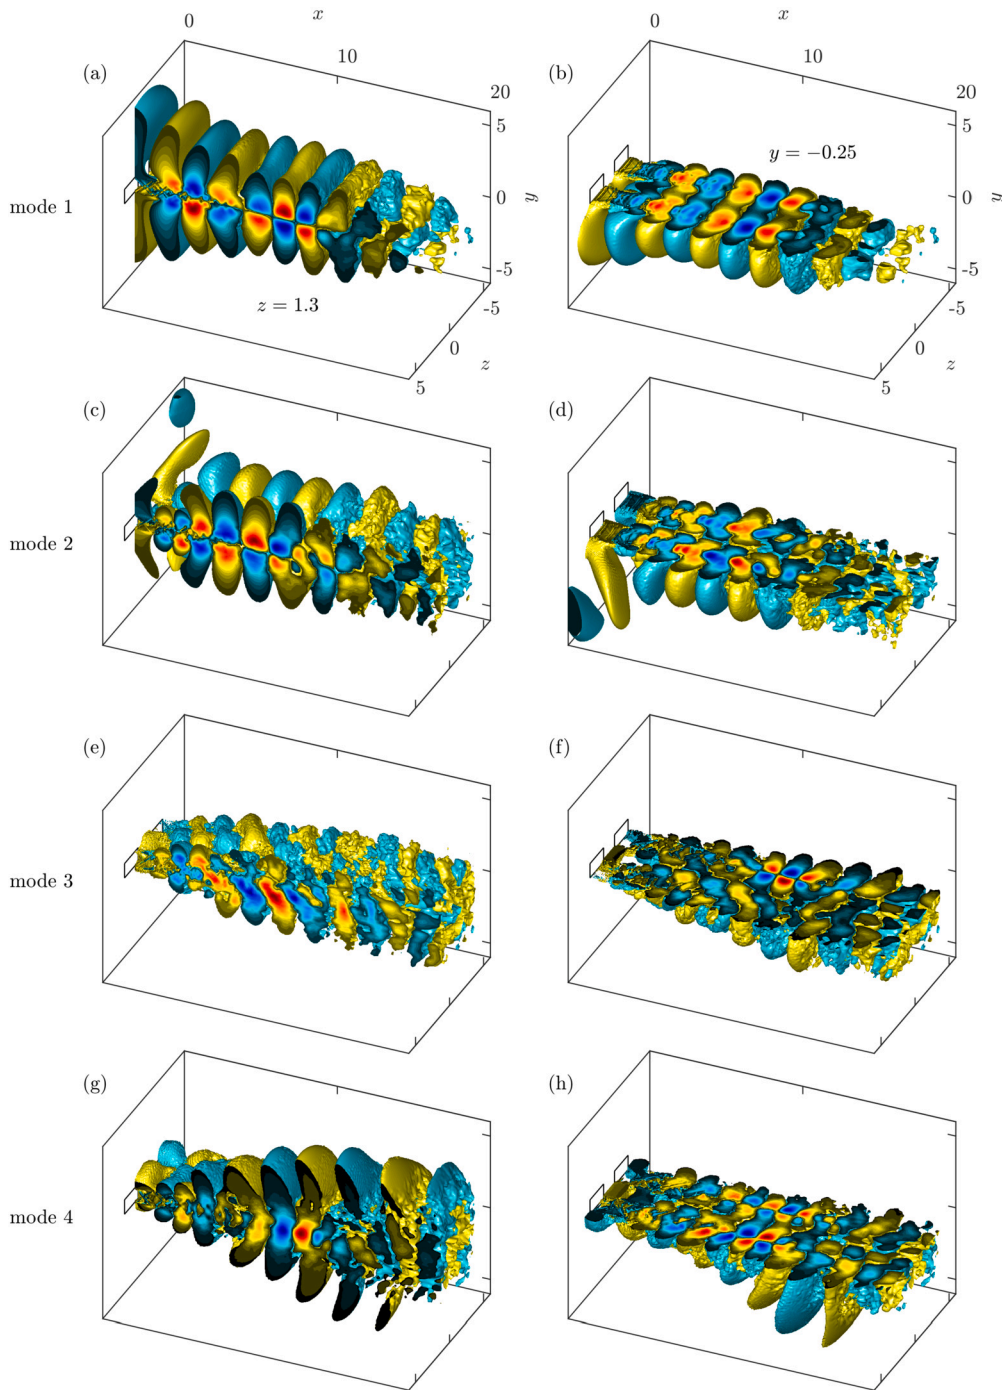


Fig. 4. SPOD modes of the twin-rectangular jet at frequency $f = 0.21$: (a,b) mode 1; (c,d) mode 2; (e,f) mode 3; (g,h) mode 4. Isosurfaces of $\text{Re}\{\phi_p\} = \pm 0.0005$ are shown, along with cross-sections at $z = 1.3$ (left column) and $y = -0.25$ (right column). The corresponding SPOD eigenvalues are highlighted in Fig. 3.

with global observations using the laws of physics to create a complete and consistent global climate and weather dataset. We obtained this data from the Climate Data Store (CDS) [38].

In particular, we consider the horizontal speed of air moving towards the east on 37 pressure levels (i.e., vertical levels) spanning from January 1940 to December 2022 [38]. This quantity is also referred to as U component of the horizontal wind velocity, and its unit is meters per second. The dataset contains 727,584 time snapshots on a 3D grid of dimension $1440 \times 721 \times 37$. The total size of this dataset is 51,745 GB, corresponding to 27.9 trillion data points and 103,490 GB in memory when using single-precision floating-point, as for this example, and 199,452 GB in double-precision. In Fig. 5, we depict the

horizontal wind velocity adopted for pressure level 1 (top left), 12 (top right), 24 (bottom left), and 37 (bottom right). dataset.

The SPOD algorithm uses 10-year data blocks, that corresponds to a block size of $N_f = 87,600$ time snapshots, resulting in a total of $L = 8$ blocks, where we used 0% overlapping.

This SPOD configuration was chosen to capture low-frequency atmospheric modes. More specifically, the analysis of this dataset aims to capture the quasi-biennial oscillation (QBO) that has an approximate period of 2 to 2.5 years, as reported in [6]. This atmospheric oscillation is characterized by quasi-periodic reversals of the zonal-mean zonal winds in the equatorial stratosphere – see also [51]. QBO has important implications on teleconnections, influencing weather patterns

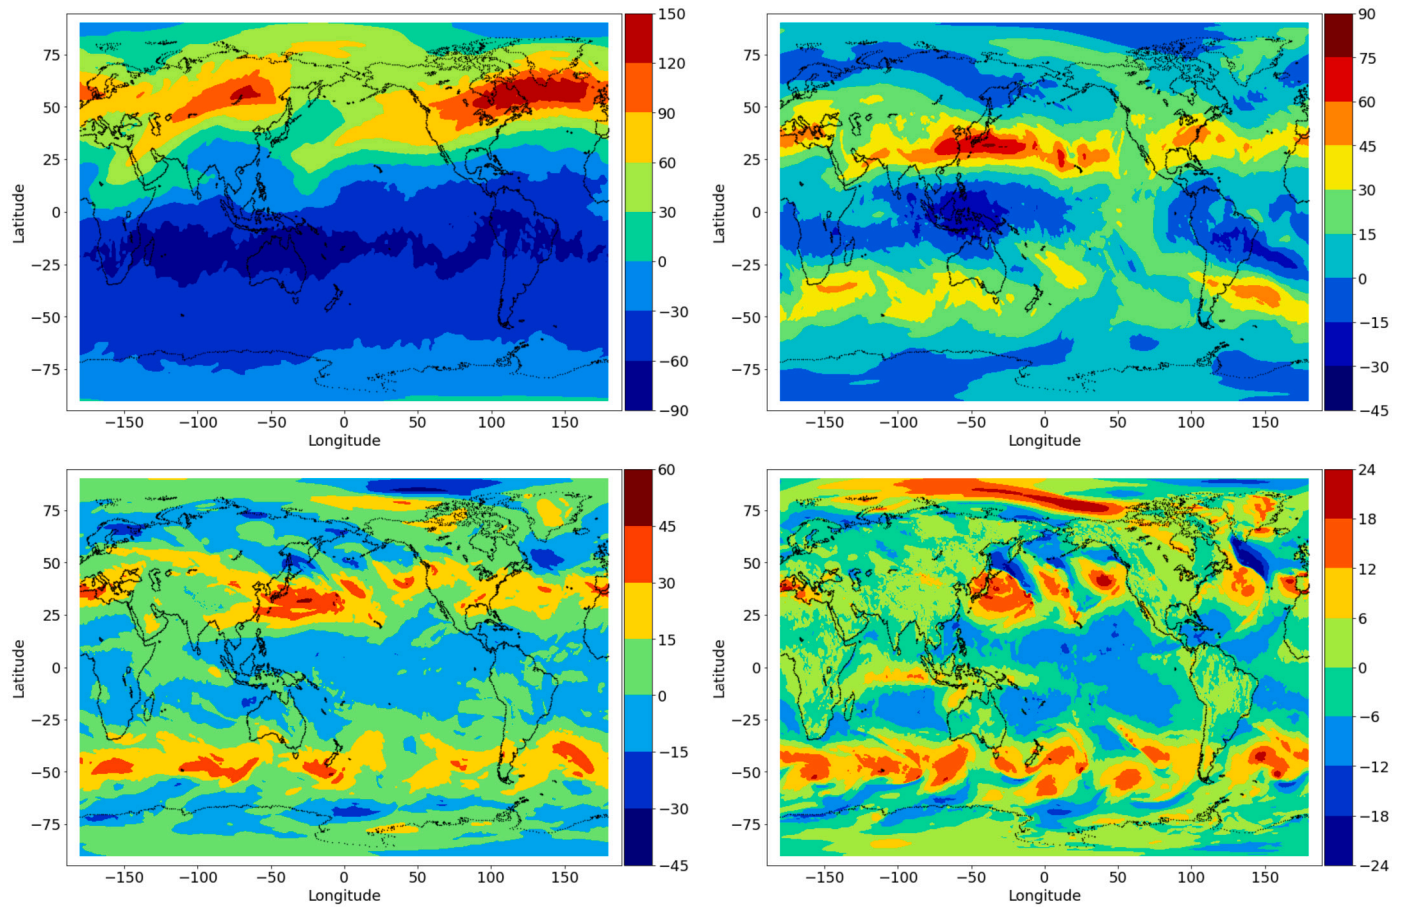


Fig. 5. U component of the wind velocity for pressure (i.e., vertical) levels 1 (top left), 12 (top right), 24 (bottom left), and 37 (bottom right), at midnight (00:00) on the 1st of January 2010. Level 1 corresponds to a pressure of 1 millibars, level 12 to 125 millibars, level 24 to 600 millibars, and level 37 to 1000 millibars.




Fig. 6. Eigenvalue spectra vs. period (in days). The pink vertical line denotes the peak associated to the QBO, whose associated mode is depicted in Fig. 7. We can also notice other high-frequency peaks, related to yearly, sub-yearly, daily and sub-daily patterns.

in the Northern Hemisphere, and the tropics (including tropical precipitation) [52]. Its most distinctive feature is a latitudinal band of the U component of the wind velocity in the tropical region ($\pm 20^\circ$ latitude).

Indeed, Fig. 6 shows the eigenvalue spectra for the U component of the wind velocity, where the leading eigenvalue show a prominent

peak at period $T = 912.5$ days. This peak corresponds to the QBO, and shows as this phenomenon exhibits low-rank behavior, since the energy of the leading eigenvalue is remarkably separated from the other eigenvalues. Fig. 7 shows the highly-coherent three-dimensional structure of the leading mode. This manifests as a latitudinal band of the U component of the wind velocity in the tropical region, that is the signature of QBO. The results are consistent with those in [6], albeit the data adopted there was ERA-20C [53], that has significantly coarser spatial and temporal resolution, but longer coverage (from 1900 to 2010). In our results, we also observe several high-frequency peaks in the eigenvalues (periods shorter than 1 day) thanks to the improved temporal resolution of our dataset.

As for the example presented in section 4.1, a comprehensive discussion of QBO dynamics is outside the scope of this work. Yet, we remark that the analysis outlined in this section would have been extremely challenging without the parallel (distributed) implementation presented in this work.

5. Scalability

The scalability results were obtained on the geophysical dataset introduced and analyzed in section 4.2. Here, we discuss the scalability setup (section 5.1), strong (section 5.2) and weak (section 5.3) scalability studies conducted using that dataset.

5.1. Scalability setup

To test the scalability of the PySPOD library, we used the Shaheen II supercomputer, a Cray XC40 system hosted by King Abdullah University of Science and Technology (KAUST) [54]. Each of Shaheen's

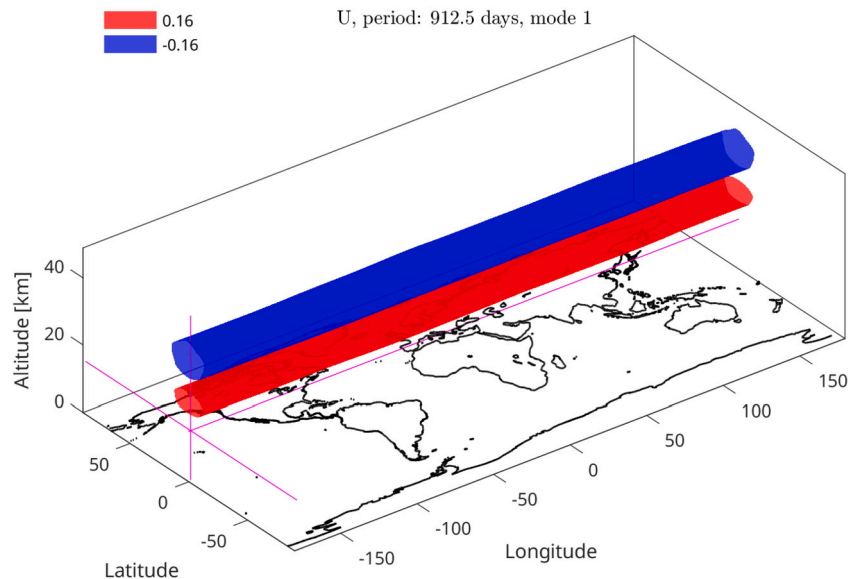


Fig. 7. Real part of the leading three-dimensional SPOD mode for the U component of the wind velocity; period of 912.5 days.

6,174 nodes features two Intel Haswell (Xeon E5-2698v3) CPUs with 16 cores each and 128 GB of memory. The nodes are connected via a Cray Aries interconnect with Dragonfly topology. While the nodes were allocated in exclusive mode by the scheduler, the network was shared with other users, as is typical in a production environment. To alleviate the memory capacity bottleneck and maximize the I/O bandwidth, we run PySPOD with 4 MPI processes per compute node. We do not use multithreading.

For storage, we utilized the Cray Sonexion® 2000 Storage System, which offers over 16 PB of usable capacity using 72 Scalable Storage Units (SSUs), 144 Object Storage Services (OSSs) and 144 Object Storage Targets (OSTs). The system features 5,988 4 TB disks. The theoretical performance of this storage exceeds 500 GB/s [55], and published application benchmarks have shown that applications such as WRF can achieve a bandwidth of 25–35 GB/s, including communication, or 65 GB/s of raw I/O performance measured on aggregators when using 144 OSTs and 4 MPI ranks per node. Similar results have been obtained with NPB benchmarks [56].

We stored both datasets on the Lustre storage system described in section 5.1. Since many processes may need to access each file and the datasets are large, we used a stripe count of 144. However, we did not use striping for output to limit each process to only communicate with one Object Storage Target (OST) and reduce contention, given that each process writes a single file. In both cases, we used the default stripe size of 1 MB.

We report the following timings for PySPOD: time for I/O (combined reading and writing of 5 leading modes to disk), computation of the DFT (step 2 in Fig. 1), computation of the inner product and associated matrix \hat{M} (step 3 in Fig. 1), eigenvalue decomposition (step 5 in Fig. 1), and computation of the SPOD modes (just prior to step 6 in Fig. 1). Given the variability of performance on Dragonfly networks [57] and shared file systems [58], we report arithmetic averages based on 5 repetitions.

We employed PySPOD release 2.0.0, along with the following packages: mpi4py 3.1.4, netCDF4 1.6.3, NumPy 1.24.2, SciPy 1.10.1, and Xarray 2023.2.0. Cray-provided modules for Python 3.10.1 and Cray MPICH 7.7.18.

5.2. Strong scalability

To perform the strong scalability analysis, we used a 2D version of the dataset presented in section 4.2, corresponding to the 10 hPa




Fig. 8. Strong scalability of PySPOD using horizontal speed of air moving towards the east data from January 1940 to December 2022, using from 256 to 8,192 processes. Dashed lines represent ideal scalability for each component.

pressure level [59]. This dataset contains 727,584 time snapshots on a 1440×721 grid and is 1,407 GB in size when stored on the disk in netCDF format. The data set contains a total of 755.4 billion data points in single-precision floating point, which amounts to 5,628 GB in memory once stored in double-precision floating-point for this study. As we scale from 256 up to 8,192 processes (64 to 2,048 nodes), this corresponds to between 21.99 GB and 0.68 GB, or 4,055 and 126 spatial points over time per process. This provides a broad range of scenarios. In all cases, all the temporal data is split into 1-year (8,760 snapshots) blocks. This choice of block size is different to that of section 4.2 for the sake of using consistent configurations in both strong and weak scalability analyses.

As shown in Fig. 8, the outcomes differ depending on the task that is being performed (different curve colors in Fig. 8). In partic-

ular, we achieved a satisfactory speedup for I/O as the number of processes increased (teal curve), with the peak average read bandwidth of 73.76 GB/s (4,096 processes) and the maximum individual measured bandwidth of 88.04 GB/s. Even though the read speed peaks at 4,096 processes, the combined I/O time (reading and writing) is minimally lower when using 2,048 processes. The reported timings include communication (i.e., the second phase of the two-phase I/O) and should be compared to the 25–35 GB/s figure quoted in section 5.1. We consider the results extremely competitive on this hardware.

The FFT calculation (violet curve), the second most expensive part of the SPOD algorithm, exhibited surprising behavior. The speedup is measured at $50 \times$ when the number of processes is increased by $32 \times$ (from 256 to 8,192 processes). This phenomenon has been previously observed on the same supercomputer [60] and can be explained by cache effects, frequency throttling, and possibly problem size-dependent NumPy optimizations.

The computation of matrix \tilde{M} (brown curve), a collective communication-heavy routine, stops scaling past 2,048 cores; however, at this point, it only takes around 7 seconds. Similarly, the eigenvalue computation (orange curve), which is problem size-dependent and not parallelized (see rationale in section 3.1), is represented by a horizontal line, i.e., its cost depends on the dimensions of the data and not the number of processes. On the other hand, the calculation of SPOD modes (green curve) scaled well despite its cost being measured in single seconds.

In terms of overall relative efficiency,¹ we calculate it to be 74% when using 2,048 cores and 38% or lower once increasing to 4,096 cores and beyond. Even though the efficiency is only 15% when using 8,192 processes, we note that in this scenario, the total runtime is only around 127 seconds. The shortest runtime was achieved when using 4,096 processes and was around 98 seconds, 4 seconds faster than when using 2,048 processes. Based on the presented performance characteristics and emphasizing that I/O is the dominating cost, we suggest using as many processes as required to saturate the I/O bandwidth. Increasing the dedicated resources past the I/O saturation point will significantly reduce the parallel efficiency.

5.3. Weak scalability

For the weak scalability analysis, we used the ERA5 dataset containing the horizontal U component of wind speed on 37 pressure levels described in section 4.2. As mentioned there, the total size of this data set is 51,745 GB, and we used between 10 and 80 years of data, corresponding to between 3.4 and 26.8 trillion data points stored using 6,273 GB to 49,863 GB of disk space. The corresponding memory required in double-precision floating-point that was adopted in this study ranged from 25,092 GB to 199,452 GB. To maintain a constant load per process, we use 10 years of data per 2,048 processes (512 nodes) and scale up to 16,384 processes (4,096 nodes), where we use 80 years of the horizontal U component of wind speed data. Similarly to the strong scalability analysis, we split the data temporally into 1-year blocks. As shown in Fig. 9, we observe satisfactory I/O behavior (teal curve). The bandwidth peaks at 80.41 GB/s when using 12,288 cores. DFT (violet curve) and modes (green curve) computation also remain efficient as more data and processes are used. However, the computation of matrix \tilde{M} shows worse scaling and its timing displays significant variation in repeated executions (brown curve). In an environment with shared interconnect, this result points to the increased collective communication cost at scale. Additionally, the eigenvalue computation (orange curve), a serial component, becomes more expensive as the problem size grows.

Overall, the efficiency is 64% when using 6,144 cores. Despite the decrease in efficiency to 38% when using 16,384 cores, the overall runtime for this scenario is under 20 minutes, which we find acceptable

¹ The relative parallel efficiency is calculated as $E_p = (256 \times T_{256}) / (P \times T_p)$, where T_p is the wall-clock time corresponding to P processes.




Fig. 9. Weak scalability of PySPOD using the hourly horizontal speed of air data on 37 pressure levels data. 10 years of data per 2,000 processes, between January 1940 and December 2020 (when using 16,384 processes).

given the sheer size of the dataset (49,863 GB on disk and 199,452 GB in memory).

6. Discussion and conclusions

The new parallel SPOD algorithm allows modal decompositions that were extremely challenging if not impossible with the serial algorithms available. We were able to compute SPOD decompositions up to 199 TB using HPC platforms, and exploiting the scalability and performance of the parallel algorithm. In particular, the key novel aspect is the I/O handling dictated by the smart data layout that was devised and implemented. This allowed preserving all time operations (more specifically the DFT), and trivially distribute across MPI ranks the spatial component of the data. The results reported in section 4 show the power of the package in providing results on big data, enabled by the scalability performance shown in section 5. The latter were possible thanks to an efficient implementation of I/O, that also allowed for a reduction in terms of memory consumption. The new package may allow unlocking new physics from big data that was not possible to analyze before.

CRedit authorship contribution statement

Marcin Rogowski: Conceptualization, Data curation, Formal analysis, Investigation, Methodology, Software, Visualization, Writing – original draft, Writing – review & editing. **Brandon C.Y. Yeung:** Data curation, Formal analysis, Software, Writing – original draft, Writing – review & editing. **Oliver T. Schmidt:** Methodology, Writing – original draft, Writing – review & editing. **Romit Maulik:** Methodology, Writing – original draft, Writing – review & editing. **Lisandro Dalcin:** Conceptualization, Data curation, Formal analysis, Funding acquisition, Investigation, Methodology, Project administration, Resources, Software, Supervision, Validation, Visualization, Writing – original draft, Writing – review & editing. **Matteo Parsani:** Funding acquisition, Project administration, Resources, Supervision, Writing – original draft, Writing – review & editing. **Gianmarco Mengaldo:** Conceptualization, Data curation, Formal analysis, Funding acquisition, Investigation, Methodology, Project administration, Resources, Software, Supervision, Validation, Visualization, Writing – original draft, Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data for the turbulent jet is confidential. Atmospheric data is from ERA5 where the scripts for downloading can be made available. Note the data size is of the order of several terabytes.

Acknowledgements

Marcin Rogowski, Lisandro Dalcin, and Matteo Parsani acknowledge support from King Abdullah University of Science and Technology (KAUST) award BAS/1/1663-01-01. Brandon C. Y. Yeung and Oliver T. Schmidt gratefully acknowledge support from Office of Naval Research award N00014-23-1-2457, under the supervision of Dr. Steve Martens. Romit Maulik was supported by U.S. DOE ASCR Award Data-intensive Scientific Machine Learning: DE-FOA-2493. Gianmarco Mengaldo was supported by MOE Tier 2 grant 22-5191-A0001-0: "Prediction-to-Mitigation with Digital Twins of the Earth's Weather". The authors are thankful to the KAUST Supercomputing Laboratory for their computing resources. LES calculations were carried out on the "Onyx" Cray XC40/50 system in ERDC DSRC, using allocations provided by DoD HPCMP.

References

- [1] G. Berkooz, P. Holmes, J.L. Lumley, The proper orthogonal decomposition in the analysis of turbulent flows, *Annu. Rev. Fluid Mech.* 25 (1993) 539–575.
- [2] R. Maulik, G. Mengaldo, PyParSVD: a streaming, distributed and randomized singular-value-decomposition library, in: 2021 7th International Workshop on Data Analysis and Reduction for Big Scientific Data (DRBSD-7), IEEE, 2021, pp. 19–25.
- [3] P.J. Schmid, Dynamic mode decomposition of numerical and experimental data, *J. Fluid Mech.* 656 (2010) 5–28.
- [4] J.L. Lumley, The structure of inhomogeneous turbulent flows, in: A.M. Yaglom, V.I. Tatarski (Eds.), *Atmospheric Turbulence and Radio Propagation*, 1967, pp. 166–178.
- [5] A. Towne, O.T. Schmidt, T. Colonius, Spectral proper orthogonal decomposition and its relationship to dynamic mode decomposition and resolvent analysis, *J. Fluid Mech.* 847 (2018) 821–867.
- [6] O.T. Schmidt, G. Mengaldo, G. Balsamo, N.P. Wedi, Spectral empirical orthogonal function analysis of weather and climate data, *Mon. Weather Rev.* 147 (2019) 2979–2995.
- [7] O.T. Schmidt, T. Colonius, Guide to spectral proper orthogonal decomposition, *AIAA J.* 58 (2020) 1023–1033.
- [8] A. Lario, R. Maulik, O.T. Schmidt, G. Rozza, G. Mengaldo, Neural-network learning of SPOD latent dynamics, *J. Comput. Phys.* 468 (2022) 111475.
- [9] O.T. Schmidt, Spectral proper orthogonal decomposition using multitaper estimates, *Theor. Comput. Fluid Dyn.* 36 (2022) 741–754.
- [10] G. Mengaldo, R. Maulik, PySPOD: a python package for spectral proper orthogonal decomposition (SPOD), *J. Open Sour. Softw.* 6 (2021) 2862, <https://doi.org/10.21105/joss.02862>.
- [11] D. Kosambi, *Statistics in function space*, in: DD Kosambi, Springer, 2016, pp. 115–123.
- [12] M. Loeve, *Probability Theory*, Courier Dover Publications, 2017.
- [13] J.L. Lumley, *Stochastic Tools in Turbulence*, Courier Corporation, 2007.
- [14] P. Welch, The use of fast Fourier transform for the estimation of power spectra: a method based on time averaging over short, modified periodograms, *IEEE Trans. Audio Electroacoust.* 15 (1967) 70–73.
- [15] L. Sirovich, Turbulence and the dynamics of coherent structures. I. Coherent structures, *Q. Appl. Math.* 45 (1987) 561–571.
- [16] A. Nekkanti, O.T. Schmidt, Frequency-time analysis, low-rank reconstruction and denoising of turbulent flows using SPOD, *J. Fluid Mech.* 926 (2021).
- [17] T. Sayadi, P.J. Schmid, Parallel data-driven decomposition algorithm for large-scale datasets: with application to transitional boundary layers, *Theor. Comput. Fluid Dyn.* 30 (2016) 415–428.
- [18] L. Dalcin, R. Paz, M.A. Storti, MPI for python, *J. Parallel Distrib. Comput.* 65 (2005) 1108–1115.
- [19] L. Dalcin, Y.-L.L. Fang, mpi4py: status update after 12 years of development, *Comput. Sci. Eng.* 23 (2021) 47–54.
- [20] P. Crandall, R. Aydt, A. Chien, D. Reed, Input/Output characteristics of scalable parallel applications, in: *Supercomputing '95: Proceedings of the 1995 ACM/IEEE Conference on Supercomputing*, 1995, p. 59.

- [21] J.M. Del Rosario, R. Bordawekar, A. Choudhary, Improved parallel I/O via a two-phase run-time access strategy, *Comput. Archit. News* 21 (1993) 31–38.
- [22] Q. Liu, J. Logan, Y. Tian, H. Abbasi, N. Podhorszki, J.Y. Choi, S. Klasky, R. Tchoua, J. Lofstead, R. Oldfield, M. Parashar, N. Samatova, K. Schwan, A. Shoshani, M. Wolf, K. Wu, W. Yu, Hello ADIOS: the challenges and lessons of developing leadership class I/O frameworks, *Concurr. Comput., Pract. Exp.* 26 (2014) 1453–1473, <https://doi.org/10.1002/cpe.3125>.
- [23] J. Bent, G. Grider, B. Kettering, A. Manzanares, M. McClelland, A. Torres, A. Torrez, Storage challenges at Los Alamos National Lab, in: *2012 IEEE 28th Symposium on Mass Storage Systems and Technologies (MSST)*, 2012, pp. 1–5.
- [24] R. Thakur, W. Gropp, E. Lusk, On implementing MPI-IO portably and with high performance, in: *Proceedings of the Sixth Workshop on I/O in Parallel and Distributed Systems*, 1999, pp. 23–32.
- [25] R. Thakur, W. Gropp, E. Lusk, Optimizing noncontiguous accesses in MPI-IO, *Parallel Comput.* 28 (2002) 83–105.
- [26] The HDF Group, Hierarchical data format version 5, <http://www.hdfgroup.org/HDF5>, 2000–2010.
- [27] J. Li, W. Liao, A. Choudhary, R. Ross, R. Thakur, W. Gropp, R. Latham, A. Siegel, B. Gallagher, M. Zingale, Parallel netCDF: a high-performance scientific I/O interface, in: *SC '03: Proceedings of the 2003 ACM/IEEE Conference on Supercomputing*, 2003, p. 39.
- [28] W.F. Godoy, N. Podhorszki, R. Wang, C. Atkins, G. Eisenhauer, J. Gu, P. Davis, J. Choi, K. Gernaschewski, K. Huck, A. Huebl, M. Kim, J. Kress, T. Kurc, Q. Liu, J. Logan, K. Mehta, G. Ostrouchov, M. Parashar, F. Poeschel, D. Pugmire, E. Suchyta, K. Takahashi, N. Thompson, S. Tsutsumi, L. Wan, M. Wolf, K. Wu, S. Klasky, ADIOS 2: the adaptable input output system. A framework for high-performance data management, *SoftwareX* 12 (2020) 100561, <https://doi.org/10.1016/j.softx.2020.100561>.
- [29] C. Xu, S. Byna, V. Venkatesan, R. Sisneros, O. Kulkarni, M. Chaarawi, K. Chadalavada, LIOPProf: exposing Lustre file system behavior for I/O middleware, in: *2016 Cray User Group Meeting*, 2016.
- [30] P.M. Dickens, J.S. Logan, A high performance implementation of MPI-IO for a Lustre file system environment, *Concurr. Comput., Pract. Exp.* 22 (2010).
- [31] T. Balle, P. Johnsen, Improving I/O performance of the weather research and forecast (WRF) model, in: *Proceedings of the Cray User Group Meeting*, London, United Kingdom, 2016.
- [32] W. Yu, J.S. Vetter, H. Sarp Oral, Performance characterization and optimization of parallel I/O on the Cray XT, in: *2008 IEEE International Symposium on Parallel and Distributed Processing*, 2008, pp. 1–11.
- [33] M. Howison, Q. Koziol, D. Knaak, J. Mainzer, J. Shalf, Tuning HDF5 for Lustre file systems, <https://www.osti.gov/biblio/1050648>, 2010.
- [34] B. Behzad, H.V.T. Luu, J. Huchette, S. Byna, Prabhat, R. Aydt, Q. Koziol, M. Snir, Taming parallel I/O complexity with auto-tuning, in: *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis, SC '13*, Association for Computing Machinery, New York, NY, United States, 2013.
- [35] B. Behzad, S. Byna, S.M. Wild, M. Prabhat, M. Snir, Improving parallel I/O autotuning with performance modeling, in: *Proceedings of the 23rd International Symposium on High-Performance Parallel and Distributed Computing, HPDC '14*, Association for Computing Machinery, New York, NY, United States, 2014, pp. 253–256.
- [36] H. Luu, M. Winslett, W. Gropp, R. Ross, P. Carns, K. Harms, M. Prabhat, S. Byna, Y. Yao, A multipatform study of I/O behavior on petascale supercomputers, in: *Proceedings of the 24th International Symposium on High-Performance Parallel and Distributed Computing, HPDC '15*, Association for Computing Machinery, New York, NY, United States, 2015, pp. 33–44.
- [37] Y. Yu, D.H. Rudd, Z. Lan, N.Y. Gnedin, A. Kravtsov, J. Wu, Improving parallel IO performance of cell-based AMR cosmology applications, in: *2012 IEEE 26th International Parallel and Distributed Processing Symposium*, 2012, pp. 933–944.
- [38] H. Hersbach, B. Bell, P. Berrisford, G. Biavati, A. Horányi, J. Muñoz Sabater, J. Nicolas, C. Peubey, R. Radu, I. Rozum, D. Schepers, A. Simmons, C. Soci, D. Dee, J.-N. Thépaut, ERA5 hourly data on pressure levels from 1940 to present, Copernicus Climate Change Service (C3S) Climate Data Store (CDS), <https://doi.org/10.24381/cds.bd0915c6>. (Accessed 19 June 2023), 2018.
- [39] M. Parsani, R. Boukharfane, I.R. Nolasco, D.C. Del Rey Fernández, S. Zampini, B. Hadri, L. Dalcin, High-order accurate entropy-stable discontinuous collocated Galerkin methods with the summation-by-parts property for compressible CFD frameworks: scalable ssdc algorithms and flow solver, *J. Comput. Phys.* 424 (2021) 109844, <https://doi.org/10.1016/j.jcp.2020.109844>.
- [40] D. Moxey, C.D. Cantwell, Y. Bao, A. Cassinelli, G. Castiglioni, S. Chun, E. Juda, E. Kazemi, K. Lackhove, J. Marcon, et al., Nektar++: enhancing the capability and application of high-fidelity spectral/hp element methods, *Comput. Phys. Commun.* 249 (2020) 107110.
- [41] G.A. Brès, F.E. Ham, J.W. Nichols, S.K. Lele, Unstructured large-eddy simulations of supersonic jets, *AIAA J.* 55 (April 2017) 1164–1184, <https://doi.org/10.2514/1.J055084>.
- [42] H. Hersbach, B. Bell, P. Berrisford, S. Hirahara, A. Horányi, J. Muñoz-Sabater, J. Nicolas, C. Peubey, R. Radu, D. Schepers, et al., The ERA5 global reanalysis, *Q. J. R. Meteorol. Soc.* 146 (2020) 1999–2049.
- [43] B.C.Y. Yeung, O.T. Schmidt, G.A. Brès, Three-dimensional spectral POD of supersonic twin-rectangular jet flow, *AIAA Paper* 2022–3345, June 2022, <https://doi.org/10.2514/6.2022-3345>.

- [44] G.A. Brès, B.C.Y. Yeung, O.T. Schmidt, A. Esfahani, N. Webb, M. Samimy, T. Colonius, Towards large-eddy simulations of supersonic jets from twin rectangular nozzle with plasma actuation, *AIAA Paper 2021–2154*, August 2021, <https://doi.org/10.2514/6.2021-2154>.
- [45] M. Samimy, N. Webb, A. Esfahani, R. Leahy, Perturbation-based active flow control in overexpanded to underexpanded supersonic rectangular twin jets, *J. Fluid Mech.* 959 (2023) 1–34, <https://doi.org/10.1017/jfm.2023.139>.
- [46] B.-T. Chu, On the energy transfer to small disturbances in fluid flow (part I), *Acta Mech.* 1 (September 1965) 215–234, <https://doi.org/10.1007/BF01387235>.
- [47] J. Jeun, A. Karnam, G.J. Wu, S.K. Lele, F. Baier, E.J. Gutmark, Aeroacoustics of twin rectangular jets including screech: large-eddy simulations with experimental validation, *AIAA J.* 60 (November 2022) 6340–6360, <https://doi.org/10.2514/1.J060895>.
- [48] G.A. Brès, S.T. Bose, C.B. Ivey, M. Emory, F. Ham, GPU-accelerated large-eddy simulations of supersonic jets from twin rectangular nozzle, *AIAA Paper 2022–3001*, June 2022, <https://doi.org/10.2514/6.2022-3001>.
- [49] O.T. Schmidt, A. Towne, G. Rigas, T. Colonius, G.A. Brès, Spectral analysis of jet turbulence, *J. Fluid Mech.* 855 (2018) 953–982, <https://doi.org/10.1017/jfm.2018.675>.
- [50] P. Jordan, T. Colonius, Wave packets and turbulent jet noise, *Annu. Rev. Fluid Mech.* 45 (2013) 173–195, <https://doi.org/10.1146/annurev-fluid-011212-140756>.
- [51] M. Baldwin, L. Gray, T. Dunkerton, K. Hamilton, P. Haynes, W.J. Randel, J.R. Holton, M. Alexander, I. Hirota, T. Horinouchi, et al., The quasi-biennial oscillation, *Rev. Geophys.* 39 (2001) 179–229.
- [52] L.J. Gray, J.A. Anstey, Y. Kawatani, H. Lu, S. Osprey, V. Schenzinger, Surface impacts of the quasi biennial oscillation, *Atmos. Chem. Phys.* 18 (2018) 8227–8247.
- [53] P. Poli, H. Hersbach, D.P. Dee, P. Berrisford, A.J. Simmons, F. Vitart, P. Laloyaux, D.G. Tan, C. Peubey, J.-N. Thépaut, et al., Era-20c: an atmospheric reanalysis of the twentieth century, *J. Climate* 29 (2016) 4083–4097.
- [54] B. Hadri, S. Kortas, S. Feki, R. Khurram, G. Newby, Overview of the KAUST’s Cray X40 System – Shaheen II, in: *Proceedings of the Cray User Group Meeting*, Chicago, United States, 2015.
- [55] B. Hadri, S. Kortas, R. Fiedler, G.S. Markomanolis, Regression testing on Shaheen Cray XC40: implementation and lessons learned, in: *Proceedings of the Cray Users Group Meeting (CUG2017)*, 2017.
- [56] G.S. Markomanolis, B. Hadri, R. Khurram, S. Feki, Scientific applications performance evaluation on burst buffer, in: J.M. Kunkel, R. Yokota, M. Taufer, J. Shalf (Eds.), *High Performance Computing*, Springer International Publishing, Cham, 2017, pp. 701–711.
- [57] S. Chunduri, K. Harms, S. Parker, V. Morozov, S. Oshin, N. Cherukuri, K. Kumaran, Run-to-run variability on Xeon Phi based Cray XC systems, in: *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, SC ’17*, Association for Computing Machinery, New York, NY, United States, 2017.
- [58] J. Lofstead, F. Zheng, Q. Liu, S. Klasky, R. Oldfield, T. Kordenbrock, K. Schwan, M. Wolf, Managing variability in the IO performance of petascale storage systems, in: *SC ’10: Proceedings of the 2010 ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*, 2010, pp. 1–12.
- [59] H. Hersbach, B. Bell, P. Berrisford, G. Biavati, A. Horányi, J. Muñoz Sabater, J. Nicolas, C. Peubey, R. Radu, I. Rozum, D. Schepers, A. Simmons, C. Soci, D. Dee, J.-N. Thépaut, ERA5 hourly data on single levels from 1940 to present, Copernicus Climate Change Service (C3S) Climate Data Store (CDS), <https://doi.org/10.24381/cds.adbb2d47>. (Accessed 19 June 2023), 2018.
- [60] L. Dalcin, M. Mortensen, D.E. Keyes, Fast parallel multidimensional FFT using advanced MPI, *J. Parallel Distrib. Comput.* 128 (2019) 137–150, <https://doi.org/10.1016/j.jpdc.2019.02.006>.