



Delft University of Technology

## A pencil-distributed finite-difference solver for extreme-scale calculations of turbulent wall flows at high Reynolds number

Diez Sanhueza, Rafael; Peeters, Jurriaan; Costa, Pedro

### DOI

[10.1016/j.cpc.2025.109811](https://doi.org/10.1016/j.cpc.2025.109811)

### Publication date

2025

### Document Version

Final published version

### Published in

Computer Physics Communications

### Citation (APA)

Diez Sanhueza, R., Peeters, J., & Costa, P. (2025). A pencil-distributed finite-difference solver for extreme-scale calculations of turbulent wall flows at high Reynolds number. *Computer Physics Communications*, 316, Article 109811. <https://doi.org/10.1016/j.cpc.2025.109811>

### Important note


To cite this publication, please use the final published version (if applicable).  
Please check the document version above.

### Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

### Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.  
We will remove access to the work immediately and investigate your claim.



## Computational Physics

## A pencil-distributed finite-difference solver for extreme-scale calculations of turbulent wall flows at high Reynolds number

Rafael Diez Sanhueza<sup>ID,\*</sup>, Jurriaan Peeters, Pedro Costa*Process & Energy Department, Delft University of Technology, Leeghwaterstraat 39, 2628 CB Delft, the Netherlands*

## ARTICLE INFO

The review of this paper was arranged by Prof. Peter Vincent

## Keywords:

Direct numerical simulation  
Distributed Poisson solver  
GPU acceleration  
High-performance computing

## ABSTRACT

We present a computational method for extreme-scale simulations of incompressible turbulent wall flows at high Reynolds numbers. The numerical algorithm extends a popular method for solving second-order finite differences Poisson/Helmholtz equations using a pencil-distributed parallel tridiagonal solver to improve computational performance at scale. The benefits of this approach were investigated for high-Reynolds-number turbulent channel flow simulations, with up to about 80 billion grid points and 1024 GPUs on the European flagship supercomputers Leonardo and LUMI. An additional GPU porting effort of the entire solver had to be undertaken for the latter. Our results confirm that, while 1D domain decompositions are favorable for smaller systems, they become inefficient or even impossible at large scales. This restriction is relaxed by adopting a pencil-distributed approach. The results show that, at scale, the revised Poisson solver is about twice as fast as the baseline approach with the full-transpose algorithm for 2D domain decompositions. Strong and weak scalability tests show that the performance gains are due to the lower communication footprint. Additionally, to secure high performance when solving for wall-normal implicit diffusion, we propose a reworked flavor of parallel cyclic reduction (PCR) that is split into pre-processing and runtime steps. During pre-processing, small sub-arrays with independent 1D coefficients are computed by parallel GPU threads, without any global GPU communication. Then, at runtime, the reworked PCR enables a fast solution of implicit 1D diffusion without computational overhead. Our results show that the entire numerical solver, coupled with the PCR algorithm, enables extreme-scale simulations with 2D pencil decompositions, which do not suffer performance losses even when compared to the best 1D slab configurations available for smaller systems.

## 1. Introduction

Turbulent flows at high Reynolds numbers are among the most complex and prevalent problems in engineering and physics. While numerous flows at low Reynolds numbers may be studied using simple analytical or numerical models, many flows found in nature and industry operate at high Reynolds numbers in a turbulent regime. These flows exhibit complex behavior, which is difficult to predict using existing correlations or upscaled models. Fundamental understanding of turbulence that can improve engineering models requires direct numerical simulations (DNS), where the chaotic and multi-scale flow dynamics are fully resolved up to the smallest temporal and spatial scales. While this has a large computational cost, the exponential growth in computing power during the last decades, along with the development of efficient numerical methods, have enabled the simulation of flows at increasingly high Reynolds numbers. Indeed, following the developments since the first

DNS of isotropic turbulence by Orszag and Patterson [1] in 1972, it is now possible to perform large-scale simulations with trillions of grid points in modern supercomputers [2,3].

We are experiencing yet another breakthrough, thanks to the recent proliferation of general-purpose GPU-based supercomputers [4,5]. GPUs are known to perform well in tasks that only require simple arithmetic operations or RAM access patterns [6], which are common in computational fluid dynamics (CFD). They have much higher throughput than CPUs by allowing many parallel threads to perform the same operation per clock cycle. Thus, when successfully ported, GPU-accelerated solvers can easily outperform multi-core CPU solvers [7,8], enabling the numerical solution to complex problems at much lower costs. However, large-scale CFD problems need to operate at scale, in a distributed-memory paradigm where the data is distributed among many GPUs. This introduces new challenges, as many-GPU systems are more prone to feature performance bottlenecks associated with intra and internode commu-

\* Corresponding author.

E-mail address: [R.G.DiezSanhueza-1@tudelft.nl](mailto:R.G.DiezSanhueza-1@tudelft.nl) (R. Diez Sanhueza).

<https://doi.org/10.1016/j.cpc.2025.109811>

Received 9 February 2025; Received in revised form 8 July 2025; Accepted 4 August 2025

**Table 1**

Specifications for the GPU nodes in the Leonardo and LUMI supercomputers [10,9,11,12]. The data formally corresponds to the Booster partition in Leonardo, and the LUMI-G nodes for LUMI. The abbreviation “Gb” refers to Gigabit, whereas “GB”/“TB” corresponds to Gigabyte/Terabyte.

	Leonardo	LUMI
GPU model	NVIDIA A100	AMD MI250X
Number of devices	4 GPUs	8 GCDs (2 per GPU)
Inter-node bandwidth	100 Gb/s	100 Gb/s <sup>†</sup>
Intra-node bandwidth	800 Gb/s	400 – 1600 Gb/s <sup>†</sup>
Internal memory bandwidth	1.6 TB/s	1.6 TB/s <sup>†</sup>
HBM	64 GB	64 GB <sup>†</sup>
FP64 Peak Performance	9.7 TFLOPs	23.95 TFLOPs <sup>†</sup>
	19.5 TFLOPs with Tensor Cores	

<sup>†</sup> Data for each GCD in LUMI.

nication, which may require adjustments in the numerical algorithm. Let us consider the current European pre-exascale flagship supercomputers Leonardo and LUMI. A summary of the GPU specifications for both supercomputers can be found in Table 1. In the case of LUMI, it is important to highlight that each MI250x is split into two Graphics Compute Dies (GCDs)<sup>1</sup> [9], each with roughly similar characteristics as a NVIDIA A100 GPU in terms of memory and processing power. The internal memory bandwidths for both AMD and NVIDIA GPUs are much faster than intra- and inter-node communication. Therefore, algorithms that minimize device-to-device communication can be optimal, even if they slightly increase the number of arithmetic operations or HBM (High Memory Bandwidth) usage per device. Additionally, opportunities for code optimization can be found by selecting algorithms that replace inter-node communication with faster intra-node data transfers.

Several works have performed large-scale turbulent flow simulations using multi-GPU configurations. Compressible flow solvers, for instance, contain many fully explicit calculations that can be readily parallelized using GPUs. The DNS solver *STREAMS* [13] can use multi-GPU systems to simulate compressible wall-bounded flows, taking into account complex effects such as shock-wave interactions. In *URANOS* [14], a compressible flow solver is developed for large-scale simulations using various modeling frameworks, and several possible choices for the discretization schemes. In incompressible flow solvers, GPU porting for distributed-memory calculations at scale faces an extra challenge. Typically, a major performance bottleneck is solving a large linear system associated with the pressure Poisson equation to ensure incompressibility. Nevertheless, several recent works have shown great progress in the development of multi-GPU solvers. Focusing on spectral or finite-difference approaches, an example is the AFiD-GPU code [15] for large-scale simulations of wall-bounded flows using multi-GPU (or multi-CPU) configurations. Another example is the CaNS code, which is used in the present study [16,17]. CaNS is an incompressible DNS solver, which is compatible with various types of boundary conditions for canonical flow cases in rectangular grids, such as isotropic turbulence or wall-bounded flows. This solver is compatible with both multi-GPU and multi-core CPU architectures, and has been recently re-reported to GPUs porting using OpenACC and the hardware-adaptive *cuDecomp* library for GPU communications at scale [18]. This library allows pencil-distributed solvers that require collective transpose operations to perform runtime autotuning and determine the optimal domain decomposition and GPU communication backend. A simple FFT-based finite-differences numerical solver like CaNS requires two types of communication operations: halo exchanges and transposes. Halo exchanges are relatively simple, standard operations, where each task exchanges boundary values with its neighbors. Transposes are more expensive *all-to-all* collective operations, where 3D data of a field is redistributed among MPI tasks such


that all cells aligned in a specific dimension are local to a single process/task. This is important while performing, for instance, fast Fourier-based transformations in spectral Poisson/Helmholtz solvers, since FFT algorithms require frequent access to the arrays being transformed. Naturally, all transpose operations involving 3D arrays are expensive, and often the major performance bottleneck.

Second-order FFT-based finite-difference solvers such as those used in AFiD and CaNS require performing FFTs along two directions, and the solution of a resulting tridiagonal system along the other direction, which is typically the wall-normal one in the case of wall-bounded flows with one inhomogeneous direction. The solution of the tridiagonal system has been typically performed using transpose operations, such that the whole system is local to each task and serially solved. However, there is possible room for improvement here by exploiting a parallel tridiagonal solver that avoids this collective operation. Notably, an interesting approach was presented by László et al. [19] and exploited in [20–23], which showed compelling performance gains at scale. In short, this method uses a hybrid Thomas–parallel cyclic reduction (PCR) algorithm that effectively converts the tridiagonal system to be solved in parallel into a series of smaller systems that can be solved independently, coupled to a smaller problem to be solved collectively for the first and last unknowns of each small system [22].

Most works exploring PCR in this context have adopted a 1D parallelization [21,23], with slabs parallel to directions of FFT-based synthesis. This is efficient and was proven to work up to a certain scale. However, as the flow Reynolds number increases, it becomes impossible to resort to a 1D parallelization. As an example, Fig. 1 presents the total memory requirement in a DNS solver (CaNS), to simulate turbulent channel flows at increasing friction Reynolds number ( $Re_\tau$ ), as well as the total size of a single  $n_x \times n_y$  wall-parallel slice. Expectedly, as  $Re_\tau$  increases, in addition to stricter time steps restrictions, the number of grid cells increases roughly as  $N_{x,y} \propto Re_\tau$  in the streamwise ( $x$ ) and spanwise ( $y$ ) directions of the channel flow, and  $N_z \propto Re_\tau^{3/4}$  in the wall-normal direction ( $z$ ) [24,25]. Hence, as the Reynolds number increases, the thickness of a wall-parallel slab that fits a fixed amount of grid points (e.g., dictated by the RAM constraints of a GPU or CPU device) becomes ever thinner, until it becomes impossible to decompose the domain further. This is particularly problematic in wall-bounded turbulence, where the number of grid points along the wall-parallel directions should be larger than in the wall-normal one [24]. Yet, the same is bound to happen in other turbulent flows (e.g., homogeneous isotropic turbulence) at sufficiently high Reynolds number.

Consequently, even with the ever-increasing memory capacity of GPUs, for DNS of high Reynolds number flows with this type of approach, one may be bound to adopt a 2D pencil-like domain decomposition. Leveraging a less communication-intensive approach for solving the Poisson equation, while retaining a pencil-distributed decomposition, is precisely the motivation of the present work. One alternative within this context is to replace serial TDMA (tridiagonal matrix algorithms) by parallel methods, with a lower communication footprint.

<sup>1</sup> In this work, we consider each GCD in LUMI to be an independent GPU (unless explicitly noted), since they are exposed to the user’s software as separate devices.



**Fig. 1.** Estimates of the total memory requirements for double-precision DNS runs of turbulent channel flows on a  $L_x \times L_y \times L_z = 12.8 \times 6.4 \times 2$  domain at different friction Reynolds numbers ( $Re_\tau$ ), as well as the size of a single slice  $n_x \times n_y$  of the DNS domain in HBM for all points in the streamwise (x) and spanwise (y) directions. The calculations are performed using the domain partition algorithms of the DNS solver CaNS [17] and the parallel decomposition library *cuDecomp* [18]. The symbols denote the limits where a 1D slab decomposition becomes impossible for GPUs with 64 GB and 128 GB of memory, respectively, which is typical of current high-end HPC GPUs, along with the total GPU memory of the largest supercomputer as of 2025: El Capitan [26]. This marks an upper bound of the maximum Reynolds number that could be investigated with current computational resources:  $Re_\tau \approx 50,000$ .

We develop such an approach based on a PCR-TDMA method (named P-TDMA hereafter), and test it on the CaNS solver with a focus on many-GPU calculations at scale. To run our solver on an AMD-based system like LUMI, an additional porting effort was undertaken, which we will also describe here. We report the revised solver's performance and scalability for large-scale DNS of turbulent channel flow at high Reynolds number. Moreover, we show that particular care should be taken for wall-normal implicit diffusion, to secure high performance at scale. Our results show an almost 2 $\times$  speedup at scale for the Poisson solver with 2D decomposition, which significantly improves the overall solver performance on large-scale GPU-based supercomputers.

This manuscript is organized as follows. Section 2 describes the governing equations, the discretization scheme, and the implementation of the Poisson/Helmholtz solver, including the cross-platform effort to run on AMD-based systems. Then Section 3 discusses the study and the scalability benchmarks. Finally, Section 4 presents the conclusions.

## 2. Methodology

### 2.1. Governing equations and numerical discretization

The current numerical framework solves the incompressible Navier-Stokes equations,

$$\nabla \cdot \mathbf{u} = 0, \quad (1)$$

$$\partial_t \mathbf{u} + (\mathbf{u} \cdot \nabla) \mathbf{u} = -\nabla p + \nu \nabla^2 \mathbf{u}, \quad (2)$$

where  $\mathbf{u}$  and  $p$  correspond to the fluid velocity vector and the pressure scaled by the fluid density;  $\nu$  is the fluid kinematic viscosity. The numerical scheme is based on an incremental pressure correction scheme or fractional-step method [27,28], where a prediction velocity  $\mathbf{u}^*$  is first calculated by integrating the momentum equation in time, and continuity is imposed using a correction pressure  $\Phi$ , which is obtained from the solution of a Poisson equation. The equations were solved on a rectangular box using a structured Cartesian grid with a staggered (MAC) arrangement for the velocity and pressure grid cells. As per the restriction of the Poisson solver, we employ uniform spacing along two Cartesian directions (x,y) and non-uniform spacing in the third spatial direction (z). Second-order finite differences are used

for spatial discretization. This has several advantages with respect to higher-order methods, such as being computationally efficient while still enabling simulations with similar fidelity as spectral discretization methods in practice [29], and being flexible and easily extended with numerical techniques for handling complex geometries like the immersed-boundary method [30–32], or multi-fluid flows [33]. Finally, Wray's low-storage Runge-Kutta scheme [34] is used for temporal discretization. The numerical scheme is presented below in semi-discrete form:

$$\mathbf{u}^* = \mathbf{u}^k + \Delta t (\alpha_k (\mathcal{A}\mathbf{u}^k + \nu \mathcal{L}\mathbf{u}^k) + \beta_k (\mathcal{A}\mathbf{u}^{k-1} + \nu \mathcal{L}\mathbf{u}^{k-1}) - \gamma_k \mathcal{G}p^{k-1/2}), \quad (3)$$

$$\mathcal{L}\Phi = \frac{D\mathbf{u}^*}{\gamma_k \Delta t}, \quad (4)$$

$$\mathbf{u}^{k+1} = \mathbf{u}^* - \gamma_k \Delta t \mathcal{G}\Phi, \quad (5)$$

$$p^{k+1/2} = p^{k-1/2} + \Phi, \quad (6)$$

where  $\alpha$ ,  $\beta$ , and  $\gamma$  refer to the coefficients of the RK3 scheme, which are given by:  $\alpha = \{8/15, 5/12, 3/4\}$ ,  $\beta = \{0, -17/60, -5/12\}$  and  $\gamma = \alpha + \beta$ ; the index  $k$  refers to the RK3 sub-iteration index  $k = \{0, 1, 2\}$ , and  $\Delta t$  is the time step. For flows at very low Reynolds numbers, or highly refined grids, the time step size  $\Delta t$  can be prohibitively small if diffusive terms are integrated in time explicitly. In such cases, it may be preferable to perform an implicit discretization of the diffusion terms at the cost of solving an extra Helmholtz equation per velocity component, as illustrated below:

$$\mathbf{u}^{**} = \mathbf{u}^k + \Delta t (\alpha_k \mathcal{A}\mathbf{u}^k + \beta_k \mathcal{A}\mathbf{u}^{k-1} + \gamma_k (-\mathcal{G}p^{k-1/2} + \nu \mathcal{L}\mathbf{u}^k)), \quad (7)$$

$$\mathbf{u}^* - \gamma_k \frac{\nu \Delta t}{2} \mathcal{L}\mathbf{u}^* = \mathbf{u}^{**} - \gamma_k \frac{\nu \Delta t}{2} \mathcal{L}\mathbf{u}^k, \quad (8)$$

$$\mathcal{L}\Phi = \frac{D\mathbf{u}^*}{\gamma_k \Delta t}, \quad (4)$$

$$\mathbf{u}^{k+1} = \mathbf{u}^* - \gamma_k \Delta t \mathcal{G}\Phi, \quad (9)$$

$$p^{k+1/2} = p^{k-1/2} + \Phi - \gamma_k \frac{\nu \Delta t}{2} \mathcal{L}\Phi. \quad (10)$$

Note that, indeed, Eq. (8) is a Helmholtz equation for the prediction velocity  $\mathbf{u}^*$ , which is implicit in the three spatial directions (x,y,z). While this equation can be solved efficiently using the fast direct methods presented in Section 2.2, the computational overhead is still considerable. Fortunately, in many cases, the time step constraints for  $\Delta t$  are due to fine grid spacing only along the non-uniform grid direction (here, z). This is particularly true for wall-bounded flows with one inhomogeneous direction, such as pipes or channels, which require fine grid spacing near the walls [15,16]. In these cases, one can discretize only the wall-normal diffusion term implicitly, and replace Eq. (8) with a one-dimensional system per velocity component:


$$\mathbf{u}^* - \gamma_k \frac{\nu \Delta t}{2} \mathcal{L}_z \mathbf{u}^* = \mathbf{u}^{**} - \gamma_k \frac{\nu \Delta t}{2} \mathcal{L}_z \mathbf{u}^k, \quad (11)$$

where  $\mathcal{L}_z$  denotes the discrete Laplacian term associated with the z direction. This is numerically much cheaper, as the second-order finite-difference discretization of this equation requires the solution of a simple tridiagonal system.

### 2.2. Numerical solution of the Poisson/Helmholtz equation

#### 2.2.1. Fourier-based synthesis

The solution of the Poisson equation for the pressure comprises some of the numerical algorithm's most computation and communication-intensive steps. Here, Eqs. (4) and (8) are solved using the method of eigenfunctions expansions, which allows for fast, direct solutions by leveraging the FFT algorithm [35,16]. After performing a Fourier-based synthesis of the Poisson/Helmholtz equation along directions x and y, the following system of tridiagonal equations can be obtained along the non-uniform grid direction z for a grid cell with index  $i, j, k$ :



**Fig. 2.** Schematic representation of a FFT-based linear solver using a 2D pencil decomposition for the DNS domain. The black arrows indicate the global transpose operations for the data stored among different MPI tasks. The color of each 2D pencil represents a different MPI task. Forward (*fwd.*) operations are first performed from left to right following the direction of the black arrows. Then, all transpose operations are reversed, and inverse/backward (*bwd.*) Fourier-based transforms are performed, as indicated by the gray arrows. Note that the first transpose operation is often implemented as two consecutive transposes:  $z \rightarrow y$  and  $y \rightarrow x$ . (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

$$(\lambda_i/\Delta x^2 + \lambda_j/\Delta y^2) \tilde{\Phi}_{i,j,k} + (\eta_{k-1} \tilde{\Phi}_{i,j,k-1} + \eta_k \tilde{\Phi}_{i,j,k} + \eta_{k+1} \tilde{\Phi}_{i,j,k+1}) = \tilde{f}_{i,j,k}, \quad (12)$$

where the tilde ( $\tilde{\cdot}$ ) denotes two successive discrete Fourier-based (i.e., Fourier/cosine/sine) transforms applied to a variable along the  $x$  (index  $i$ ) and  $y$  (index  $j$ ) directions. Note that each  $(i, j)$  pair corresponds to a tridiagonal system along the non-uniform direction ( $z$ , index  $k$ ). The coefficients  $\lambda_i$  and  $\lambda_j$  are the second-order accurate eigenvalues (or modified wavenumbers); see, e.g., [35];  $\Delta x$  and  $\Delta y$  correspond to the uniform grid spacing in the  $x$  and  $y$  directions, whereas the set of coefficients  $\eta$  represent the finite-difference discretization of the  $\mathcal{L}_z$  operator along  $z$ . While the method of eigenfunctions expansions aims at exploiting the FFT algorithm, it still allows for multiple combinations of boundary conditions representative of different classes of canonical turbulent flows, from isotropic turbulence to several boundary-free and wall-bounded shear flows.

After obtaining  $\tilde{\Phi}_{i,j,k}$  from Eq. (12), the final solution ( $\Phi_{i,j,k}$ ) is easily computed from the inverse Fourier-based synthesis. The numerical methods and algorithms to solve Eq. (12) are the key parts of the present work and will be discussed next in Section 2.2.2. Finally, there are noteworthy nuances in implementing fast real-to-complex/complex-to-real (Fourier) and real-to-real (sine/cosine) transforms on GPUs in a unified framework, which we describe in Appendix A.

### 2.2.2. Original distributed-memory solution

The numerical solution of the second-order finite-difference Poisson/Helmholtz equations in rectangular grids is solved using FFT-based methods in a distributed-memory setting. The original approach used to solve this problem is described below,<sup>2</sup> and illustrated in Fig. 2. The following steps are taken:

1. Compute right-hand side term  $d_{i,j,k}$  of eq. (12) in  $z$ -aligned pencils.
2. Transpose data to  $x$ -aligned pencils and perform  $N_y N_z$  Fourier-based transforms along  $x$ .
3. Transpose data to  $y$ -aligned pencils and perform  $N_x N_z$  Fourier-based transforms along  $y$ .
4. Transpose data to  $z$ -aligned pencils and solve the resulting  $N_x N_y$  tridiagonal systems of equations along  $z$ .
5. Perform the reciprocal transpose operation as in step 4.

<sup>2</sup> It is important to note that, while the code CaNS allows for an arbitrary default pencil orientation (i.e., outside the pressure solver), we start from  $Z$ -aligned pencils since this minimizes the number of collective communications when solving the momentum equation with  $z$ -implicit diffusion. Starting from  $x$ -aligned pencils would avoid transpose operations in the Poisson solver, but many additional transpose operations would be required for inverting a tridiagonal system per velocity component.

6. Perform the reciprocal inverse transforms and transpose operation as in step 3.
7. Perform the reciprocal inverse transforms and transpose operation as in step 2, to obtain the final solution in  $z$ -aligned pencils.

Here,  $N_{x/y/z}$  are the local number of grid cells along  $x/y/z$  during the different steps of the algorithm for each MPI task. The transpose operations are an *all-to-all* collective, which may be very expensive. Within this approach, the FFT-based transforms and solution of the tridiagonal systems can be trivially mapped to different parallel (GPU) threads. Note that, whenever a 1D slab-like parallelization is possible, some of the transpose operations shown above would turn into a no-op, making it often desirable. In this regard, the best-performing slab configurations are those partitioned along  $y$ . This is convenient, since each GPU can perform Fourier transformations along the  $x$ -direction, and solve tridiagonal systems of equations along the  $z$ -direction, without performing additional collective operations. Even with 1D implicit diffusion, only one pair of transposes is required per step: the  $x \leftrightarrow y$  transposes shown in Fig. 2.


While the approach presented in Fig. 2 enables GPU-accelerated DNS of fluid flows on many CPUs/GPUs, the transposing operations may result in a major performance loss. This may be particularly problematic in modern GPU-based systems at scale, as the inter-node bandwidth is orders of magnitude slower than the GPU memory bandwidth or the intra-node communication. Hence, we need an approach that: (1) keeps a 2D parallelization, which is unavoidable at scale, and (2) reduces the amount of data used in collective communications to a reasonable minimum. We will explain this approach below.

### 2.2.3. Solution with parallel tridiagonal solver

In this approach, we split the computational domain along the  $z$  direction, and exploit a parallel tridiagonal solver as in [19,21,36] to circumvent the large *all-to-all* operations in the previous section. This approach is shown in Fig. 3. This algorithm starts with a cyclic reduction step, such that only information regarding the boundaries of every slice in the  $z$ -direction must be communicated to other MPI tasks. Subsequently, a tri-diagonal system for the boundary values is constructed. While an *all-to-all* type of collective is still needed, internal data points are not communicated, which can drastically reduce the communication overhead. The steps of this algorithm are summarized as follows:

1. Compute right-hand side term  $d_{i,j,k}$  of eq. (12) in  $x$ -aligned pencils.
2. Perform  $N_y N_z$  Fourier-based transforms along  $x$  (see Fig. 3).
3. Transpose data to  $y$ -aligned pencils and perform  $N_x N_z$  Fourier-based transforms along  $y$ .
4. Perform  $N_x N_y$  cyclic reductions along  $z$ , pack  $2N_x N_y$  boundary values, and transpose the packed boundaries to  $z$ -aligned pencils.





**Fig. 3.** Modified parallel tridiagonal solver for large-scale DNS. The first two sub-images follow the same conventions as Fig. 2, where the color of each partition indicates a different MPI task. In the third and fourth sub-images, the process of cyclic reduction is highlighted, by applying a different color to the boundary values for each partition.

5. Solve  $N_x N_y$  reduced systems of tridiagonal equations along  $z$ , with local size  $2p_z$ .
6. Transpose data to  $y$ -aligned pencils, unpack  $2N_x N_y$  boundary values, and reconstruct the internal solution fields.
7. Perform the reciprocal inverse transforms and transpose operation as in step 3.
8. Perform the reciprocal inverse transforms as in step 2.

In the previous steps,  $N_{x/y/z}$  is again the local grid size for each MPI rank in each Cartesian direction, whereas  $p_z$  is the number of divisions of the computational domain along the  $z$ -direction, which would correspond to  $p_z = 4$  in Fig. 3. The details of the parallel tridiagonal algorithm, and the modifications proposed in this work for the computation of its internal coefficients, are explained in Sections 2.2.4. Clearly, the amount of data transferred among MPI tasks is substantially reduced [21]. The tridiagonal system of equations found in the right-side of Fig. 3 have a size of  $2p_z$ , where  $p_z$  is the number of partitions of the computational domain along the  $z$ -direction. Therefore, the parallel tridiagonal solver should be efficient as long as  $2p_z \ll N_z$ , where  $N_z$  is the total number of grid points in the  $z$ -direction.

Interestingly, in the parallel tridiagonal solver, increasing the number of lateral divisions ( $p_y$ ) favors strong scalability: When  $p_y$  is increased, the size of the boundaries (per MPI task) is reduced as  $2n_x n_y / p_y$ . Therefore, doubling  $p_y$  halves the data communicated per task, leading to excellent scalability. In contrast, increasing the number of vertical partitions ( $p_z$ ) does not reduce the MPI workload, and thus it is unfavorable for scalability. This is particularly relevant for 1D slab configurations where  $p_y = 1$  and  $p_z$  is the total of GPUs. Still, 1D slab configurations are optimal when  $p_z \ll N_z$ .

#### 2.2.4. Parallel tridiagonal algorithm

Numerous approaches may be considered to parallelize the Thomas algorithm to solve a tridiagonal system (see, e.g., the survey in [21]). Here we adopt the method proposed by [19], which uses cyclic reduction, combined with a Thomas algorithm for a reduced system. We have illustrated the approach in Fig. 3, and summarize it below.

First, the distributed tridiagonal systems of this form (cf. eq. (12))

$$a_k \phi_{k-1} + b_k \phi_k + c_k \phi_{k+1} = d_k, \quad (13)$$

are locally reduced to a problem where inner unknowns within the computational subdomain are only a function of the values at its top and bottom boundary:

$$a'_k \phi_0 + \phi_k + c'_k \phi_{m-1} = d'_k, \quad (14)$$

using a cyclic reduction step. The original set of coefficients and right-hand-side ( $a, b, c, d$ ) are then reduced to the ( $a', c', d'$ ), where the main diagonal is normalized to have unit weight. Algorithm 1 describes this approach for completeness, and more details can be found in [19,21,36].

Second, the sets of values ( $a', c', d'$ ) at the boundaries of every domain  $k = \{0, m-1\}$  can be grouped and transposed in a collective operation (recall Fig. 3). Then, the standard Thomas algorithm solves the reduced systems of tridiagonal equations for the boundary values of all subdomains. Finally, the boundary data for  $\phi_{i,j,k}$  can be globally transposed, and the values of  $\phi_{i,j,k}$  in the interior of every sub-domain can be reconstructed using eq. (14).

**Algorithm 1** Cyclic reduction step of the parallel tridiagonal algorithm [19,21], where  $\mathbf{d}$  corresponds to the right-hand-side of the system (see Eq. (12)).

---

```

1: Step 1: Initialization
2: Input:  $\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{d}$ 
3:  $\mathbf{a}_0 \leftarrow \mathbf{a}_0 / \mathbf{b}_0$ ;  $\mathbf{c}_0 \leftarrow \mathbf{c}_0 / \mathbf{b}_0$ ;  $\mathbf{d}_0 \leftarrow \mathbf{d}_0 / \mathbf{b}_0$ 
4:  $\mathbf{a}_1 \leftarrow \mathbf{a}_1 / \mathbf{b}_1$ ;  $\mathbf{c}_1 \leftarrow \mathbf{c}_1 / \mathbf{b}_1$ ;  $\mathbf{d}_1 \leftarrow \mathbf{d}_1 / \mathbf{b}_1$ 
5: for  $i = 2, \dots, m-1$  do
6:    $r \leftarrow 1 / (\mathbf{b}_i - \mathbf{a}_i \mathbf{c}_{i-1})$ 
7:    $\mathbf{d}_i \leftarrow r(\mathbf{d}_i - \mathbf{a}_i \mathbf{d}_{i-1})$ 
8:    $\mathbf{c}_i \leftarrow r \mathbf{c}_i$ 
9:    $\mathbf{a}_i \leftarrow -r \mathbf{a}_i \mathbf{a}_{i-1}$ 
10: end for
11: for  $i = m-3, \dots, 1$  do
12:    $\mathbf{d}_i \leftarrow \mathbf{d}_i - \mathbf{c}_i \mathbf{d}_{i+1}$ 
13:    $\mathbf{a}_i \leftarrow \mathbf{a}_i - \mathbf{c}_i \mathbf{a}_{i+1}$ 
14:    $\mathbf{c}_i \leftarrow -\mathbf{c}_i \mathbf{c}_{i+1}$ 
15: end for
16:  $r \leftarrow 1 / (1 - \mathbf{a}_1 \mathbf{c}_0)$ 
17:  $\mathbf{d}_0 \leftarrow r(\mathbf{d}_0 - \mathbf{c}_0 \mathbf{d}_1)$ 
18:  $\mathbf{a}_0 \leftarrow r \mathbf{a}_0$ 
19:  $\mathbf{c}_0 \leftarrow -r \mathbf{c}_0 \mathbf{c}_1$ 
20:  $\mathbf{b} = 1$ 
21:
22: Step 2: Solve reduced system of equations for boundary values
23:
24: Step 3: Reconstruct the solution in-place
25: Input:  $\mathbf{a}, \mathbf{c}, \mathbf{d}, \mathbf{x}_0, \mathbf{x}_{m-1}$ 
26:  $\mathbf{d}_0 \leftarrow \mathbf{x}_0$ 
27:  $\mathbf{d}_{m-1} \leftarrow \mathbf{x}_{m-1}$ 
28: for  $i = 1, \dots, m-2$  do
29:    $\mathbf{d}_i \leftarrow \mathbf{d}_i - \mathbf{a}_i \mathbf{x}_0 - \mathbf{c}_i \mathbf{x}_{m-1}$ 
30: end for

```

---

A few important notes should be considered to secure parallel performance at scale when combining the pressure Poisson equation with  $z$ -implicit diffusion. While the same computational approach may be taken for solving both cases, the straightforward implementation of Algorithm 1 would be far from optimal in both cases. Note that: (1) the reduced tridiagonal system for the Poisson equation is time-invariant (eq. (4)), with a problem that changes for each  $(i, j)$  index, yet (2) the  $z$ -implicit matrix is constant for each  $(i, j)$  index, but time-dependent (eq. (11)). Hence, a key optimization for the Poisson equation is to

perform the transpose operations associated with the reduced system coefficients ( $a'$ ,  $b'$ ,  $c'$ ) only once as an initialization step (recall the penultimate step in Fig. 2), as only  $d'$  varies with the right-hand-side of eq. (4). Regarding implicit  $z$  diffusion, ( $a'$ ,  $b'$ ,  $c'$ ) are time-dependent, but identical for all  $(i, j)$  indexes mapped to different GPU threads to solve the equations along  $z$ . Thus, rather than communicating ( $a'$ ,  $b'$ ,  $c'$ ) through MPI operations, it is much faster to have each GPU computing the global ( $a'$ ,  $b'$ ,  $c'$ ) coefficients corresponding to all MPI tasks aligned in the  $z$ -direction, and locally copy the values pertaining to its own subdomain. This is done on the GPUs with unnoticeable computational overhead, and effectively avoids expensive MPI communication operations.

Accordingly, to efficiently handle implicit  $z$  diffusion, Algorithm 1 was re-derived in a flavor that splits the solution into an initialization with pre-computed coefficients and a runtime step. This approach is presented in Algorithm 2. A major advantage of this approach is that only the array  $d'$  is modified in-place at runtime. This marks a large contrast to Algorithm 1, which requires thread-private arrays (or memory buffers) to track intermediate changes in the arrays ( $a'$ ,  $c'$ ). This change is particularly relevant when solving for implicit 1D diffusion, since only 1D arrays with precomputed ( $a'$ ,  $b'$ ,  $c'$ ) coefficients can handle the solution process.

From a mathematical perspective, the new algorithm is derived by analyzing the cyclic reduction process, and carefully tracking which references to the ( $a'$ ,  $b'$ ,  $c'$ ) arrays can be replaced by either their input or output values. After making this distinction, it becomes evident that the reduction process for the array  $d'$  does not depend on intermediate values for ( $a'$ ,  $b'$ ,  $c'$ ) being over-written. Therefore, it is natural to split the process into initialization and runtime stages. Moreover, it is important to highlight that only the initial values of ( $a'$ ,  $c'$ ) are used. In the DNS solver, the variables ( $a'$ ,  $c'$ ) always correspond to 1D vectors, even for Poisson or Helmholtz solvers. As a result, storing the initial values of ( $a$ ,  $c$ ) creates a negligible performance overhead.

### 2.3. Implementation

As previously noted, the approaches we have presented are implemented in the open-source code CaNS [16,17]. CaNS is written in Modern Fortran, and since its version 2.0 offloads data and computation to GPUs using OpenACC. At the fine-grained parallelization level of the TDMA implementations, each  $i, j$  index is assigned to a thread, which serially performs operations along the third domain direction. The CPU implementation uses the *2DECOMP&FFT* library [37] to perform transposes. Conversely, the multi-GPU implementation uses the *cuDecomp* library for pencil-distributed calculations at scale that feature transposes and halo exchanges [18] is used in CaNS. The main advantage of *cuDecomp* is its runtime autotuning capabilities, which allows to confidently select a well-performing combination of 2D processor grid and communication backend (with several low-level implementations of the transposing algorithm in CUDA-aware MPI, NCCL, or NVSHMEM). *cuDecomp*'s flexibility allowed for a very straightforward implementation of the communication operations that can be visualized in Fig. 2. Indeed, the *all-to-all* operations needed to communicate the boundary values for each sub-group of slices along the  $z$ -direction could be replaced by the existing transpose operations available in the *cuDecomp* or *2DECOMP* libraries.

The distributed-memory implementation of CaNS, using *cuDecomp* and *cuFFT*, allowed for very efficient calculations at scale on NVIDIA-based systems. However, in the present work we decided to benchmark our approach on the supercomputers Leonardo (NVIDIA-based) and LUMI (AMD-based). We summarize our implementation approach for the latter, which we plan to incorporate in the CaNS public repository in the near future.

Regarding the verification of the implementation, CaNS has been extensively validated in the past [16,17]. Therefore, the correctness of the current implementation can be trivially verified, since only the parallel tridiagonal solvers have been touched. An explicit comparison

**Algorithm 2** Alternative flavor of the cyclic reduction method proposed in this work, with a separation between initialization and runtime operations. The variables (**A, B, C**) correspond to the original coefficients of the tridiagonal equations, whereas (**a, b, c**) are the modified coefficients after cyclic reduction.

---

```

1: Step 1: Initialization
2: Input: A, B, C, a, b, c
3:  $a_1 \leftarrow A_1$ 
4:  $b_1 \leftarrow B_1$ 
5: for  $i = 2, \dots, m-1$  do
6:    $b_i \leftarrow B_i - A_i C_{i-1}/b_{i-1}$ 
7:    $a_i \leftarrow -A_i a_{i-1}/b_{i-1}$ 
8: end for
9:  $c_{m-1} \leftarrow C_{m-1}$ 
10:  $c_{m-2} \leftarrow C_{m-2}$ 
11: for  $i = m-3, \dots, 1$  do
12:    $a_i \leftarrow a_i - C_i a_{i+1}/b_{i+1}$ 
13:    $c_i \leftarrow -C_i c_{i+1}/b_{i+1}$ 
14: end for
15:  $a_0 \leftarrow A_0$ 
16:  $b_0 \leftarrow B_0 - C_0 a_1/b_1$ 
17:  $c_0 \leftarrow -C_0 c_1/b_1$ 
   ▷ Note: To avoid GPU divisions in Steps 2 and 4, optimized implementations
   can store  $1/b$  instead of  $b$ .
18:
19: Step 2: Runtime reduction
20: Input: A, B, C, a, b, c, d
21: for  $i = 2, \dots, m-1$  do
22:    $d_i \leftarrow d_i - A_i d_{i-1}/b_{i-1}$ 
23: end for
24: for  $i = m-3, \dots, 0$  do
25:    $d_i \leftarrow d_i - C_i d_{i+1}/b_{i+1}$ 
26: end for
27:
28: Step 3: Solve reduced system of equations for boundary values
29:
30: Step 4: Reconstruct the solution in-place
31: Input: a, b, c, d,  $x_0, x_{m-1}$ 
32:  $d_0 \leftarrow x_0$ 
33:  $d_{m-1} \leftarrow x_{m-1}$ 
34: for  $i = 1, \dots, m-2$  do
35:    $d_i \leftarrow (d_i - a_i x_0 - c_i x_{m-1})/b_i$ 
36: end for

```


---

with respect to the output of the full-transpose method is accurate up to machine precision for the modified subroutines, which verifies the correctness of the implementation.

#### 2.3.1. Many-GPU implementation on LUMI

On LUMI, the Cray Fortran compiler is readily compatible with OpenACC, and GPU-aware MPI is supported to perform data transfer among GPU devices (or GCDs). However, since LUMI has AMD cards, some work was required to port the transpose and halo exchange algorithms for LUMI. One approach would be to adjust *cuDecomp* such that the NVIDIA-specific features of the library are masked out of the build workflow. In the present work, we took a different route and developed a cross-platform communication library based on Fortran and OpenACC named *diezDecomp*. Its source code is available on GitHub under an MIT license [38]. This implementation uses OpenACC and GPU-aware MPI to perform transpose and halo exchange operations. During transpose operations, *diezDecomp* has been optimized to pack and unpack data related to different MPI ranks simultaneously (in parallel GPU threads), and it supports conversions between different indexing orders (e.g.,  $x/y/z$  to  $y/x/z$ ). As further verification, the performance of the *diezDecomp* library was tested in the Leonardo supercomputer, achieving nearly identical running times as the *cuDecomp* library.

Finally, it is important to highlight that the current distributed Poisson/Helmholtz solver is able to work with various types of FFT libraries, such as *cuFFT* [39], *hipFFT* [40], or even in-house FFT implementations.



**Fig. 4.** Strong scalability chart for a wall-bounded flow with grid size  $(N_x \times N_y \times N_z) = (7168 \times 7168 \times 1594)$ , which roughly corresponds to a friction Reynolds number of  $Re_\tau \approx 5000$ . The variable  $p_z$  corresponds to the number of divisions along the  $z$  direction for the pencil decomposition scheme. The abbreviation “FTM” refers to the original DNS solver, which was based on the full-transpose method. All simulations were performed in the Leonardo and LUMI supercomputers. Please note that P-TDMA method with 1D slabs (filled black triangles) cannot be run with 1024 GPUs/GCDs, due to an insufficient number of grid points in the  $z$ -direction:  $N_z/p_z < 2$ .

In this context, our porting effort uses the *hipFFT* library for simulations in AMD GPUs using the bindings provided by the *hipFort* project [41], whereas *cuFFT* is enabled for NVIDIA GPUs via the CUDA toolkit [6]. Since both FFT libraries have similar APIs, our implementation uses CPP macros to switch between libraries depending on the target platform.

### 3. Results

#### 3.1. Strong and weak scalability

We consider a large-scale turbulent plane channel flow setup, which exercises all important steps presented in this paper, including  $z$ -implicit diffusion. Fig. 4 presents a strong scalability test for a  $N_x \times N_y \times N_z = 7168 \times 7168 \times 1594$  grid, containing approximately 80 billion grids points. This corresponds to a channel with a domain size of  $L_x \times L_y \times L_z = 12.8 \times 6.4 \times 2$ , with a friction Reynolds number  $Re_\tau \approx 5000$  [42]. Note that, in many GPU-resident workloads, a strong scaling test likely shows performance deterioration, as the occupancy of each GPU is being lowered [18]. The tests were performed on the Leonardo and LUMI supercomputers, which again, feature NVIDIA A100 and AMD MI250X GPU cards, respectively. In the figure, we compare the full transpose method (FTM) to the approach with the parallelized tridiagonal matrix algorithm (P-TDMA), for both 1D (slab) decompositions (as in Fig. 2 without decomposing along  $x$ , and in Fig. 3 without decomposing along  $y$ ) and for 2D pencil decomposition schemes with different levels of process decomposition along  $z$ :  $p_z = 2$  and 128. The number of partitions along the  $y$ -direction ( $p_y$ ) are set from the total number of GPUs:  $n_g = p_y \times p_z$ . These configurations are intended to show how the solver scales in the two limits where the 2D pencil decomposition has either the least number of partitions along  $z$ , or very high values. Clearly, the P-TDMA algorithm is far more efficient than FTM when 2D pencil decompositions are considered. This is explained by Figs. 2 and 3, because the P-TDMA algorithm transposes a significantly smaller amount of data than the FTM method. Focusing on the performance of the Poisson solver (bottom panels of Fig. 4), there is a consistent 1.5× improvement in wall-clock time per step. While strong

scaling deterioration is expected, it is interesting to note that some of the curves show reasonably mild deviations from the ideal scaling curve.

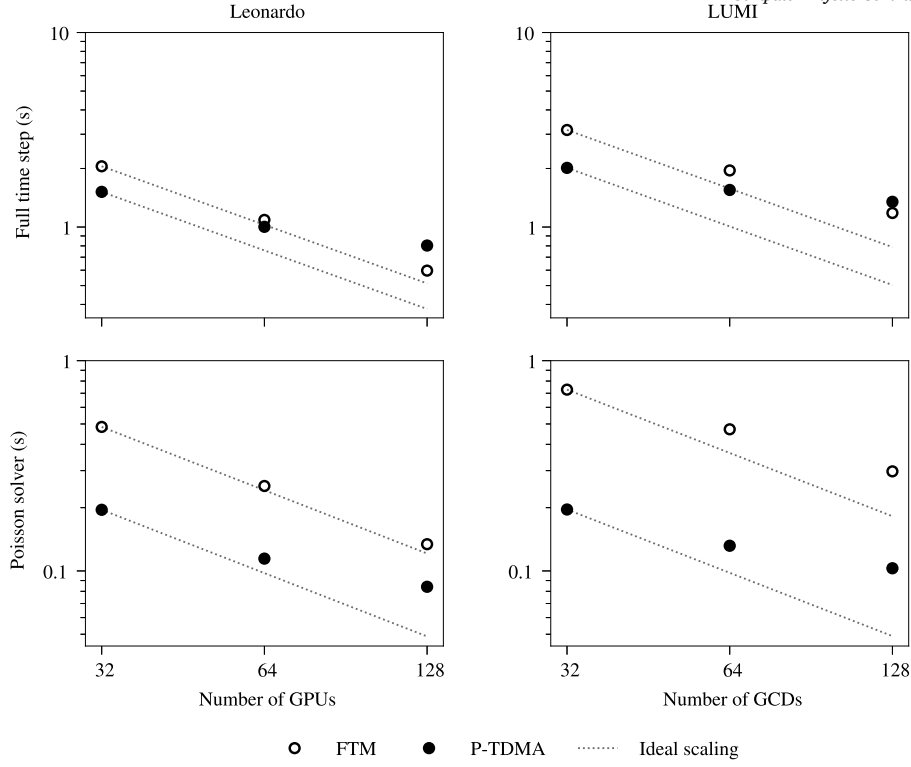
The P-TDMA method shows great efficiency for  $p_z = 2$ , and almost matches the best-performing 1D slab configurations in the FTM case. This is expected, as the P-TDMA method only transfers data along one small boundary in the  $z$ -direction when  $p_z = 2$ , with a computational cost roughly similar to a halo exchange.

Let us now consider the difference between the wall-clock time of the Poisson solver and the full time steps. The most important difference here is the solution of three separate systems of 1D implicit diffusion equations in the wall-clock time of the full time step. Therefore, configurations that transfer more data when solving for 1D implicit diffusion tend to have worse performance, such as the FTM approach with a 2D pencil decomposition.

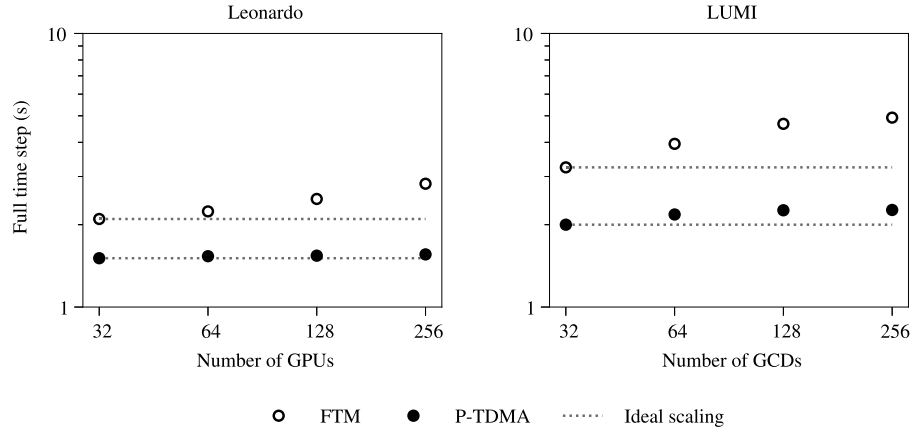
The wall-time per steps for the Poisson solver with 1024 GPUs/GCDs converge to a similar value for both the best FTM configurations and the P-TDMA cases with 2D pencil decompositions, since their performance is dominated by the transpose operations in the  $x$ - $y$  directions. Interestingly, one can notice a major performance gain for the P-TDMA when  $p_z = 128$  on LUMI, which corresponds to a special case where 50% of the required data after the  $x$ - $y$  transpose is already present in the GPU, requiring only a local copy.

In the case of the P-TDMA algorithm with 1D slabs, it can be observed in Fig. 4 that the running times remain nearly constant when the number of GPUs/GCDs increases from 256 to 512. This is due to the P-TDMA method having boundaries with a fixed size of  $(N_x \times N_z \times 2)$ , and thus, the data transferred per GPU/GCD remains equal as the number of processes is increased. Despite this, the P-TDMA method with 1D slabs still outperforms the FTM approach, when the number of grid points along the  $z$ -direction is very large for every 1D slab. This is illustrated in Fig. 5, where a strong scalability benchmark is presented for 1D slabs for  $Re_\tau \approx 2500$  and a grid size of  $N_x \times N_y \times N_z = 3200 \times 3456 \times 900$ . In this benchmark, the P-TDMA method outperforms the FTM approach with 32 GPUs/GCDs, yet its performance decreases as the number of processes grows. For cases at lower Reynolds numbers, the performance improvements can be expected to grow in favor of the P-TDMA method,





**Fig. 5.** Strong scalability chart for a wall-bounded flow with grid size  $(N_x \times N_y \times N_z) = (3200 \times 3456 \times 900)$ , which roughly corresponds to a friction Reynolds number of  $Re_\tau \approx 2,500$ . All runs correspond to 1D slab configurations, as in Fig. 2 without decomposing along  $x$ , and in Fig. 3 without decomposing along  $y$ .



**Fig. 6.** Weak scalability analysis for DNS cases with a fixed grid size  $(N_x \times N_y \times N_z) = (3456 \times 3072 \times 30)$  per GPU device. For the P-TDMA method, the pencil decomposition scheme only considers one partition along the  $y$ -direction ( $p_y = 1$ ), and a number of  $z$ -partitions ( $p_z$ ) equal to the number of GPUs. The FTM approach considers the 1D slab configuration as in Fig. 2 without decomposing along  $x$ .


since the number of grid points per GPU along  $z$  increases for the 1D slabs.

The results from the weak scaling tests are shown in Fig. 6. Weak scaling is the most important scaling indicator in large-scale GPU-resident DNS, as optimal resource usage requires maximizing GPU occupancy, which is kept fixed in these tests. We fixed the local domain sizes to about 318.5 million grid points per GPU, by considering a grid with  $3456 \times 3072 \times 30$  points per GPU, which saturates the GPUs (GCDs) on Leonardo (LUMI). For the P-TDMA method, the number of partitions  $p_z$  along the  $z$ -direction is equal to the number of GPUs/GCDs, since this corresponds to the most challenging scenario for the weak scaling analysis; for FTM, we use the optimal configuration for 1D slabs. Clearly, the performance of the P-TDMA approach is superior, not only being approximately 2× faster in wall-clock time per step, but also in terms of weak scaling. An 8-fold increase in the GPUs/GCDs used results in

a 3% (13%) performance degradation on Leonardo (LUMI), while the full transpose method shows a major weak scaling loss of 35% (52%) on Leonardo (LUMI). This shows that, indeed, the proposed improvements are important for efficient wall turbulence simulations at scale.

### 3.2. Breakdown of parallel performance

Let us now understand in more detail the results presented in the previous sections. Fig. 7, compares the workload distribution of the current approach and the previous algorithm based on the full-transpose method. The DNS cases chosen for comparison are simulations with 1024 GPUs/GCDs, shown in the strong scalability chart (Fig. 4). Consistently with the previous observations, the performance differences in the supercomputers we have tested show similar trends overall. The 2D pencil decomposition considered for comparison is  $(p_y \times p_z) = (512, 2)$ ,



**Fig. 7.** Comparison between the GPU timer profiles for the DNS solver with the P-TDMA and FTM methods on Leonardo and LUMI. For each profiled operation, the left bars correspond to the average wall-time per step (*mean*), while the right bars are the aggregated results considering the best elapsed times per operation (*min*). The data was extracted from the DNS runs with 1024 GPUs (Leonardo) or GCDs (LUMI), which are presented in the strong scalability chart (Fig. 4) for a fixed grid size  $(N_x \times N_y \times N_z) = (7168 \times 7168 \times 1594)$ . For the DNS solver with P-TDMA, the data further corresponds to the case with  $(p_y \times p_z) = (512 \times 2)$  partitions along the  $y$  and  $z$  directions, which has the fastest running times in the strong scalability chart. Both GPU timer profiles correspond to the averaged results for a single RK3 sub-step within the DNS solvers. Please note that the  $y \leftrightarrow z$  transposes for the P-TDMA algorithm transfer much less data, and thus their impact on the plotted budgets is minor.

which has the lowest running times overall for the P-TDMA algorithm. Not surprisingly, the P-TDMA approach is much faster than the full-transpose algorithm while solving for the pressure-Poisson equation. However, the P-TDMA algorithm is slower when solving for 1D implicit diffusion alongside each velocity component  $(u, v, w)$ . This is also expected, as again, the FTM uses an initial  $z$ -aligned decomposition, where the full tridiagonal problems are local to each MPI task. Both algorithms show similar overhead associated with halo exchanges, which is also expected.

### 3.2.1. Average timings

When inspecting the contributions to the mean total wall-time per step in Fig. 7, it becomes clear that *all-to-all* operations are the main performance bottleneck of the full-transpose method, whereas the P-TDMA approach is far less communication-intensive. Interestingly, the Poisson solver with the P-TDMA algorithm spends 81.7% (89.6%) of the time on Leonardo (LUMI) performing local transposes along the  $x$ - $y$  plane, which are required between the Fourier-based transforms described in Section 2.2. The *all-to-all* operations associated with the parallel tridiagonal solver only require 0.2% (0.6%) of the computing time in Leonardo (LUMI), as this collective operation only communicates boundary values for each MPI task (recall the penultimate step in Fig. 2). Note that, while the time spent performing  $x \rightarrow y$  transposes can be (partially) reduced by minimizing the  $p_y$  divisions along the  $y$ -direction (e.g.,  $p_y = 2$ ), the overall performance degrades in Fig. 4 when the number of vertical partitions  $p_z$  is increased for this case. This trend is valid for both the FTM approach and the P-TDMA methods, since the cost of performing  $y$ - $z$  transposes increases as  $p_z$  grows. Therefore, a careful trade-off

must be considered. Moreover, the optimal DNS domain decomposition is not only problem-dependent but also hardware-dependent, which makes runtime tuning of the computational setup very relevant [18].

Finally, in Fig. 7, it can be noticed that the P-TDMA approach also performs  $x \rightarrow y$  transposes when solving for implicit 1D diffusion, with a small but noticeable communication footprint. This is not strictly necessary, since cyclic reduction can be directly performed using the initial  $x$ -aligned pencils (Fig. 3, left), and then a direct  $x \rightarrow z$  transpose could be used to obtain  $z$ -aligned pencils [15]. This direct transpose is less trivial to implement and is not featured in *cuDecomp* or *2DECOMP&FFT*. Yet, since the *diezDecomp* communication library supports  $x \rightarrow z$  transposes for any desired 2D decomposition, we tested the performance gains of this direct transpose. The results of this additional benchmark are shown in Appendix B for LUMI. While the computational overhead of the additional  $x \rightarrow y$  transpose is small for the cases shown in Fig. 7 with the decomposition  $(p_y \times p_z) = (512 \times 2)$ , we identified that other pencil decompositions suffered from higher performance losses. For instance, the benchmark shows that the DNS case with  $(p_y \times p_z) = (8 \times 128)$  is about 55% faster during the calculation of  $z$ -implicit diffusion when using direct  $x \rightarrow z$  transposes. This results in savings of about 18% in the total wall-clock time per step.

### 3.2.2. Best elapsed times per operation

To better explain the differences between systems, we included in Fig. 7 the aggregated best elapsed times per operation (right bars). Remarkably, the total running times for the FTM (P-TDMA) method are decreased by about 28% (37%) for LUMI, and by 8% (9%) for Leonardo.

It appears that the average wall-clock time per step on LUMI is significantly suboptimal compared to that on Leonardo.

Significant performance differences between these machines for *all-to-all* collective operations have been reported in De Sensi et al. [11], where Leonardo outperformed LUMI by almost 50% in goodput benchmarks at scale. We argue that the analysis considering the best elapsed times per operation resembles a goodput benchmark better, since these results contain the smallest levels of network delays possible. Interestingly, when inspecting the aggregated best elapsed times (*min*) in Fig. 7, it can be noticed that LUMI is only 31% (45%) slower than Leonardo for the total running times of the FTM (P-TDMA) method, consistently with the trends in De Sensi et al. [11]. We verified that this trend is robust across other cases (not shown).

#### 4. Conclusions

We have presented a numerical approach for GPU-based massively-parallel DNS of turbulent wall flows with one inhomogeneous direction. Using the CaNS solver as base, we extended it with a parallel tridiagonal algorithm for solving the pressure Poisson equation, and to handle implicit integration of the wall-normal diffusion term. To achieve this, we adopted a recently-proposed approach for solving distributed tridiagonal systems [19,21], and implemented it in a pencil-decomposed framework. Allowing for two-dimensional decompositions is key, as slab-decomposed approaches are bound to breakdown for DNS at sufficiently high Reynolds number.

Carefully handling *z*-implicit diffusion was key to secure the improved parallel performance at scale. To this end, we have proposed a re-worked flavor of the original parallel cyclic reduction – TDMA approach presented in László et al. [19]. We have shown that, by re-working the algorithm into a pre-processing and runtime step, one can easily solve the three linear systems per time iteration associated with this implicit discretization.

We have tested the different approaches at scale, using up to 1 024 GPUs/GCDs on the supercomputers Leonardo and LUMI. The results of the scalability tests reveal that the new distributed Poisson solver shows compelling performance gains for 2D pencil decompositions, being approximately twice faster in the LUMI and Leonardo supercomputers than the original CaNS version based on the full-transpose approach using 1 024 GPUs/GCDs. A detailed analysis of the GPU timer profiles reveals that the performance improvements are largely due to the reduced size of the global all-to-all transpose operations among MPI tasks. The scalability of the DNS solver was also tested in large-scale simulations of wall-bounded flows, benchmarking the performance of entire physical time steps. At scale, the new approach was found to be approximately 1.5× faster in the LUMI and Leonardo supercomputers with a 2D pencil decomposition of  $(p_y \times p_z) = (512 \times 2)$  while completing entire physical time steps. The observed performance differences between the two machines were understood by inspecting the best recorded times per algorithm step, showing that LUMI runs experienced higher latency than Leonardo. In general, we find that minimizing the number of  $p_z$  partitions in 2D pencil decompositions reduces the running times for the DNS solver with either the full-transpose method or the parallel tridiagonal algorithm. This is attributed to the reduced cost of performing transposes in the *y*-*z* directions. Additionally, we highlight that the DNS solver, coupled with the parallel tridiagonal algorithm, can be configured to work with 2D pencil decompositions achieving identical running times as the most optimized 1D slab configurations available for medium-scale systems.

Regarding implementation, while the underlying numerical solver works on Leonardo out-of-the-box, some work was required to successfully run it on LUMI. As a by-product of the present effort, a cross-platform communication library *diezDecomp* was developed for halo exchanges and any-to-any transpose operations between MPI ranks with mismatched local problem sizes. While we could have achieved the same

by modifying *cuDecomp*, we found some advantages in having a simpler library in Modern Fortran as an alternative with fewer dependencies.

Overall, this approach will enable DNS of turbulent wall flows at unprecedented scales, helping to bridge the gap between current setups that can be studied using first-principles simulations, and important applications in environmental and engineering systems.

#### CRediT authorship contribution statement

**Rafael Diez Sanhueza:** Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Jurriaan Peeters:** Writing – review & editing, Writing – original draft, Validation, Supervision, Software, Resources, Project administration, Investigation, Conceptualization. **Pedro Costa:** Writing – review & editing, Writing – original draft, Validation, Supervision, Software, Resources, Project administration, Methodology, Investigation, Data curation, Conceptualization.

#### Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Pedro Costa reports financial support was provided by European High Performance Computing Joint Undertaking through the grants no. EHPC-EXT-2022E01-054 and no. EHPC-DEV-2024D04-039. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

#### Acknowledgements

The performance tests on the Leonardo supercomputer, based at CINECA (Italy), were enabled by the EuroHPC grant no. EHPC-EXT-2022E01-054. Access to the LUMI supercomputer was granted through the project no. EHPC-DEV-2024D04-039, and the NWO Large Compute Grant no. EINF-10608. We also acknowledge NWO for providing access to the Snellius supercomputer, based at SURF (Netherlands). We thank Josh Romero, Massimiliano Fatica, and Sergio Pirozzoli for the insightful discussions. Finally, we would also like to thank the two anonymous Reviewers for their careful review, which contributed to the improvement of the manuscript.

#### Appendix A. GPU implementation of Fourier-based transform

While a real-to-complex Fourier transform of a signal  $\mathbf{x}$  with  $n$  numbers has  $n/2 + 1$  complex numbers, the imaginary parts of the first and last elements (for even  $n$ ) are zero. Let us consider the output of the real-to-complex Fourier transform of

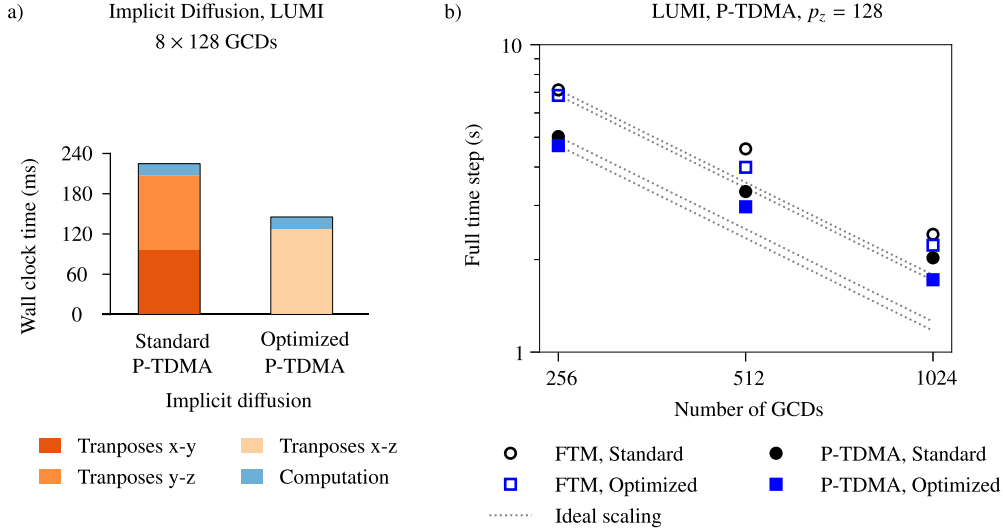
$$\mathbf{x} = [x_0 \ x_1 \ \dots \ x_{n-1}], \quad (\text{A.1})$$

given by

$$\tilde{\mathbf{x}} = [\tilde{x}_0^r \ \tilde{x}_0^i \ \tilde{x}_1^r \ \tilde{x}_1^i \ \dots \ \tilde{x}_{[n/2+1]}^r \ \tilde{x}_{[n/2+1]}^i]; \quad (\text{A.2})$$

$\tilde{\mathbf{x}}$  has  $[n/2 + 1]$  elements, with  $[ ]$  denoting the integer floor operation. Since each complex number is represented by two real ones,  $\tilde{\mathbf{x}}$  is represented by  $(2 [n/2 + 1])$  real numbers, with  $\tilde{x}_0^i = 0$ , and  $\tilde{x}_{[n/2+1]}^i = 0$  for even  $n$ . Hence, the real-to-complex transform can be uniquely represented by a set of  $n$  numbers. This property is explored in several FFT packages (e.g., *FFTPACK*, and the *half-complex* format of *FFTW* used in the CaNS code for CPU-based runs). Unfortunately, popular GPU-based FFT libraries like *cuFFT*, *hipFFT*, or *MKL* do not support this format [39,40,43].

Representing the output of the real-to-complex transforms in arrays of size  $n$  is desirable, as it allows us to handle the output of a real-to-complex transform in the same manner as a real-to-real transform,



**Fig. B.8.** Comparison between the standard P-TDMA approach and the optimized version using direct  $x \rightarrow z$  transpose operations for implicit 1D diffusion (subplot a) using  $8 \times 128$  GCDs and a grid size of  $(N_x \times N_y \times N_z) = (7168 \times 7168 \times 1594)$ . The subplot (b) corresponds to a strong scalability test for the entire DNS solver using the same grid size, but a different number of GCDs in the LUMI supercomputer.

greatly simplifying the implementation of different transform types in the Poisson solver. Hence, since the first GPU version of the CaNS code [17],  $\tilde{x}$  is packed in the following format:

$$\tilde{x}^r = \begin{cases} \begin{bmatrix} \tilde{x}_0^r & \tilde{x}_0^i & \dots & \tilde{x}_{[n/2+1]}^r & \tilde{x}_{[n/2+1]}^i \end{bmatrix} & \text{if } n \text{ is even,} \\ \begin{bmatrix} \tilde{x}_0^r & \tilde{x}_0^i & \dots & \tilde{x}_{[n/2+1]}^r & \tilde{x}_{[n/2+1]}^i \end{bmatrix} & \text{if } n \text{ is odd.} \end{cases} \quad (\text{A.3})$$

It is easily seen that both cases have  $n$  real elements. This operation has  $\mathcal{O}(1)$  time complexity, while re-arranging the signal such as  $[\tilde{x}_0^r, \dots, \tilde{x}_1^i, \dots]$ , would have  $\mathcal{O}(n)$  complexity. With this cheaper re-arrangement of the arrays, the Fourier eigenvalues  $\lambda_i$  and  $\lambda_j$  in eq. (12) must be consistently re-ordered to comply with this format. This is an inexpensive operation that is performed during the initialization of the Poisson/Helmholtz solver.

Finally, the reciprocal unpacking operations are done for performing the inverse complex-to-real transform to have an input array with  $[n/2 + 1]$  elements, resulting in an output signal with  $n$ -elements in the correct order.

## Appendix B. Performance gains from direct $x \rightarrow z$ transposes with implicit 1D diffusion

When the GPU profiling results from Fig. 7 are analyzed, it can be noticed that the implicit 1D diffusion solver performs two consecutive transposes in the  $x$ - $y$  and  $y$ - $z$  directions, which will be denoted as  $x \rightarrow y \rightarrow z$  in this section. This is suboptimal. Ideally, the cyclic reduction process should be performed using the original  $x$ -aligned pencils for the velocity components  $u/v/w$ , and then direct  $x \rightarrow z$  transposes should be used before solving the reduced systems of tridiagonal equations.

To better understand the benefits of performing a direct  $x \rightarrow z$  transpose instead of two consecutive transposes, we used the *diezDecomp* library created for the LUMI porting effort. The implementation intersects the  $x/y/z$  bounds of different MPI tasks, with no strong restrictions, and thus it is trivial to implement any variant of  $x \rightarrow z$  transpose. This allowed for an implementation of this more complex communication operation with minimal changes in the DNS code.

While avoiding  $x \rightarrow y \rightarrow z$  transposes has a small impact in the GPU profiling results shown in Fig. 7, we identified other DNS cases

where the impact of this transpose sequence was much higher. For instance, the DNS cases with  $(p_y \times p_z) = (8 \times 128)$  have a much higher MPI workload for the parallel tridiagonal solver (due to the reduced size of  $p_y$ ), and thus they benefit more from removing  $x \rightarrow y \rightarrow z$  transposes. In Fig. B.8, the results of the scalability tests using  $x \rightarrow z$  transposes are presented for a 2D pencil decomposition with  $p_z = 128$  vertical partitions. The configuration  $p_z = 128$  was chosen, since its parallel tridiagonal solver works with larger arrays and the unnecessary  $x \rightarrow y$  transposes have a significant impact in the results previously shown in Fig. 4. In the subplot (a), it can be seen that the running times for the implicit 1D diffusion solver are 55% slower when successive  $x$ - $y$  and  $y$ - $z$  transposes are used. The scalability chart at the right reveals that the system with  $x \rightarrow z$  transposes is more efficient in large-scale simulations, reducing the running times of the entire DNS solver by 18% for the P-TDMA algorithm with 1024 GCDs.

## Data availability

Data will be made available on request.

## References

- [1] S.A. Orszag, G.S. Patterson, Numerical simulation of three-dimensional homogeneous isotropic turbulence, *Phys. Rev. Lett.* 28 (1972) 76–79.
- [2] P.K. Yeung, K. Ravikumar, Advancing understanding of turbulence through extreme-scale computation: intermittency and simulations at large problem sizes, *Phys. Rev. Fluids* 5 (2020) 110517.
- [3] T. Ishihara, T. Gotoh, Y. Kaneda, Study of high-Reynolds number isotropic turbulence by direct numerical simulation, *Annu. Rev. Fluid Mech.* 41 (2009) 165–180.
- [4] D. Reed, D. Gannon, J. Dongarra, HPC forecast: cloudy and uncertain, *Commun. ACM* 66 (2023) 82–90.
- [5] J. Nickolls, W.J. Dally, The GPU computing era, *IEEE MICRO* 30 (2010) 56–69.
- [6] CUDA Fortran for scientists and engineers, M. Fatica, G. Ruetsch (Eds.), *CUDA Fortran for Scientists and Engineers*, Morgan Kaufmann, Boston, 2014.
- [7] F. Salvadore, M. Bernardini, M. Botti, GPU accelerated flow solver for direct numerical simulation of turbulent flows, *J. Comput. Phys.* 235 (2013) 129–142.
- [8] Y. Kim, D. Ghosh, E.M. Constantinescu, R. Balakrishnan, GPU-accelerated DNS of compressible turbulent flows, *Comput. Fluids* 251 (2023) 105744.
- [9] AMD Instinct MI250X Accelerators, AMD, <https://www.amd.com/en/products/accelerators/instinct/mi200/mi250x.html>, 2025.
- [10] NVIDIA A100 Tensor Core GPU: Unprecedented acceleration at every scale, Nvidia, <https://www.nvidia.com/content/dam/en-zz/Solutions/Data-Center/a100/pdf/nvidia-a100-datasheet-nvidia-us-2188504-web.pdf>, 2025.
- [11] D. De Sensi, L. Pichetti, F. Vella, T. De Matteis, Z. Ren, L. Fusco, M. Turisini, D. Cesarini, K. Lust, A. Trivedi, D. Roweth, F. Spiga, S. Di Girolamo, T. Hoefler, Exploring

- gpu-to-gpu communication: insights into supercomputer interconnects, in: SC24: International Conference for High Performance Computing, Networking, Storage and Analysis, 2024, pp. 1–15.
- [12] G. Schieffer, R. Shi, S. Markidis, A. Herten, J. Faj, I. Peng, Understanding data movement in amd multi-gpu systems with infinity fabric, in: Proceedings of the SC '24 Workshops of the International Conference on High Performance Computing, Network, Storage, and Analysis, SC-W '24, IEEE Press, 2025, pp. 567–576.
- [13] M. Bernardini, D. Modesti, F. Salvatore, S. Pirozzoli, STREAmS: a high-fidelity accelerated solver for direct numerical simulation of compressible turbulent flows, *Comput. Phys. Commun.* 263 (2021) 107906.
- [14] F. De Vanna, F. Avanzi, M. Cogo, S. Sandrin, M. Bettencourt, F. Picano, E. Benini, URANOS: a GPU accelerated Navier-Stokes solver for compressible wall-bounded flows, *Comput. Phys. Commun.* 287 (2023) 108717.
- [15] X. Zhu, E. Phillips, V. Spandan, J. Donners, G. Ruetsch, J. Romero, R. Ostilla-Mónico, Y. Yang, D. Lohse, R. Verzicco, M. Fatica, R.J. Stevens, AFID-GPU: a versatile Navier-Stokes solver for wall-bounded turbulent flows on GPU clusters, *Comput. Phys. Commun.* 229 (2018) 199–210.
- [16] P. Costa, A FFT-based finite-difference solver for massively-parallel direct numerical simulations of turbulent flows, *Comput. Math. Appl.* 76 (2018) 1853–1862.
- [17] P. Costa, E. Phillips, L. Brandt, M. Fatica, GPU acceleration of CaNS for massively-parallel direct numerical simulations of canonical fluid flows, in: Development and Application of Open-Source Software for Problems with Numerical PDEs, *Comput. Math. Appl.* 81 (2021) 502–511.
- [18] J. Romero, P. Costa, M. Fatica, Distributed-memory simulations of turbulent flows on modern GPU systems using an adaptive pencil decomposition library, in: Proceedings of the Platform for Advanced Scientific Computing Conference, PASC '22, Association for Computing Machinery, New York, NY, USA, 2022.
- [19] E. László, M. Giles, J. Appleyard, Manycore algorithms for batch scalar and block tridiagonal solvers, *ACM Trans. Math. Softw.* 42 (2016).
- [20] M. Yang, J.-H. Kang, K.-H. Kim, O.-K. Kwon, J.-I. Choi, PaScaLTDMA 2.0: a multi-GPU-based library for solving massive tridiagonal systems, *Comput. Phys. Commun.* 290 (2023) 108785.
- [21] K.-H. Kim, J.-H. Kang, X. Pan, J.-I. Choi, PaScaLTDMA: a library of parallel and scalable solvers for massive tridiagonal systems, *Comput. Phys. Commun.* 260 (2021) 107722.
- [22] Z. Gong, G. Deng, C. An, Z. Wu, X. Fu, A high order finite difference solver for simulations of turbidity currents with high parallel efficiency, *Comput. Math. Appl.* 128 (2022) 21–33.
- [23] M. Yang, G. Oh, T. Xu, J. Kim, J.-H. Kang, J.-I. Choi, Multi-GPU-based real-time large-eddy simulations for urban microclimate, *Build. Environ.* 245 (2023) 110856.
- [24] S. Pirozzoli, P. Orlandi, Natural grid stretching for DNS of wall-bounded flows, *J. Comput. Phys.* 439 (2021) 110408.
- [25] S.B. Pope, Turbulent flows, *Meas. Sci. Technol.* 12 (2001) 2020–2021.
- [26] TOP500 list - November 2024, <https://www.top500.org/lists/top500/list/2024/11>, 2024.
- [27] A.J. Chorin, A numerical method for solving incompressible viscous flow problems, *J. Comput. Phys.* 2 (1967) 12–26.
- [28] J. Kim, P. Moin, Application of a fractional-step method to incompressible Navier-Stokes equations, *J. Comput. Phys.* 59 (1985) 308–323.
- [29] P. Moin, R. Verzicco, On the suitability of second-order accurate discretizations for turbulent flow simulations, in: Vortical Structures and Wall Turbulence, *Eur. J. Mech. B, Fluids* 55 (2016) 242–245.
- [30] E. Fadlun, R. Verzicco, P. Orlandi, J. Mohd-Yusof, Combined immersed-boundary finite-difference methods for three-dimensional complex flow simulations, *J. Comput. Phys.* 161 (2000) 35–60.
- [31] W.P. Breugem, B.J. Boersma, Direct numerical simulations of turbulent flow over a permeable wall using a direct and a continuum approach, *Phys. Fluids* 17 (2005) 025103.
- [32] M. Uhlmann, An immersed boundary method with direct forcing for the simulation of particulate flows, *J. Comput. Phys.* 209 (2005) 448–476.
- [33] G. Tryggvason, R. Scardovelli, S. Zaleski, Direct Numerical Simulations of Gas-Liquid Multiphase Flows, Cambridge University Press, 2011.
- [34] A.A. Wray, Very low storage time-advancement schemes, Technical Report, Internal Report, NASA-Ames Research Center, Moffett Field, CA, 1986.
- [35] U. Schumann, R.A. Sweet, Fast Fourier transforms for direct solution of Poisson's equation with staggered boundary conditions, *J. Comput. Phys.* 75 (1988) 123–137.
- [36] K.-H. Kim, J.-H. Kang, X. Pan, J.-I. Choi, PaScaL-TCS: a versatile solver for large-scale turbulent convective heat transfer problems with temperature-dependent fluid properties, *Comput. Phys. Commun.* 290 (2023) 108779.
- [37] S. Rolfo, C. Flageul, P. Bartholomew, F. Spiga, S. Laizet, The 2DECOMP&FFT library: an update with new CPU/GPU capabilities, *J. Open Source Softw.* 8 (2023) 5813.
- [38] R. Diez, diezDecomp: cross-platform library for transposes and halo exchanges in extreme-scale simulations, <https://github.com/Rafael10Diez/diezDecomp>, 2025.
- [39] cuFFT API reference, <https://docs.nvidia.com/cuda/cufft>, 2024.
- [40] hipFFT documentation, <https://rocm.docs.amd.com/projects/hipFFT/en/latest>, 2024.
- [41] ROCm Development Team, hipFort: Fortran interfaces for ROCm libraries, <https://github.com/ROCm/hipfort>, 2025, Version 0.6.1 (latest release as of May 2025).
- [42] M. Lee, R.D. Moser, Direct numerical simulation of turbulent channel flow up to  $Re_\tau \approx 5200$ , *J. Fluid Mech.* 774 (2015) 395–415.
- [43] Intel oneAPI Math Kernel Library (oneMKL), <https://www.intel.com/content/www/us/en/developer/tools/oneapi/oneapi.html>, 2024.