Contents lists available at ScienceDirect





journal homepage: www.elsevier.com/locate/compfluid



GPU-accelerated DNS of compressible turbulent flows*

Youngdae Kim^{a,1}, Debojyoti Ghosh^{b,*}, Emil M. Constantinescu^a, Ramesh Balakrishnan^c

^a Mathematics and Computer Science Division, Argonne National Laboratory, 9700 S. Cass Ave., Lemont, 60439, IL, United States ^b Center for Applied Scientific Computing, Lawrence Livermore National Laboratory, 7000 East Ave., Livermore, 94550, CA, United States ^c Computational Science Division, Argonne National Laboratory, 9700 S. Cass Ave., Lemont, 60439, IL, United States

ARTICLE INFO

MSC:

76F65

65Y05

76F05

35030

65M06

GPUs

Keywords:

ABSTRACT

Navier-Stokes equations WENO schemes Compressible flows Direct numerical simulation

This paper explores strategies to transform an existing CPU-based high-performance computational fluid dynamics solver, HyPAR, for compressible flow simulations on emerging exascale heterogeneous (CPU+GPU) computing platforms. The scientific motivation for developing a GPU-enhanced version of HyPAR is to simulate canonical turbulent flows at the highest resolution possible on such platforms. We show that optimizing memory operations and thread blocks results in 200x speedup of computationally intensive kernels compared with a CPU core. Using multiple GPUs and CUDA-aware MPI communication, we demonstrate both strong and weak scaling of our GPU-based HyPAR implementation on the NVIDIA Volta V100 GPUs. We simulate the decay of homogeneous isotropic turbulence in a triply periodic box on grids with up to 1024³ points (5.3 billion degrees of freedom) and on up to 1,024 GPUs. We compare the wall times for CPU-only and CPU+GPU simulations. The results presented in the paper are obtained on the Summit and Lassen supercomputers at Oak Ridge and Lawrence Livermore National Laboratories, respectively.

1. Introduction

Turbulence is characterized by the seemingly chaotic, yet correlated, motion of fluid elements (in both the Eulerian and Lagrangian sense) over a range of spatial and temporal scales. The distinguishing features of turbulence, however, are the lack of clear separation between the scales of motion (unlike in the kinetic theory of gases), the large number of degrees of freedom, and the large range of length and time scales. Although it is accepted that the dynamics of turbulence is described by the Navier-Stokes equations (NSE), it is prohibitively expensive to simulate all the length and time scales of motion. For many cases of interest, it is practically impossible to simulate turbulence without appropriate models to account for the effect of the small scales on the evolution of the turbulent flow field. As a result, various theoretical statistical models have been proposed to explain the underlying mechanisms in turbulent flow by decomposing the velocity \vec{U} , pressure p, and temperature T fields into the average and fluctuating

quantities, $A(\vec{r}, t) = \bar{A}(\vec{r}, t) + a(\vec{r}, t)$, where \bar{A} and a respectively denote the ensemble average $\langle A \rangle$ and fluctuating *a* terms in the decomposition. Approaches for formulating statistical models of turbulence [1], such as direct interaction approximation [2,3] and eddy-damped quasi-normal Markovianization [4], proposed evolution equations for higher-order correlations that are obtained by taking moments of the NSE. In these equations, the temporal evolution of correlations, on the left-hand side of the NSE, calls for modeling higher-order moments on the right-hand side of the NSE, where again the strong spatiotemporal correlations between the velocity components $\langle u_i^m u_i^n \rangle$ and between other flow quantities have precluded attempts to close the hierarchy of the moment equations and devise a universal model for turbulent flows.

The availability of massively parallel computing platforms offers a route for direct numerical simulation (DNS) of the Navier-Stokes equations for canonical flows on simple geometries and increasing Reynolds numbers. One such canonical flow that continues to be simulated on

https://doi.org/10.1016/j.compfluid.2022.105744

Received 30 June 2022; Received in revised form 15 October 2022; Accepted 25 November 2022 Available online 28 November 2022

0045-7930/Published by Elsevier Ltd.

This research was supported by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of the U.S. Department of Energy Office of Science and the National Nuclear Security Administration, and by the U.S. Department of Energy, Office of Science under contract number DE-AC02-06CH11357. Y. Kim and R. Balakrishnan were partially supported by the ECP/ProxyApps sub-project. E. Constantinescu was partially supported by U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research (ASCR). Part of this work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract No. DE-AC52-07NA27344.

Corresponding author.

youngdaekim26@gmail.com (Y. Kim), ghosh5@llnl.gov (D. Ghosh), emconsta@anl.gov (E.M. Constantinescu), bramesh@anl.gov E-mail addresses: (R. Balakrishnan).

¹ Present address: ExxonMobil Technology and Engineering Company, Annandale, NJ, 08801, United States. This work was done while the first author was at Argonne National Laboratory as a postdoctoral appointee.

newer parallel computing platforms is that of forced isotropic turbulence in a triply periodic box. Using innovative methods to stir the flow (to ensure that the turbulence does not decay), pseudo-spectral solvers for incompressible flows have been used to simulate homogeneous isotropic turbulence (HIT) on boxes consisting of 12288^3 grid points to achieve a maximum Reynolds number $Re_{\lambda} = 1300$ (where λ denotes the Taylor microscale) [5]. These DNS flow fields can then be used to validate existing theories and inform the development of subgrid models by exploring the general two-point, two-time, space–time velocity correlation, defined by

$$\mathcal{R}(\vec{x},t) = \langle u_i(\vec{x},t)u_i(\vec{x}+\vec{r},t+\tau) \rangle, \tag{1}$$

where $u_i, u_j(i, j = 1, 2, 3)$ denotes the fluctuating velocity components, $\vec{x} + \vec{r} = (x_1 + \delta x_1, x_2 + \delta x_2, x_3 + \delta x_3)$ denotes the distance between two points in the flow, τ denotes the temporal window, and $\langle . \rangle$ denotes the ensemble average. The flow is said to be decorrelated when the correlation decreases with increasing $|\vec{r}|$ and τ and vanishes above a critical distance $|\vec{r}|_d$ and τ_d that denote the decorrelation length and time, respectively. Similar expressions can be defined to correlate other quantities, such as the pressure (p) and temperature (T) between two points and at two different times.

While a considerable body of literature exists for HIT simulations of incompressible flows, HIT simulations of compressible flows are not as extensive. DNS of compressible turbulence requires solving the energy equation in addition to the momentum equations. Another complication that arises in simulating compressible flows is the appearance of shocks and shocklets in compressible HIT, where the flow properties change abruptly across the shocklets. Hence, shockcapturing numerical methods are necessary that also resolve all the turbulent scales of motion with minimal dissipation and dispersion errors. Needed are high accuracy, spectral resolution, and the ability to capture unsteady shocklets/shock waves. Weighted essentially nonoscillatory (WENO) schemes [6-9] use solution-dependent reconstruction to achieve high-order accuracy for smooth flows and yield nonoscillatory solution across shocks and other discontinuities. Although WENO schemes of very high orders have been designed [10,11], they suffer from relatively poor spectral resolution and are ill-suited for the DNS of turbulent flows. Several strategies were introduced to improve their dissipation and dispersion properties. One strategy is improving the calculation of the nonlinear weights [11-15] such that they attain their optimal values when the flow is smooth but not wellresolved. Another strategy is improving the spectral resolution of the underlying finite-difference scheme, for example, bandwidth-optimized schemes [16], dissipation-relation-preserving scheme [17], and minimized dispersion schemes with controllable dissipation [18-21]. Given the high spectral resolution of compact or Padé-type schemes [22], a related strategy has been using compact finite-difference schemes as the underlying linear discretization; examples include weighted compact schemes (WCNS) [23-25] and CRWENO [26-30] schemes. As an alternative to the WENO schemes, targeted essentially nonoscillatory (TENO) schemes were introduced [31] that modify the stencil selection procedure to reduce the dispersion and dissipation errors at high wavenumbers; they have been successfully applied to the DNS of compressible turbulent flows [32]. Many efforts combine these ideas in hybrid schemes, for example, hybrid central-upwind finite-difference schemes [33] and hybrid compact-WENO [34-36].

GPUs require a fundamentally different programming paradigm from that of CPUs. Some memory operations that are efficient on CPUs could cause a serious performance degradation on GPUs. The past decade has seen several GPU implementations of computational fluid dynamics (CFD) codes for simulating turbulent flows. Linear, highorder finite-difference schemes were implemented on GPUs and applied to the DNS of turbulent flows, for example, 4th-order [37] and 6thorder [38,39] central schemes. These studies reported speedups by factors of 10–20 when compared with a CPU-only implementation. High-order WENO schemes have been implemented and assessed on GPUs [40-42], including on one of the fastest supercomputers [43]. One conclusion of these studies was that the WENO kernel was the most expensive component for a typical NSE code. The GPU implementation of WENO-type schemes was applied to the Favre-averaged NSE [44,45]. DNS codes based on WENO-type schemes have been implemented on GPUs to simulate supersonic turbulent flows, such as the shock-turbulence interaction [46], supersonic channel flow, shock-boundary layer interaction [47], and supersonic compressionexpansion slope flow [48]. Given the poor spectral resolution of the traditional WENO scheme, GPU implementations of more modern variants such as the TENO schemes were developed [49-51]. Given the suitability of compact/pseudo-spectral schemes, a linear 6th-order tridiagonal compact scheme with filtering was implemented for GPUs [52] to simulate the evolution of eddies. However, the algorithm suffered from an inefficient implementation of the tridiagonal solve on the GPU that did not scale well with the problem size. Heterogeneous (CPU+GPU) implementations of two shock-capturing compact schemes, the weighted compact nonlinear scheme (WCNS) and hybrid dissipative compact scheme, were developed [53] and applied to simulate flow around aircraft configurations.

The objective of this paper is a computationally efficient heterogeneous (CPU+GPU) implementation of an existing scalable, CPU-parallel shock-capturing CFD algorithm for the DNS of canonical turbulent flows. We report our strategies to accelerate an existing open-source MPI-parallel C/C++ code, HyPAR [54], with GPUs. We consider the fifth-order WENO scheme (WENO5) [7]; in subsequent publications, we will consider high-resolution alternatives that are better suited for DNS of turbulent flows, such as the compact-WENO and low-dissipation WENO schemes. Since efficient memory operations are crucial for fully utilizing the massive parallel computing capability of GPUs, our focus is on optimizing the memory access patterns of HyPAR on GPUs. In particular, we aim to avoid slow memory transactions and reduce as many wasted computational resources (such as thread blocks) and memory transactions as possible. Our optimization approach results in more than 200 times faster computation time using multiple threads on a GPU compared with a CPU core. Using multiple GPUs and CUDA-aware MPI communication, we demonstrate the scalability of our GPU-based implementation. We evaluate strong and weak scaling and investigate the impact of communication cost on the overall computation time. We also simulate the decay of HIT in a periodic domain with up to 1024³ grids points (~ 5×10^9 billion degrees of freedom) on up to 1024 GPUs and compare the wall time with MPI-only simulations on up to 8192 CPU cores. As a highlight, our current GPU-based implementation can perform 100 time steps on a grid with 1024^3 (~10⁹) points in less than 30 s using 1024 GPUs.

Our implementation presented here is among the earliest attempts at simulating scale-resolved compressible HIT on DOE leadership-class heterogeneous (*i.e.*, CPU+GPU) platforms. Recent work for incompressible turbulence at such scales includes simulations with up to ~21 billion grid points on 1280 NVIDIA P100 GPUs [55,56] and a pseudospectral algorithm for grids with up to 18,432³ points and 18,432 NVIDIA V100 GPUs [57]. The simulations presented in this paper were carried out on the Summit supercomputer at the Oak Ridge Leadership Computing Facility (OLCF) [58] (referred as "OLCF/Summit" in the rest of the paper) and the Lassen supercomputer at Lawrence Livermore National Laboratory [59] (referred as "LLNL/Lassen"). The implementation described here has been released in the latest version of HyPAR [54] under an MIT license.

The rest of the paper is organized as follows. In Section 2 we introduce our simulation: the governing equations for compressible Navier–Stokes equations, the numerical methods, and the HIT decay example. Section 3 describes our kernel design principles for efficiently implementing the existing CPU code-base on GPUs, with a specific focus on optimizing its memory access patterns. In Section 4 we report the computational performance and scalability tests of the GPU-based implementation. Section 5 describes high-resolution simulations of HIT decay on very fine grids and compares the CPU+GPU and CPU-only wall times. We summarize our conclusions in Section 6.



Fig. 1. Illustration of cell centers and cell interfaces on a grid line along a specific dimension.

2. Background

In this section we briefly introduce governing equations for compressible Navier–Stokes equations, their numerical methods, and the isotropic turbulence decay example. The example will be used to evaluate the computational performance of our GPU-based implementation in Section 4.

2.1. Governing equations

The nondimensional, compressible Navier–Stokes equations [60] are expressed as a system of hyperbolic-parabolic partial differential equations as

$$\frac{\partial \mathbf{u}}{\partial t} + \frac{\partial \mathbf{f}}{\partial x} + \frac{\partial \mathbf{g}}{\partial y} + \frac{\partial \mathbf{h}}{\partial z} = \frac{\partial \mathbf{f}^v}{\partial x} + \frac{\partial \mathbf{g}^v}{\partial y} + \frac{\partial \mathbf{h}^v}{\partial z},\tag{2}$$

where $\mathbf{u} = [\rho, \rho u, \rho v, \rho w, e]^{\mathrm{T}}$ is the vector of conserved variables. The convective fluxes are

$$\mathbf{f} = \begin{bmatrix} \rho u \\ \rho u^2 + p \\ \rho uv \\ \rho uw \\ (e+p)u \end{bmatrix}, \ \mathbf{g} = \begin{bmatrix} \rho v \\ \rho uv \\ \rho vv \\ (e+p)v \end{bmatrix}, \ \mathbf{h} = \begin{bmatrix} \rho w \\ \rho uw \\ \rho uw \\ \rho vw \\ \rho w^2 + p \\ (e+p)w \end{bmatrix},$$
(3)

where ρ is the density, u, v, w are the Cartesian components of the velocity, p is the pressure, and e is the internal energy given by

$$e = \frac{p}{\gamma - 1} + \frac{1}{2}\rho\left(u^2 + v^2 + w^2\right),$$
(4)

and $\gamma = 1.4$ is the heat coefficient ratio. The viscous fluxes are

$$\mathbf{f}^{\nu} = \begin{bmatrix} \mathbf{0} \\ \tau_{xx} \\ \tau_{yx} \\ \tau_{zx} \\ \mathbf{v} \cdot \boldsymbol{\tau}_{x} - q_{x} \end{bmatrix}, \ \mathbf{g}^{\nu} = \begin{bmatrix} \mathbf{0} \\ \tau_{xy} \\ \tau_{yy} \\ \tau_{zy} \\ \mathbf{v} \cdot \boldsymbol{\tau}_{y} - q_{y} \end{bmatrix}, \ \mathbf{h}^{\nu} = \begin{bmatrix} \mathbf{0} \\ \tau_{xz} \\ \tau_{yz} \\ \tau_{zz} \\ \mathbf{v} \cdot \boldsymbol{\tau}_{z} - q_{z} \end{bmatrix};$$
(5)

the viscous stresses are given by

$$\tau_{ij} = \mu \frac{M_{\infty}}{Re_{\infty}} \left[\left(\frac{\partial u_i}{\partial x_j} - \frac{\partial u_j}{\partial x_i} \right) - \frac{2}{3} \frac{\partial u_k}{\partial x_k} \delta_{ij} \right],\tag{6}$$

where Re_{∞} and M_{∞} are the reference Reynolds and Mach numbers, respectively, μ is the normalized coefficient of viscosity, $\mathbf{v} \equiv (u, v, w)$ is the velocity vector, and

$$\boldsymbol{\tau}_{(\cdot)} = \left(\tau_{x}, \tau_{y}, \tau_{z}\right). \tag{7}$$

The thermal conduction terms are

$$q_i = -\frac{\mu}{(\gamma - 1)} \frac{M_{\infty}}{Re_{\infty}} \frac{\partial T}{Pr} \frac{\partial T}{\partial x_i},$$
(8)

where $T = \gamma p / \rho$ is the temperature and Pr = 0.72 is the Prandtl number.

2.2. Numerical method

The physical domain is discretized by using a three-dimensional grid with N points in each dimension. The convective terms on the left-hand side of (2) are discretized with the conservative finite-difference formulation [8,9] and the WENO5 scheme [7]. This section briefly describes the spatial discretization of the convective flux **f** along x; this can be

extended trivially to the convective terms along the other dimensions. The spatial derivatives in the viscous terms on the right-hand side (RHS) are discretized by using fourth-order central finite differences. We thus obtain an ordinary differential equation (ODE) in time; this is evolved by using the third-order total-variation-diminishing (TVD) and fourth-order Runge–Kutta methods. A more detailed description of the numerical methodology is available in prior publications [26,27,61].

We describe the upwind discretization of the convective terms by considering an arbitrary grid line along x and ignoring the other dimensions. The derivative is computed at a grid point j as

$$\frac{\partial \mathbf{f}}{\partial x}\Big|_{j} = \frac{1}{\Delta x} \left(\hat{\mathbf{f}}_{j+\frac{1}{2}} - \hat{\mathbf{f}}_{j-\frac{1}{2}} \right) + \mathcal{O}\left(\Delta x^{p}\right), \tag{9}$$

where the locations $j \pm 1/2$ denote the cell interfaces (see Fig. 1) and p is the order of the discretization scheme. The numerical flux at the cell interface, $\hat{\mathbf{f}}_{j+\frac{1}{2}}$, is computed by using the Roe scheme [62] with the Harten entropy fix [63]:

$$\hat{\mathbf{f}}_{j+\frac{1}{2}} = \frac{1}{2} \left(\hat{\mathbf{f}}_{j+\frac{1}{2}}^{L} + \hat{\mathbf{f}}_{j+\frac{1}{2}}^{R} \right) - \frac{1}{2} \left| A_{j+\frac{1}{2}} \right| \left(\hat{\mathbf{u}}_{j+\frac{1}{2}}^{R} - \hat{\mathbf{u}}_{j+\frac{1}{2}}^{L} \right),$$
(10)

where the superscripts *L*, *R* denote left- and right-biased discretizations, respectively. The dissipation matrix, $\begin{vmatrix} A_{j+\frac{1}{2}} \end{vmatrix}$, is computed from the eigensystem evaluated at the Roe-averaged state at the cell interface j + 1/2 as

$$A_{j+\frac{1}{2}} = X_{j+\frac{1}{2}} \left| A_{j+\frac{1}{2}} \right| X_{j+\frac{1}{2}}^{-1},$$
(11)

where *X* is the matrix with columns as the right eigenvectors and Λ is a diagonal matrix with the eigenvalues as its entries.

Each scalar component of $\hat{\mathbf{f}}_{j+\frac{1}{2}}^{L,R}$ and $\hat{\mathbf{u}}_{j+\frac{1}{2}}^{L,R}$ is computed by using the WENO5 scheme. We describe this method for a left-biased scalar variable, $\hat{f}_{j+\frac{1}{2}}^{L}$; the procedure for a right-biased scalar variable, $\hat{\mathbf{f}}_{j+\frac{1}{2}}^{R}$, can be obtained by reflecting the expressions around the cell interface j + 1/2. We drop the superscript L in the text below.

WENO schemes use a solution-dependent interpolation stencil selection [6,7] to achieve high-order accuracy where the solution is smooth and to avoid oscillations across discontinuities. The WENO5 scheme is constructed by identifying three third-order interpolation schemes at the cell interface; the final interpolation method is their weighted sum, where the weights depend on the smoothness of the stencils underlying the corresponding third-order interpolation. These weights approach optimal values for smooth solutions; consequently, the method achieves fifth-order accuracy. In the presence of discontinuities or sharp gradients, the weights corresponding to those stencils approach zero. The final method has a stencil biased away from the discontinuity, thus avoiding numerical oscillations.

The three third-order interpolation schemes and their optimal weights at cell interface j + 1/2 are

$$\hat{f}_{j+1/2}^{1} = \frac{1}{3}f_{j-2} - \frac{7}{6}f_{j-1} + \frac{11}{6}f_{j}, \ c_{1} = \frac{1}{10},$$
(12)

$$\hat{f}_{j+1/2}^2 = -\frac{1}{6}f_{j-1} + \frac{5}{6}f_j + \frac{1}{3}f_{j+1}, \ c_2 = \frac{6}{10}, \tag{13}$$

$$\hat{f}_{j+1/2}^3 = \frac{1}{3}f_j + \frac{5}{6}f_{j+1} - \frac{1}{6}f_{j+2}, \ c_3 = \frac{3}{10}.$$
 (14)

The fifth-order interpolation scheme at j + 1/2 is obtained as

$$\hat{f}_{j+1/2} = \sum_{k=1,2,3} c_k \hat{f}_{j+1/2}^k$$



Fig. 2. Direct numerical simulation of HIT in a triply periodic domain discretized by a 128³ grid. The solution is obtained at a final time of $t/\tau = 3$, where τ is the turbulent time scale.

$$= \frac{1}{30}f_{j-2} - \frac{13}{60}f_{j-1} + \frac{47}{60}f_j + \frac{27}{60}f_{j+1} - \frac{1}{20}f_{j+2}.$$
 (15)

The solution-dependent weights are computed as

$$\omega_k = \frac{\alpha_k}{\sum_k \alpha_k}; \ \alpha_k = \frac{c_k}{\left(\epsilon + \beta_k\right)^p}; \ k = 1, 2, 3, \tag{16}$$

where $\epsilon = 10^{-6}$ is a small number to prevent division by zero and β_k are the smoothness indicators for the stencils:

$$\beta_1 = \frac{13}{12}(f_{j-2} - 2f_{j-1} + f_j)^2 + \frac{1}{4}(f_{j-2} - 4f_{j-1} + 3f_j)^2, \tag{17}$$

$$\beta_2 = \frac{13}{12}(f_{j-1} - 2f_j + f_{j+1})^2 + \frac{1}{4}(f_{j-1} - f_{j+1})^2, \tag{18}$$

$$\beta_3 = \frac{13}{12}(f_j - 2f_{j+1} + f_{j+2})^2 + \frac{1}{4}(3f_j - 4f_{j+1} + f_{j+2})^2.$$
(19)

The WENO5 scheme is obtained by multiplying the third-order interpolation schemes by the solution-dependent weights instead of the optimal weights:

$$\hat{f}_{j+1/2} = \sum_{k=1,2,3} \omega_k \hat{f}_{j+1/2}^k$$

$$= \frac{\omega_1}{3} f_{j-2} - \frac{1}{6} (7\omega_1 + \omega_2) f_{j-1} + \frac{1}{6} (11\omega_1 + 5\omega_2 + 2\omega_3) f_j$$

$$+ \frac{1}{6} (2\omega_2 + 5\omega_3) f_{j+1} - \frac{\omega_3}{6} f_{j+2}.$$
(20)

If the solution is locally smooth, $\omega_k \rightarrow c_k$, k = 1, 2, 3, and (20) is equivalent to (15).

2.3. Example: Isotropic turbulence decay

We simulate the decay of isotropic turbulence in a periodic domain; this is a canonical turbulent flow problem [64,65] characterized by the transfer of energy from larger to smaller length scales. Our governing equations do not include a turbulence model; we thus conduct a direct numerical simulation by resolving the Kolmogorov scale. The initial solution comprises a divergence-free velocity field with random isotropic fluctuations that satisfy the following kinetic energy spectrum:

$$E(k) = 16\sqrt{\frac{2}{\pi}} \frac{u_0^2}{k_0} \left(\frac{k}{k_0}\right)^4 \exp\left[-2\left(\frac{k}{k_0}\right)^2\right],$$
(21)

where *k* is the wavenumber, *E* is the kinetic energy, u_0 is the root-meansquare turbulence intensity, and k_0 is the wavenumber corresponding to the maximum energy. The domain is a cube of length 2π with periodic boundaries. The initial density and pressure are constant ($\rho =$ 1, $p = 1/\gamma$). We specify $u_0 = 0.3$ and $k_0 = 4$, resulting in a smooth



Fig. 3. Energy spectrum of the initial and final $(t/\tau = 3)$ solutions.

turbulent flow. The Taylor microscale-based Reynolds number, $Re_{\lambda} = \rho u_0 \lambda / \mu$, is specified as 50, where λ is the Taylor microscale.

Figs. 2(a) and 2(b) show the density and vorticity magnitude contours, respectively, on a two-dimensional plane at z = 0 for a solution obtained at $t/\tau = 3$, where $\tau = \lambda/u_0$ is the turbulent time scale. This solution is obtained on a grid with 128³ points and a time step $\Delta t = 0.025$ (CFL \approx 1). Fig. 3 shows the initial and final energy spectra for the solution. As the flow evolves and the turbulence decays, we observe an energy transfer from the lower to higher wavenumbers, as expected. The $k^{-5/3}$ relationship is also plotted for reference. We use this example in our subsequent studies.

3. Implementation principles on GPUs

In this section we present our kernel design principles to efficiently implement HyPAR on GPUs, with a specific focus on optimizing its memory access patterns. Since many core operations of HyPAR are memory-bound, it is crucial for computational performance to avoid slow memory transactions and to reduce as many wasted compute resources and memory transactions as possible. To this end, we introduce the following three design principles that target optimized computational performance on GPUs: (i) implementing lexicographic thread configuration, (ii) removing data transfers between CPUs and GPUs during operation, and (iii) performing coalesced memory access. In the following subsections we introduce these principles one by one



Fig. 4. Tiling vs. lexicographic over weights computation (16).

with numerical performance comparisons with those of alternative approaches. All of the numerical experiments in this section have been performed on a single NVIDIA's V100 using CUDA.

We note that we have also tried some user-defined caching approaches, such as manually allocating dynamic shared memory (instead of relying on the underlying GPU's automatic caching mechanism) to share function values in (17)–(19) between different grid points. However, we observed no gains or even worse performance than when we relied on the given automatic caching mechanism. Therefore, we did not employ manual shared-memory management in our implementation. We believe that such a worse performance with shared memory is partly due to a significantly improved cache performance of the Volta architecture, as described in [66, See page 17].

3.1. Lexicographic thread block configuration

By default, we have a one-to-one mapping from threads of a GPU kernel into the grid points of a discretization scheme so that each thread is assigned to a single grid point. One way to implement this mapping is to assign threads in the same lexicographic order of the coordinate values of the grid points. In this case, each thread block is one-dimensional, and the order of threads within a block follows the same lexicographic order of the corresponding grid points. An alternative way to implement the mapping is to use tiling. For example, if a grid is 64×64 , we can generate thread blocks, each of which is of size 32×8 (256 threads per block). A total of 16 thread blocks are needed to cover the entire grid in this case.

The tiling approach is intuitive; however, it could cause many wasted threads when some dimension of the grid is not a multiple of the number of threads of a warp. For example, if a given grid is of size 65×64 , the previous 32×8 thread block scheme will result in some thread blocks having many inactive threads. In this case, thread blocks covering grid points with 65 as its *x*-coordinate have only 8 active threads among 256 threads. Such an unfavorable grid naturally occurs when we perform computation using ghost points, which are added to the existing grid scheme.

Fig. 4 compares the computation time between the tiling and lexicographic thread block configurations, obtained from a kernel computing the weights (16). The grid employed is of size $65 \times 64 \times 64$, and the thread block configuration for the tiling approach is 32×8 . As we see in the figure, the tiling approach shows a slower computation time than that of the lexicographic configuration. Similar results were obtained with different block configurations for tiling.

3.2. No data transfers between CPUs and GPUs

When we solve the three-dimensional Navier-Stokes equations (2) using the WENO5 scheme, our profiling results,² described in Fig. 5 show that about 90% of the computation time has been spent in the five core operations: WENO - computing weights (16) Interp - calculation of the flux at cell interfaces (20), Upwind - computing the upwind flux using Roe's scheme (10), Parabolic - computing the viscous terms on the RHS of (2), and Derivative - computing the derivatives in the viscous terms using fourth-order central finite differences. Since



Fig. 6. Computation and data transfer times for WENO.

these operations are defined for each grid point independently, they can be computed in parallel, providing an opportunity for acceleration on GPUs. Our initial implementation of computing weights (WENO) on GPUs shows about two orders of magnitude of faster computation time than on CPUs.

However, if we include data transfer time to initiate the computation on the device and return the results back to the host, the overall computation time significantly degrades. In Fig. 6 we present the computation times on CPUs and GPUs and data transfer time for one function call of the WENO routine for $64 \times 64 \times 64$ grid. Although the computation time has improved by more than 400 times on GPUs, the time for data transfer dominates the overall computation time, making it even slower than on CPUs. The high data transfer time is attributed to the large memory size for storing weights, which is around 400 MB.

To avoid such significant degradation, we implement all operations that process data entirely on GPUs, except for the initial data read and return of a final solution. Since data required for each GPU kernel is processed only on GPUs, they can be readily found in the GPU's global memory. Therefore, we do not have to transfer any data between CPUs and GPUs during operation, removing one of the main performance bottlenecks. This allows us to recover the fast computation time on GPUs as shown in Fig. 6. We note that we also used CUDA-aware MPI for MPI implementations to support multiple GPUs so that we do not have to communicate via staging through CPUs.

3.3. Coalesced memory access

In Fig. 7(a) we show the memory layout of our original CPU-based code, where we use double precision for floating-point computation. In the figure, each coordinate value of $(\bar{x}, \bar{y}, \bar{z})$ denotes the maximum grid coordinate value of the corresponding dimension, and the superscript/subscript on v represents the index of a grid point/variable, respectively. The five variables for each grid point are located one after the other in memory. Because grid points are allocated in memory in lexicographic order and are accessed sequentially along with their five variables, memory accesses exhibit locality in this case, leading to fast memory operations through a high cache hit ratio.

However, this memory layout gives rise to 40-byte strided memory accesses on GPUs, as shown in Fig. 8(a), potentially incurring more memory transactions than needed. Unlike the CPU case, 32 threads in a warp are executed in lockstep on GPUs. Since each grid point has five variables of size 8 bytes each (double precision), a warp accessing the first variable of 32 grid points will result in 10 128-byte memory transactions footnoteThe granularity of L1 cache update is 128-byte. assuming that the accessed data have not been found in cache. Bus utilization (the ratio of data actually read to fetched) is 20% (= $256/1280 \times 100$) in this case.

² We used gprof available at https://sourceware.org/binutils/docs/gprof.



(a) Old memory layout (CPU-friendly).



(b) New memory layout (GPU-friendly).

Fig. 7. Memory layout: old vs. new.



(b) 8-byte strided memory access for the new layout

Fig. 8. Different strides for different memory layouts.



Fig. 9. Effect of different strides on memory bandwidth.

In contrast, if we allocate each variable of grid points one after the other as shown in Fig. 7(b), we will need only 2 128-byte memory transactions with 100% bus utilization (= $256/256 \times 100$), as shown in Fig. 8(b). Therefore, coalesced memory access provides more efficient memory transactions.

Although bus utilization is low in the former case, later accesses to other variables might be served via cache, resulting in similar overall memory access time to that of the latter case. However, this expectation depends on the lifetime of fetched data in cache. The amount of fetched data is proportional to the number of grid points. Since cache size is limited (128 kB for L1 cache on NVIDIA's Volta architecture), for larger numbers of grid points the likelihood of the eviction of data in cache becomes higher, reducing their lifetime in cache.

Fig. 9 demonstrates the effect of the aforementioned two strided accesses on memory bandwidth with respect to varying grid sizes. In the figure, the *x*-label represents the size of each dimension; for example, 16 represents 16 = |x| = |y| = |z|. Since we have the same number

of grid points along each dimension, the total number of grid points is $16 \times 16 \times 16$ in this case. The memory bandwidth was measured on the kernel to compute the weights (16). We observe that the speedup of 8-byte strided access with respect to 40-byte strided access becomes larger as we increase the grid size. If data were frequently found in cache for 40-byte strided access, the bandwidth ratio between those two access methods would be similar and stay constant. However, the speedup is amplified as we increase the grid size, suggesting that there have been many more wasted memory transactions for 40-byte strided access than those of 8-byte strided access.

4. Computational performance

In this section we present the computational performance of our GPU-based implementation on the isotropic turbulence decay example described in Section 2.3. In Section 4.2 we compare its computation time with that of the CPU-based implementation. We demonstrate the scalability of our implementation using multiple GPUs in Section 4.3.

4.1. Setting

We use the same parameters for the isotropic turbulence decay example as described in Section 2.3 except for time step, which we set to 0.002. Since we are interested mainly in the speedup of the computation time of our GPU-based implementation, we perform 100 time steps and measure the wall time for all experiments in this section.

The experiments in this section were performed on OLCF/Summit; each compute node has 6 NVIDIA Volta V100 GPUs and 2 POWER9 CPU sockets with 22 physical cores each and 4 hardware threads on each physical core.

4.2. One GPU vs. one CPU core

Fig. 10 presents the wall times of one GPU and one CPU core, respectively, for executing 100 time steps over varying grid sizes. From the results, we make two observations. The first observation is that up until 64^3 the grid size is not large enough to fully utilize the massive parallel computing capability of GPUs. Although the grid size becomes 8 times larger when we change its size from 32^3 to 64^3 , the computation time on GPUs increases by just 4.46x. This behavior is in contrast to the proportional increase of wall times on CPUs as we increase the grid size. The disproportional increase of the computation time on GPUs between grids with 32^3 and 64^3 points implies that there are many idle computation units when we solve over grids with 32^3 points or less.

Once we reach a grid with 64^3 points, however, the computational resources on GPUs start to saturate. This is verified by checking the computation time between grids of sizes 64^3 and 128^3 : the wall time increases exactly at the same rate as the increase of the grid size in this case. As shown in the figure, for grids having more than 64^3 points this proportionality is maintained, implying that we are extensively making use of GPUs. In the case where multiple GPUs are employed as in Section 4.3, our observation provides a guideline for determining



Fig. 10. One GPU vs. one CPU core for 100 time steps.



Fig. 11. Strong scaling over a grid with 256³ points.

the appropriate number of grid points to be assigned to each GPU that prevents underutilization.

We note that a grid with 208^3 points is the maximum size that HyPAR can run on one GPU because of the 32 GB memory limit on V100. For V100 with 16 GB of memory, the maximum size is around 160^3 . We also note that we did not measure the computation time on the CPU for the largest grid, since it is expected to be more than 5 h based on its consistent increase in computation time as the grid size increases.

Our second observation is that the computation time with multiple threads on one GPU is about 200 times faster than that on one CPU core. The speedup is largely attributed to the acceleration of the WENO routine: we achieved about 400 times faster computation time compared with that of the CPU-based code. Since it takes about 45% of the total computation time on CPUs as shown in Fig. 5, the improved computation time of the WENO routine significantly affects the overall acceleration on GPU. However, not all routines showed such large gains. For the Interp and Upwind routines, the gains were about a factor of 100 and 250, respectively. The reason is that they are computationally less involving than the WENO routine.

We note that the accelerations for smaller grids are not as large as those for larger grids as shown in Fig. 10, since they are not large enough to fully utilize GPUs, as we have seen in the preceding observation.

4.3. Multiple GPUs using MPI

We demonstrate scalability of HyPAR by evaluating strong and weak scaling over varying numbers of GPUs and grid points. Communication Computers and Fluids 251 (2023) 105744

Table	1
	

Grid	shape	per	GPU	for	decomposing	2563	grid

# GPUs	Grid shape per GPU
8	$128 \times 128 \times 128$
16	$64 \times 128 \times 128$
32	$64 \times 64 \times 128$
64	$64 \times 64 \times 64$
128	$32 \times 64 \times 64$
256	$32 \times 32 \times 64$
512	$32 \times 32 \times 32$



Fig. 12. Weak scaling over up to 64 GPUs and 640³ grid.



Fig. 13. Ratio of communication to computation cost.

between GPUs was performed via CUDA-aware MPI, which allowed us to directly communicate between memories of multiple GPUs on the same or different nodes without staging them through CPUs.

Fig. 11 presents our strong-scaling results. We fix the number of grid points to 256^3 and compute the speedup over up to 512 GPUs. Since a grid with 256^3 points cannot be loaded on a single GPU because of its memory limit, the *x*-axis starts from 8 GPUs so that each GPU is assigned to 128^3 grid points. Table 1 presents the grid shape per GPU we used for decomposing 256^3 grid for our strong-scaling experiments.

Table 2

Simulation resources used and wall times per time step (seconds) for GPU-enabled and CPU-only simulations of HIT decay on grids with 512^3 and 1024^3 points.

Grid	# nodes	GPU-enabled sime	ulation	MPI-only simulation		
		Number of GPUs	Wall time	Number of MPI ranks	Wall time	
512 ³	16	64 (4 × 4 × 4)	0.42	512 (8 × 8 × 8)	5.5	
	128	512 (8 × 8 × 8)	0.083	4096 (16 × 16 × 16)	0.72	
1024 ³	128	512 (8×8×8)	0.42	4096 (16 × 16 × 16)	5.5	
	256	1024 (16×8×8)	0.23	8192 (32 × 16 × 16)	2.8	

From the results, we observe that our GPU implementation scales well with the number of GPUs. However, we observe that the increase rate of the speedup diminishes as the number of GPUs increases. The reason for this decrease is that the MPI communication cost becomes a dominating factor to the overall cost.

The wall time for 8 GPUs was 24.62 s, and it took about 3.14 s for the MPI communication. In this case, the communication to computation percentage ratio was 12% (= $3.14/24.62 \times 100$). For 64 GPUs, however, the ratio became 45% (= $2.40/5.24 \times 100$), and it was 58% (= $1.11/1.90 \times 100$) for 512 GPUs. Since our implementation on GPUs is significantly faster than the CPU-based implementation, the impact of the communication cost quickly becomes substantial. We note that the diminishing rate seemed to be accelerated when we started to employ 32^3 grids (for 128–512 GPUs). We think this is partly because GPUs were not fully utilized for such small grids, as we discussed in Section 4.2.

For weak-scaling tests, we used three different grid schemes per GPU: $64 \times 64 \times 64$, $128 \times 128 \times 128$, and $160 \times 160 \times 160$. The size of the original grid increases proportionally to the number of GPUs: for 64 GPUs with $160 \times 160 \times 160$ ($128 \times 128 \times 128$ and $64 \times 64 \times 64$) grid per GPU, the shape of the original grid is $640 \times 640 \times 640$ ($512 \times 512 \times 512$ and $256 \times 256 \times 256$); hence, 4 GPUs are assigned to each dimension.

Fig. 12 presents our weak-scaling results in terms of efficiency. A similar trend of diminishing efficiency for larger numbers of GPUs has been observed, as with our strong-scaling results. The diminishing rate largely depends on the ratio of communication to computation cost. As we see in Fig. 13, the ratio is higher for 64³ grid, 43% for 64 GPUs, than those for 128³ and 160³ grids, 27% and 23% for 64 GPUs, respectively. Therefore, the 64³ grid showed a worse efficiency than those of the 128³ and 168³ grids. In Fig. 13, "Comp" denotes the pure computation time on each GPU. In all cases, the computation time on individual GPUs was almost the same.

5. HIT decay simulation

In this section we simulate the decay of isotropic turbulence on two large grids: 512^3 points (6.7×10^8 degrees of freedom) and 1024^3 (5.3×10^9) points. The simulation setup is the same as described in Section 2.3.³ The results presented here are obtained on LLNL/Lassen; each node has 4 NVIDIA Volta V100 GPUs and 2 IBM Power9 CPUs with 40 cores available for computations, where each core is capable of 4 hardware threads. The computational performances of GPU-enabled simulations and MPI-only simulations are compared by maximizing the parallelism on the same number of compute nodes. The 512^3 grid is simulated on 16 and 1,28 nodes; thus GPU-enabled simulations use 64 and 512 GPUs, respectively, and MPI-only simulations use 512 and 4096 MPI ranks, respectively. The 1024^3 grid is simulated on 128 and 2,56 nodes; GPU-enabled simulations use 4096 and 8192 MPI ranks, respectively, Even though each CPU core is capable of 4 hardware threads, the MPI-only simulations have only 1 on a core; we observed performance degradation with 2 or 4 MPI ranks per core.

Table 2 summarizes the resources used for these simulations and the observed wall times per time step for the 3-stage total-variationdiminishing (TVD) Runge–Kutta method in seconds (this does not include the time for file I/O). Overall, we see an excellent speedup for GPU-enabled simulations. While the 512³ grid simulation on 1,28 nodes is around 8 times faster with the GPU, all the other cases show that the GPU code is 12–13 times faster than the MPI code. We note that for GPU-enabled simulations, the CPU cores do not do any meaningful computations other than the initial transfer of the solution to the GPU and the transfer of the solution back to the CPU for writing to disk when required. The MPI-only simulations do not use the GPUs at all. Thus, both these approaches do not use the compute nodes to their maximum potential.

Fig. 14(a) shows the kinetic energy spectra at the final time $t/\tau = 3$ for solutions obtained on with 64³, 128³, 256³, 512³, and 1024³ points. The time step corresponds to a CFL of around 0.4 for all the simulations. We verified that the difference between the solutions obtained with the GPU-enabled code and the MPI code was within machine roundoff (10⁻¹⁵). The energy spectra show that the finer grids are able to resolve the higher wavenumbers (smaller length scales), as expected. Fig. 14(b) shows the time evolution of the energy spectrum for the solution on the 1024³-points grid; as the flow evolves, kinetic energy is transferred to the smaller length scales. This underscores the need for efficient scale-resolving simulations. Figs. 15(a) and 15(b) show the Schlieren and vorticity magnitude plots for the solution obtained on the grid with 1024³ points; the flow is dominated with small-scale structures as well as localized regions of very high gradients.

6. Conclusions

In this work we explore the development of HyPAR [54], a CPUbased C/C++ high-performance code, to run efficiently on GPU-based computing platforms. The naive use of the SIMD directives in an existing CPU codebase, by primarily targeting memory operations across the CPU–GPU bus, leads to severe performance degradation on GPUs because GPUs require a fundamentally different programming paradigm from that of CPUs. A bottom-up approach that focused on forming the equations directly on GPUs significantly reduces traffic between CPUs and GPUs, across the PCIe bus/NVLink, and improves the performance of the code significantly.

This work focuses on the three-dimensional compressible Navier– Stokes equations discretized by a nonlinear fifth-order WENO scheme in space and an explicit Runge–Kutta method in time. This combination of numerical discretization schemes allows us to exploit high potential parallelism and scalability while maintaining high numerical accuracy and spectral resolution. We demonstrate our results on a canonical turbulent flow problem, the decay of isotropic turbulence in a periodic domain characterized by the transfer of energy from larger to smaller length scales. We conduct DNS studies by resolving the Kolmogorov scale on a grid of size 1024^3 (~5 billion degrees of freedom).

The implementation of a high-order shock-capturing algorithm in this paper is among the earliest for simulating HIT in compressible flows on heterogeneous platforms. We carry out simulations on up to 1024 GPUs with CUDA-aware MPI communication across the nodes. Our results demonstrate that the optimizations proposed in this paper lead to more than 200x reduction in computation time using multiple threads on a GPU compared with a single CPU core. The scalability of the GPU-accelerated code is evaluated by using strong and weak scaling, and we investigate the impact of communication cost on the overall computation time. We compare our GPU-enabled and MPI-only implementations on DOE leadership-class supercomputers with multicore CPUs and GPUs, and we observe that the GPU-enabled code performs around 12–13 times faster on the these architectures. The strategies presented in this paper extend to conventional shock-capturing schemes,

³ This specific example is available in the HyPAR repository [54] in the subdirectory Examples/3D/NavierStokes3D/DNS_Isotropic TurbulenceDecay_CUDA.



(a) Energy spectrum for HIT decay at $t/\tau = 3$ for (b) Evolution of the energy spectrum on a grid with various grid sizes. 1024^3 points.





(a) Schlieren $(\|\nabla \rho\|_2)$ at $z = \pi$ and $t/\tau = 3$.

(b) Vorticity magnitude at z/=0 and $t/\tau = 3$.

Fig. 15. Isotropic turbulence decay: Solution in the x-y plane at z = 0 on a grid with 1024^3 points.

including low-dissipation, bandwidth-optimized WENO schemes that are better suited to DNS studies. In ongoing work, we are exploring nonlinear compact schemes, such as the CRWENO or hybrid compact-WENO schemes with higher spectral resolution. However, they require the direct solution of tridiagonal systems of equations; this presents additional challenges for GPUs.

CRediT authorship contribution statement

Youngdae Kim: GPU (CUDA) Implementation & code optimization, Numerical experiments, Writing – original draft. Debojyoti Ghosh: Conceptualization, Methodology, Numerical experiments and simulations, Code documentation & release. Emil M. Constantinescu: Numerical methods, Time integration, WENO methods, Supervision, Writing – review & editing. Ramesh Balakrishnan: Theory and CFD Methodology, Numerical experiment, Turbulent flow simulations, Supervision, 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 will be made available on request.

References

- Wallace JM. Space-time correlations in turbulent flow: A review. Theor Appl Mech Lett 2014;4(2):022003. http://dx.doi.org/10.1063/2.1402203.
- Kraichnan RH. The structure of isotropic turbulence at very high Reynolds numbers. J Fluid Mech 1959;5(4):497—543. http://dx.doi.org/10.1017/ S0022112059000362.

- [3] Kraichnan RH. Kolmogorov's hypotheses and Eulerian turbulence theory. Phys Fluids 1964;7(1723). http://dx.doi.org/10.1063/1.2746572.
- [4] Lesieur M, Ossia S. 3D isotropic turbulence at very high Reynolds numbers: EDQNM study. J Turbul 2000;1(1):007.
- [5] Buaria D, Pumir A, Bodenschatz E. Self-attenuation of extreme events in Navier–Stokes turbulence. Nature Commun 2020;11(1):5852.
- [6] Liu X-D, Osher S, Chan T. Weighted essentially non-oscillatory schemes. J Comput Phys 1994;115(1):200–12. http://dx.doi.org/10.1006/jcph.1994.1187.
- [7] Jiang G-S, Shu C-W. Efficient implementation of weighted ENO schemes. J Comput Phys 1996;126(1):202–28. http://dx.doi.org/10.1006/jcph.1996.0130.
- [8] Shu C-W, Osher S. Efficient implementation of essentially non-oscillatory shockcapturing schemes. J Comput Phys 1988;77(2):439–71. http://dx.doi.org/10. 1016/0021-9991(88)90177-5.
- Shu C-W, Osher S. Efficient implementation of essentially non-oscillatory shockcapturing schemes, II. J Comput Phys 1989;83(1):32–78. http://dx.doi.org/10. 1016/0021-9991(89)90222-2.
- [10] Balsara DS, Shu C-W. Monotonicity preserving weighted essentially nonoscillatory schemes with increasingly high order of accuracy. J Comput Phys 2000;160(2):405–52. http://dx.doi.org/10.1006/jcph.2000.6443.
- [11] Wu C, Wu L, Li H, Zhang S. Very high order WENO schemes using efficient smoothness indicators. J Comput Phys 2021;432:110158. http://dx.doi.org/10. 1016/j.jcp.2021.110158.
- [12] Henrick AK, Aslam TD, Powers JM. Mapped weighted essentially nonoscillatory schemes: Achieving optimal order near critical points. J Comput Phys 2005;207(2):542–67. http://dx.doi.org/10.1016/j.jcp.2005.01.023.
- [13] Borges R, Carmona M, Costa B, Don WS. An improved weighted essentially non-oscillatory scheme for hyperbolic conservation laws. J Comput Phys 2008;227(6):3191–211. http://dx.doi.org/10.1016/j.jcp.2007.11.038.
- [14] Yamaleev NK, Carpenter MH. Third-order Energy Stable WENO scheme. J Comput Phys 2009;228(8):3025–47. http://dx.doi.org/10.1016/j.jcp.2009.01. 011.
- [15] Yamaleev NK, Carpenter MH. A systematic methodology for constructing highorder energy stable WENO schemes. J Comput Phys 2009;228(11):4248–72. http://dx.doi.org/10.1016/j.jcp.2009.03.002.
- [16] Martin MP, Taylor EM, Wu M, Weirs VG. A bandwidth-optimized WENO scheme for the effective direct numerical simulation of compressible turbulence. J Comput Phys 2006;(220):270–289, 220 (2006).
- [17] Sjögreen B, Yee H. Accuracy consideration by DRP schemes for DNS and LES of compressible flow computations. Comput & Fluids 2017;159:123–36. http: //dx.doi.org/10.1016/j.compfluid.2017.09.017.
- [18] Sun Z-S, Ren Y-X, Larricq C, ying Zhang S, cheng Yang Y. A class of finite difference schemes with low dispersion and controllable dissipation for DNS of compressible turbulence. J Comput Phys 2011;230(12):4616–35. http://dx.doi. org/10.1016/j.jcp.2011.02.038.
- [19] Sun D, Qu F, Bai J, Yan C. An effective low dissipation method for compressible flows. Aerosp Sci Technol 2020;100:105757. http://dx.doi.org/10.1016/j.ast. 2020.105757.
- [20] Li Y, Chen C, Ren Y-X. A class of high-order finite difference schemes with minimized dispersion and adaptive dissipation for solving compressible flows. J Comput Phys 2022;448:110770. http://dx.doi.org/10.1016/j.jcp.2021.110770.
- [21] Deng X, hua Jiang Z, Vincent P, Xiao F, Yan C. A new paradigm of dissipationadjustable, multi-scale resolving schemes for compressible flows. J Comput Phys 2022;466:111287. http://dx.doi.org/10.1016/j.jcp.2022.111287.
- [22] Lele SK. Compact finite difference schemes with spectral-like resolution. J Comput Phys 1992;103(1):16–42. http://dx.doi.org/10.1016/0021-9991(92)90324-R.
- [23] Subramaniam A, Wong ML, Lele SK. A high-order weighted compact high resolution scheme with boundary closures for compressible turbulent flows with shocks. J Comput Phys 2019;397:108822. http://dx.doi.org/10.1016/j.jcp.2019. 07.021.
- [24] Chen R, Wu L, You Y. Development of a new type of weighted compact scheme. J Comput Phys 2021;446:110662. http://dx.doi.org/10.1016/j.jcp.2021.110662.
- [25] Hiejima T. A high-order weighted compact nonlinear scheme for compressible flows. Comput & Fluids 2022;232:105199. http://dx.doi.org/10.1016/j. compfluid.2021.105199.
- [26] Ghosh D, Baeder JD. Compact reconstruction schemes with weighted ENO limiting for hyperbolic conservation laws. SIAM J Sci Comput 2012;34(3):A1678–706. http://dx.doi.org/10.1137/110857659.
- [27] Ghosh D, Baeder JD. Weighted non-linear compact schemes for the direct numerical simulation of compressible, turbulent flows. J Sci Comput 2014;61(1):61–89. http://dx.doi.org/10.1007/s10915-014-9818-0.
- [28] Peng J, Shen Y. Improvement of weighted compact scheme with multi-step strategy for supersonic compressible flow. Comput & Fluids 2015;115:243–55. http://dx.doi.org/10.1016/j.compfluid.2015.04.012.
- [29] Peng J, Shen Y. A novel weighting switch function for uniformly highorder hybrid shock-capturing schemes. Internat J Numer Methods Fluids 2017;83(9):681–703. http://dx.doi.org/10.1002/fld.4285.
- [30] Fernández-Fidalgo J, Nogueira X, Ramírez L, Colominas I. An a posteriori, efficient, high-spectral resolution hybrid finite-difference method for compressible flows. Comput Methods Appl Mech Engrg 2018;335:91–127. http://dx.doi.org/10.1016/j.cma.2018.02.013.

- [31] Fu L, Hu XY, Adams NA. A family of high-order targeted ENO schemes for compressible-fluid simulations. J Comput Phys 2016;305:333–59. http://dx.doi. org/10.1016/j.jcp.2015.10.037.
- [32] Lusher DJ, Sandham ND. Assessment of low-dissipative shock-capturing schemes for the compressible taylor–green vortex. AIAA J 2021;59(2):533–45. http://dx. doi.org/10.2514/1.J059672.
- [33] Chakravarthy VK, Arora K, Chakraborty D. A simple hybrid finite volume solver for compressible turbulence. Internat J Numer Methods Fluids 2015;77(12):707–31. http://dx.doi.org/10.1002/fld.4000.
- [34] Adams NA, Shariff K. A high-resolution hybrid compact-ENO scheme for shockturbulence interaction problems. J Comput Phys 1996;127(1):27–51. http://dx. doi.org/10.1006/jcph.1996.0156.
- [35] Pirozzoli S. Conservative hybrid compact-WENO schemes for shock-turbulence interaction. J Comput Phys 2002;178(1):81–117. http://dx.doi.org/10.1006/ jcph.2002.7021.
- [36] Ren Y-X, Liu M, Zhang H. A characteristic-wise hybrid compact-WENO scheme for solving hyperbolic conservation laws. J Comput Phys 2003;192(2):365–86. http://dx.doi.org/10.1016/j.jcp.2003.07.006.
- [37] Jammy SP, Jacobs CT, Sandham ND. Performance evaluation of explicit finite difference algorithms with varying amounts of computational and memory intensity. J Comput Sci 2019;36:100565. http://dx.doi.org/10.1016/j.jocs.2016. 10.015.
- [38] Salvadore F, Bernardini M, Botti M. GPU accelerated flow solver for direct numerical simulation of turbulent flows. J Comput Phys 2013;235:129–42. http: //dx.doi.org/10.1016/j.jcp.2012.10.012.
- [39] Wang Y, Yan X, Zhang J. Research on GPU parallel algorithm for direct numerical solution of two-dimensional compressible flows. J Supercomput 2021;77:10921–41. http://dx.doi.org/10.1007/s11227-021-03704-9.
- [40] Esfahanian V, Darian HM, Gohari SI. Assessment of WENO schemes for numerical simulation of some hyperbolic equations using GPU. Comput & Fluids 2013;80:260–8. http://dx.doi.org/10.1016/j.compfluid.2012.02.031, Selected contributions of the 23rd International Conference on Parallel Fluid Dynamics ParCFD2011.
- [41] Darian HM, Esfahanian V. Assessment of WENO schemes for multidimensional Euler equations using GPU. Internat J Numer Methods Fluids 2014;76(12):961–81. http://dx.doi.org/10.1002/fld.3964.
- [42] Kuo FA, Wu JS. Implementation of a parallel high-order WENO-type Euler equation solver using a CUDA PTX paradigm. J Mech 2021;37:496–512. http: //dx.doi.org/10.1093/jom/ufab016.
- [43] Huang J, Han W, Wang X, Chen W. Heterogeneous parallel algorithm design and performance optimization for WENO on the Sunway Taihulight supercomputer. Tsinghua Sci Technol 2020;25(1):56–67. http://dx.doi.org/10.26599/TST.2018. 9010112.
- [44] Antoniou A, Karantasis K, Polychronopoulos E, Ekaterinaris J. Acceleration of a finite-difference WENO scheme for large-scale simulations on many-core architectures. In: 48th AIAA aerospace sciences meeting including the new horizons forum and aerospace exposition. http://dx.doi.org/10.2514/6.2010-525.
- [45] Karantasis KI, Polychronopoulos ED, Ekaterinaris JA. High order accurate simulation of compressible flows on GPU clusters over Software Distributed Shared Memory. Comput & Fluids 2014;93:18–29. http://dx.doi.org/10.1016/ j.compfluid.2014.01.005.
- [46] Post P, di Mare F. Towards direct numerical simulations of shock-turbulence interaction in real gas flows on GPUs: Initial validation. In: di Mare F, Spinelli A, Pini M, editors. Non-ideal compressible fluid dynamics for propulsion and power. Cham: Springer International Publishing; 2020, p. 59–75.
- [47] Bernardini M, Modesti D, Salvadore F, Pirozzoli S. STREAmS: A high-fidelity accelerated solver for direct numerical simulation of compressible turbulent flows. Comput Phys Comm 2021;263:107906. http://dx.doi.org/10.1016/j.cpc. 2021.107906.
- [48] Xu D, Luo S, Song J, Liu J, Cao W. Direct numerical simulations of supersonic compression-expansion slope with a multi–GPU parallel algorithm. Acta Astronaut 2021:179:20–32. http://dx.doi.org/10.1016/j.actaastro.2020.10.047.
- [49] Hamzehloo A, Lusher DJ, Laizet S, Sandham ND. On the performance of WENO/TENO schemes to resolve turbulence in DNS/LES of high-speed compressible flows. Internat J Numer Methods Fluids 2021;93(1):176–96. http://dx. doi.org/10.1002/fld.4879.
- [50] Lusher DJ, Jammy SP, Sandham ND. OpenSBLI: Automated code-generation for heterogeneous computing architectures applied to compressible fluid dynamics on structured grids. Comput Phys Comm 2021;267:108063. http://dx.doi.org/ 10.1016/j.cpc.2021.108063.
- [51] Di Renzo M, Fu L, Urzay J. HTR solver: An open-source exascale-oriented taskbased multi-GPU high-order code for hypersonic aerothermodynamics. Comput Phys Comm 2020;255:107262. http://dx.doi.org/10.1016/j.cpc.2020.107262.
- [52] Tutkun B, Edis FO. A GPU application for high-order compact finite difference scheme. Comput & Fluids 2012;55:29–35. http://dx.doi.org/10.1016/j.compfluid. 2011.10.016.
- [53] Xu C, Deng X, Zhang L, Fang J, Wang G, Jiang Y, Cao W, Che Y, Wang Y, Wang Z, Liu W, Cheng X. Collaborating CPU and GPU for large-scale high-order CFD simulations with complex grids on the TianHe-1A supercomputer. J Comput Phys 2014;278:275–97. http://dx.doi.org/10.1016/j.jcp.2014.08.024.

Computers and Fluids 251 (2023) 105744

- [54] HyPar Repository, https://bitbucket.org/deboghosh/hypar.
- [55] Zolfaghari H, Becsek B, Nestola MG, Sawyer WB, Krause R, Obrist D. Highorder accurate simulation of incompressible turbulent flows on many parallel GPUs of a hybrid-node supercomputer. Comput Phys Comm 2019;244:132–42. http://dx.doi.org/10.1016/j.cpc.2019.06.012.
- [56] Zolfaghari H, Obrist D. A high-throughput hybrid task and data parallel Poisson solver for large-scale simulations of incompressible turbulent flows on distributed GPUs. J Comput Phys 2021;437:110329. http://dx.doi.org/10.1016/j.jcp.2021. 110329.
- [57] Ravikumar K, Appelhans D, Yeung PK. GPU acceleration of extreme scale pseudospectral simulations of turbulence using asynchronism. In: Proceedings of the international conference for high performance computing, networking, storage and analysis. SC '19, New York, NY, USA: Association for Computing Machinery; 2019, http://dx.doi.org/10.1145/3295500.3356209.
- [58] Oak Ridge Leadership Computing Facility Summit, https://www.olcf.ornl.gov/ olcf-resources/compute-systems/summit/.
- [59] Livermore Computing Lassen, https://hpc.llnl.gov/hardware/computeplatforms/lassen.

- [60] Hirsch C. Numerical computation of internal and external flows: The fundamentals of computational fluid dynamics: the fundamentals of computational fluid dynamics, vol. 1 & 2. Elsevier Science; 2007.
- [61] Ghosh D, Constantinescu EM, Brown J. Efficient implementation of nonlinear compact schemes on massively parallel platforms. SIAM J Sci Comput 2015;37(3):C354–83. http://dx.doi.org/10.1137/140989261.
- [62] Roe PL. Approximate Riemann solvers, parameter vectors, and difference schemes. J Comput Phys 1981;43(2):357–72. http://dx.doi.org/10.1016/0021-9991(81)90128-5.
- [63] Harten A. High resolution schemes for hyperbolic conservation laws. J Comput Phys 1983;49(3):357–93. http://dx.doi.org/10.1016/0021-9991(83)90136-5.
- [64] Rogallo RS. Numerical experiments in homogeneous turbulence. Tech. Rep. NASA-TM-81315, Moffett Field, CA: NASA Ames Research Center; 1981.
- [65] Mansour NN, Wray AA. Decay of isotropic turbulence at low Reynolds number. Phys Fluids (1994-Present) 1994;6(2):808–14. http://dx.doi.org/10.1063/1. 868319.
- [66] Nvidia. Nvidia tesla v100 GPU architecture. Tech. rep., 2017.