On the use of a multigrid-reduction-in-time algorithm for multiscale convergence of turbulence simulations

Stephen M. Guzik[a,*], Joshua Christopher[a], Sean Walters[a], Xinfeng Gao[a], Jacob B. Schroder[b], Robert D. Falgout[c,1]

[a] *Computational Fluid Dynamics and Propulsion Laboratory, Colorado State University, Fort Collins, USA*
[b] *Department of Mathematics and Statistics, University of New Mexico, Albuquerque, USA*
[c] *Center for Applied Scientific Computing, Lawrence Livermore National Laboratory, Livermore, USA*

Abstract

Simulations of turbulent flow present challenges in terms of accuracy and affordability on modern highly-parallel computer architectures. A multigrid-reduction-in-time algorithm is used to provide a framework for separately evolving different scales of turbulence and for parallelizing the temporal domain, thereby increasing the concurrency. It is hypothesized that the space-time locality of the small scales of turbulence can be used to circumvent difficulties in applying temporal multigrid to flows dominated by inertial physics. For algorithms that fall well short of spectral accuracy (fourth-order is used in this work) attention must be paid to the accuracy of features on scales transferred between multigrid levels. Numerical experiments were performed using implicit large-eddy simulation. Results from applying the approach to an infinite-Reynolds number Taylor-Green flow and a double-shear flow at a Reynolds number of 11650 provide strong evidence that the approach has merit. The multigrid-reduction-in-time framework can be used to parallelize the temporal domain of a high-Reynolds-number turbulent flow and permit independent convergence of different scales. Establishing this foundation allows for future research in reducing the wall-clock time to solve turbulent flows while retaining the same accuracy as sequential solvers. Current performance results from parallelizing the temporal domain are not competitive with those from sequential-in-time methods.

Keywords: Turbulence, Parallel in Time, Multigrid Reduction in Time

1. Introduction

In this study, an attempt is made to solve a turbulent flow with a parallel-in-time (PinT) algorithm. The motivating hypotheses are that, by organizing simulations of turbulent phenomena in a multigrid-reduction-in-time (MGRIT) algorithmic foundation [1], one can exploit the multiple scales of turbulence and space-time locality of finer scales to a) parallelize the temporal domain for a strongly inertial flow, and b) separately address and converge different scales of the problem. The main reason for parallelizing the temporal domain is to exploit the vast concurrency of current and future architectures, especially those with heterogeneous processing units. This reduces the wall-clock time to achieve the same answer as a sequential solve, which is the primary goal. However, the MGRIT approach provides a natural separation of resolved scales which can be useful for several purposes such as using adaptive algorithms. In addition, the ability to revisit temporal regions is a powerful capability providing the opportunity to improve a calculation that might be deemed inadequate. The impediment to using MGRIT, for which optimal theory is only well developed for elliptic and parabolic problems, is hyperbolic transport. Unfortunately, this is the nature of the inertial terms that dominate most engineering computations, including that of turbulence. Slow convergence of hyperbolic physics translates into long wall-clock times that may exceed sequential calculation, obviating the point of using MGRIT. The novelty of this research is to circumvent this impediment by appealing to

*Corresponding author
Email address: Stephen.Guzik@colostate.edu (Stephen M. Guzik)

[1] This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344 (LLNL-JRNL-827388).

the physics of turbulence. If the smaller scales of turbulence are localized in space-time, it is only necessary to calculate a localized time-history of these flow scales, allowing decomposition of the temporal domain into independent regions. The strategies espoused herein depend on the energy cascade of a turbulent flow and are not applicable to hyperbolic problems in general.

Turbulence arises when a fluid is unstable, vorticity is present, and the fluid is perturbed. The result is a chaotic distribution of velocity fluctuations describing multiple scales of vortices. Perhaps the best quantification is the energy cascade, where through a Fourier transform, one can observe the kinetic energy, E, associated with various eddy scales represented by the wavenumber, k. A representative cascade is shown in Fig. 1. In an incompressible flow, the net energy transfer is from large to small scales. Where the transfer is dominated by inertial vortical physics, it is referred to as the inertial cascade and dimensional reasoning suggests that a straight line with a slope of $-5/3$ should be observed when E is plotted versus k using logarithmic axes. At the finest scale, known as the Kolmogorov scale, the kinetic energy is finally dissipated into heat by viscosity. Turbulence is one of the most dominant and challenging problems in fluid physics: enhancing mixing in combustion devices, increasing friction drag while perhaps inhibiting separation on airfoils, or disrupting confinement in Tokamak reactors. Turbulence can be predicted with the Navier-Stokes equations as long as the mesh is finer than the Kolmogorov scale. For many applications, large inertial forces drive the Kolmogorov scale to so small a value that direct simulation of the full problem domain (e.g., a full engine combustor) is impossible on any foreseeable computer. In the meantime, the most promising technology for complex flows is large-eddy simulation (LES), where scales are separated by a filter: larger scales are solved on an appropriate grid while smaller scales are modeled. Fig. 1 illustrates the filtering concept with the grid at $k = \pi/\Delta\vec{x}$. Large scales, $\bar{u}$, are represented on the grid while small scales, u', must be modeled. How fine the mesh needs to be depends on the physics and specific problem but most engineers and scientists will still find themselves constrained by computational resources.

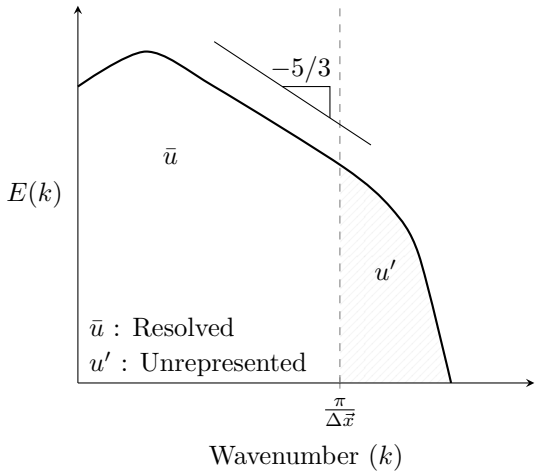

Figure 1: Spectrum of a representative turbulent energy cascade.

LES works well when rate-limiting processes happen in the resolved scales [2], such as occurs in solutions of the incompressible Navier-Stokes equations. There, well-validated subgrid models have also been devised such as the stretched-vortex subgrid-scale model developed by Pullin and collaborators [3, 4]. But add chemical reactions, and suddenly the rate-limiting processes are at the smallest scales. For these classes of problems, devising effective subgrid scale models is very challenging and it is worthwhile to consider even the validity of LES. An ability to selectively choose which scales need to be modeled or an ability to adaptively switch between modeling or resolving various scales can be of immense value. This study considers using a multigrid-reduction-in-time (MGRIT) algorithm for multiscale convergence of turbulent scales and assesses the validity of the approach for a relatively simple unbounded incompressible flow. Long-term goals involve adaptively resolving the challenging small-scale dynamics of more complex physics. Ultimately, an effective algorithm would model turbulence when the dynamics can be well predicted, and otherwise resolve the turbulence using the MGRIT algorithm in regions of space-time where there is low-confidence in the turbulence model.

Multigrid is an efficient algorithm with a theoretical O(N) convergence that is often nearly realized for relaxation problems or problems with physics governed by elliptic PDEs. The classical multigrid approach systematically employs sets of coarser grids to accelerate the convergence of iterative schemes by transferring the difficult low-frequency error modes from finer grids to high-frequency on coarser grids. The error corrections are solved on the coarser grids, where the problem size becomes much smaller, and then prolonged back to the finest mesh for solution updates. After a few cycles, the problem hopefully converges to the answer. In the context of MGRIT, the cycling between grid resolutions allows for parallelization of the temporal domain while recovering, upon convergence, the same answer as a sequentially-stepped algorithm (see Fig. 2). Although more work is done than in sequential time stepping, the additional parallelization permits use of additional processing units and can significantly reduce overall wall-clock time to solution. The extra work

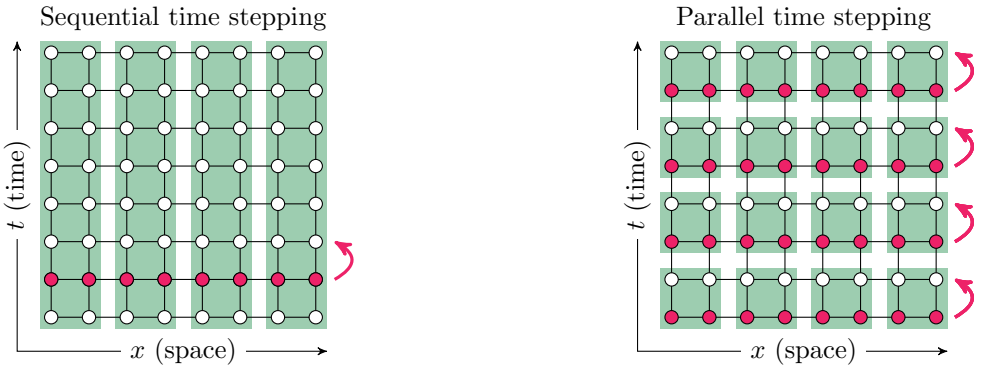

(a) Sequential time stepping.

(b) Parallel time stepping.

Figure 2: Parallel-in-time algorithms decompose and parallelize the temporal direction, increasing available data concurrency. Green boxes indicate processors and red dots indicate storage requirements in memory. While storage requirements increase, the storage per processor remains similar.

can be mitigated by reducing the computational cost of coarse-grid solutions (e.g., use an implicit method for large time-steps on the coarse grid). For some problems, such as time periodic (e.g., piston engine) or turbulent flows that have a steady-state when time averaged, MGRIT can more rapidly eradicate starting transients and easily outperform sequential time stepping. A major obstacle to the MGRIT approach is an inability to efficiently solve physics described by hyperbolic PDEs. Unfortunately, this is the nature of the inertial terms which dominate turbulent flows. For steady-state inertial problems solved with multigrid, high frequency errors are damped while low frequency errors are convected through the outer boundary [5]. But for a temporal domain, convecting error to the end time is essentially the same as sequential time-stepping. The issue here is that information flows in a wave-like manner along characteristics. The full time-history of the characteristic propagation must be solved for time-accuracy. Research is ongoing in developing new algorithms that are more amenable to solving hyperbolic PDEs using MGRIT [6, 7] or other techniques [8]. But the problem is daunting for three-dimensional compressible flows. At each point, the characteristics of the acoustic modes are spheres that expand in time and are continuously intersecting, modifying both the wavespeed and state information.

The main hypothesis of this research is that the difficulty of propagating characteristics can be circumvented by appealing to the physics of turbulence, recovering a parallel-in-time capability. Simply stated, there is no need to accurately solve the complete time-history of small turbulent scales. Rather, only the influence of fine scales on large scales is required and this influence is localized in space-time. In LES, many successful subgrid models, such as the stretched-vortex model, are in equilibrium; they provide an instantaneous snapshot of the effect of the modeled scales given an instantaneous state of resolved scales. With scale separation realized on multiple grids, one should be able to break up the temporal domain on the finest grid into localized domains that can be solved in parallel. The same goals persist as for MGRIT applied to laminar flows, in particular temporal parallelization. An accurate depiction of the coarse scales may be all that is desired and iteration or evolution of fine-scales provides their influence. It is hypothesized that the MGRIT formalism is the proper way to approach the problem, with the solution on the finest grid and error conceptually propagated on coarser grids. *A secondary hypothesis is that the MGRIT algorithm can provide a means for separately converging different scales of the problem.* Notably, on the coarsest grid level, a full sequential solve is performed in the temporal domain meaning that the time-history of coarse-scale characteristics is accurately computed. Calculations on the coarser levels can be orders of magnitude less expensive than iteration on the finest grid. Put another way, error in low-frequency content is removed on the coarse grid by convecting to the end time boundary, error in high-frequency content is removed by iterating small-scale turbulent dynamics to equilibrium.

The idea of solving a turbulent flow with a parallel-in-time algorithm has been explored previously, primarily using the Parareal algorithm [9, 10, 11]. The ability to solve turbulent flows is described as remarkable and effective for both drift wave turbulence and homogeneous isotropic turbulence. While there

are specific two-level equivalences with MGRIT[2], a notable difference is in the method design philosophy for multigrid where the goal is to solve and eradicate the error of the fine mesh on a sequence of coarser grids. In MGRIT, the coarse grid solves are recursively connected to the fine grid solves through a tau correction term. In the context of turbulence simulations, the coarsest grid solve is advecting coarse scale errors out of the temporal domain, while the tau correction term is adding the influence of fine scales. In that respect, one can think of the tau correction term the same as modeled subgrid-scale terms in a large-eddy simulation. Except in this case, the subgrid-scale terms are actually resolved on a finer mesh and solved in parallel. The present research also articulates that fine scales are evolved, at time scales much faster than that of coarse-scale dynamics, to an equilibrium versus being simply dissipated. The resulting localization in space-time permits temporal parallelization. Thus in summary, we explore our two hypotheses through the use of new modifications to MGRIT, i.e., (i) use of internal time-steps inside individual MGRIT time-steps, which allows us to separately converge different scales of the turbulence on the fine and coarse time-grids as discussed above, (ii) spectral filtering during spatial prolongation to respect scale separation, (iii) deconvolution during spatial interpolation to initialize small scales, and (iv) application of energy spectrum plots to study and verify MGRIT convergence. We note that (iv) is a common strategy adopted from the turbulence modeling community.

In the next section, the numerical apparatus used to test the hypotheses is described. This includes the mathematical model, the numerical computational fluid dynamics (CFD) algorithm, and the MGRIT algorithm. In section 3, the two test cases are defined. In section 4, strategies for solving turbulent flows with an MGRIT algorithm are explored. These strategies are used in section 5 to solve both test cases using a three-level MGRIT algorithm. The following section presents results and discussion, noting particular details and strategies before demonstrating a full simulation and assessing performance. The final section presents the conclusions.

In summary, the goal of this research is to demonstrate that the MGRIT algorithm can reproduce the sequential solution of a high Reynolds number turbulent flow, this validating the hypotheses. This research does not seek to find a more accurate algorithm for performing LES. Hence, a simplistic implicit LES (ILES) approach is used where the turbulent energy dissipation is provided by the numerical dissipation inherent to the algorithm. While crude, it is sufficient for investigating the hypotheses. The effect of applying a more advanced subgrid-scale model to the test cases considered here can be found in Walters et al. [14] (using sequential time stepping).

Performance of the MGRIT algorithm is not a focus of this article. Nonetheless, for the results generated in section 5, it was observed that the performance of the PinT algorithm compared with or slightly exceeded the performance of sequential time-stepping, while using significantly more hardware resources. To be a viable algorithm, at least an order of magnitude reduction in wall-clock time is desired. Thus, the performance of the PinT algorithm was not considered to be competitive with that of sequential time stepping. This is left to future research.

2. Method

2.1. Mathematical Model

The Navier-Stokes equations, describing conservation of mass (ρ), momentum ($\rho\vec{u}$), and energy (ρe_T), are sufficient for modeling turbulent flow,

$$\frac{\partial \rho}{\partial t} + \vec{\nabla} \cdot (\rho \vec{u}) = 0 \,, \tag{1}$$

$$\frac{\partial (\rho \vec{u})}{\partial t} + \vec{\nabla} \cdot \left(\rho \vec{u}\vec{u}^T + \mathbf{I}p\right) = \vec{\nabla} \cdot \boldsymbol{\tau} \,, \tag{2}$$

$$\frac{\partial (\rho e_T)}{\partial t} + \vec{\nabla} \cdot (\rho \vec{u} h) = \vec{\nabla} \cdot \left(\boldsymbol{\tau} \cdot \tilde{\vec{u}} - \vec{q}\right) \,, \tag{3}$$

[2]See [12] where two-level MGRIT with F-relaxation is shown to be equivalent with Parareal and see [13] where an overlapping Schwarz relaxation is shown to be equivalent to the standard MGRIT FCF-relaxation.

where the total energy $e_T = e + |\vec{u}|^2/2$, the enthalpy $h = e + p/\rho$, and, assuming a calorically perfect gas, the internal energy $e = c_v T$. The stress tensor is given by

$$\boldsymbol{\tau} = 2\mu \left(\mathbf{S} - \frac{1}{3}(\vec{\nabla} \cdot \vec{u})\mathbf{I} \right) , \quad \mathbf{S} = \frac{1}{2} \left(\vec{\nabla}\vec{u} + (\vec{\nabla}\vec{u})^T \right) , \tag{4}$$

where μ is the dynamic viscosity of the fluid. The heat flux, denoted by $\vec{q}$, is approximated by Fourier's law,

$$\vec{q} = -\kappa \vec{\nabla} T , \tag{5}$$

where κ is the thermal conductivity of the fluid. Finally, the ideal gas law, $p = \rho R T$, is used to close the equations with R being the specific gas constant.

If the mesh spacing, $\Delta\vec{x}$, can resolve the Kolmogorov scale, then turbulence is accurately simulated, and the approach is known as direct numerical simulation. For most high-Reynolds number flows, it is impractical to use a mesh so fine because of the computational expense. Instead, equations (1) to (3) can be filtered to only retain low wavenumber content. In an approach known as large-eddy simulation, the filtered scales are solved using discrete approximations to equations (1) to (3). The smaller scales must be modeled. When the equations are filtered using a Favre-averaged filter, their form remains the same except additional terms appear. For example, in the momentum equation, an additional sub-grid scale (SGS) stress appears as

$$\boldsymbol{\tau}_{\text{SGS}} = \bar{\rho} \left(\widetilde{\vec{u}\vec{u}^T} - \tilde{\vec{u}}\tilde{\vec{u}}^T \right) , \tag{6}$$

where the tilde denotes a Favre-filtered quantity ($\tilde{\vec{u}} = \overline{\rho\vec{u}}/\bar{\rho}$) and the over-bar indicates a filtered or represented scale. The first term on the right-hand-side of (6) is a filtering of the nonlinear product of $\vec{u}$ and is not known, it must be modeled (whereas $\tilde{\vec{u}}$ replaces $\vec{u}$ in the filtered equations and is solved for as part of the system state). Additional unknown terms appear in the energy equation and in state relationships between pressure and energy. Terms in the latter can be neglected for the low Mach number (nearly incompressible) flows studied herein. A description of all terms and their modeling with a stretched-vortex subgrid-scale model are described in [14] (see also [3, 15, 16] for more information on this model). In the approach of implicit LES (ILES), the additional terms are neglected and instead, characteristics of the numerical method are relied upon to produce subgrid-scale dynamics. The approach of ILES is used in this work (note that it is the authors' opinion that explicit LES is usually superior; however the simplicity afforded by ILES is taken advantage of for this study).

In one of the benchmarks that will be studied, physical dissipation is eliminated by setting $\mu = \kappa = 0$. Equations (1) to (3) are then the Euler equations which are dominated by inertial transport. While equations (1) to (3) describe the physics of compressible flow, they can be accurately applied to the low Mach number (nearly incompressible) flows considered in the test cases. The main consequence is that stably resolving negligible acoustic phenomena can impose severe restrictions on explicit time steps. Since performance aspects are not considered in this paper, this restriction is largely unimportant. If performance were a concern, better approaches include using implicit time-stepping to circumvent the time-step restrictions or using algorithms designed to solve fully incompressible equations. For test cases that are truly compressible, the hypothesis that small scales can be iterated to equilibrium would have to be separately verified, taking into account interactions between acoustics and vortices.

2.2. Computational Fluid Dynamics Algorithm

The algorithm for evolving the fluid dynamics is a fourth-order, in space and time, finite volume method implemented in a code named Chord [17, 18]. The problem domain is divided into control volumes by a structured grid and equations (1) to (3) are written in vector form and integrated over a control volume,

$$\frac{\partial}{\partial t} \int_V \mathbf{U} \, \mathrm{d}V + \int_V \vec{\nabla} \cdot \vec{\mathbf{F}} \mathrm{d}V = \int_V \vec{\nabla} \cdot \vec{\mathbf{G}} \mathrm{d}V \tag{7}$$

where $\mathbf{U} = [\rho, \rho\vec{u}, \rho e_T]^T$ is the state vector, $\vec{\mathbf{F}}$ is the inviscid flux dyad, and $\vec{\mathbf{G}}$ is the viscous flux dyad. Details of the algorithm are provided in [19] and [17]; a high-level overview follows. In cells, the primitive solution state, $\langle \mathbf{W} \rangle$, is computed from the conservative solution state, $\langle \mathbf{U} \rangle$, where the angle brackets denote

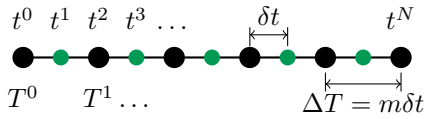

Figure 3: A two-level time grid composed of F-points in green and C-points in black. The composition of C-points and F-points form the fine time grid while, the C-points form the coarse time grid.

a quantity that is averaged over a cell. A fourth-order estimate of the primitive solution state is interpolated to faces. Next left and right values on a face are adjusted using a high-order piecewise parabolic method (PPM) limiter. On all faces, the primitive state is again made single-valued by solving a Riemann problem. The flux is computed on the faces and its divergence provides the conservative cell update. The solution is advanced in time using a standard fourth-order explicit Runge-Kutta scheme. Overall, the algorithm considers fourth-order-accurate cell- and face-averaged quantities. Non-linear calculations are performed at points (cell centers or face centers), for computing, e.g., the primitive state from the conservative state in cells or the flux from the primitive state on faces. Where necessary, first the solution is deconvolved from an averaged state, $\langle \mathbf{U} \rangle$, to a point state $\mathbf{U}$. Next a calculation is performed, e.g., the primitive state $\mathbf{W} = \mathbf{W}(\mathbf{U})$, and then the point state, $\mathbf{W}$ is convolved to obtain the average state $\langle \mathbf{W} \rangle$. A major feature of the algorithm is adaptive mesh refinement (AMR). In this work, AMR is not demonstrated, but use is made of the infrastructure for inter-level operations in the multigrid algorithm.

Where the solution state is smooth, the algorithm behaves as a fourth-order low-dissipation centered scheme where stability is primarily realized by properties of the Runge-Kutta update. Where the solution state is not smooth, the PPM limiter reduces the order of accuracy and the Riemann solve introduces upwinding, increasing dissipation in these regions. For any unresolved turbulent simulation (i.e., any large-eddy simulation), vortical structures on the scale of the grid appear as discontinuous features to the limiter, causing increased dissipation. The action of the limiter corrupts the flow features on the grid scale. However, it is not a viable option to simply turn off the limiter as some mechanism is needed to dissipate the high-frequency energy, either a SGS turbulence model (explicit LES) or numerical stability mechanism such as using the PPM limiter (implicit LES). Experience from working with the above algorithm suggests that solution features at the scale of $2\Delta\vec{x}$ are borderline accurate (roughly meaning that one may be able to use them but should test carefully) and features with a scale $\geq 4\Delta\vec{x}$ are well represented. In [14], a structural stretched-vortex SGS model is used and it was shown that it was necessary to fit the model to features on the order of at least $2\Delta\vec{x}$ to obtain reasonable results. In fact, turbulence simulations were shown to agree independent of additional numerical regularization at the grid scale, either PPM or hyper-viscosity, as long as the turbulence model was structurally fitted at a coarser length scale. In this work, coarse grid information needs to be prolonged to a finer grid in the multigrid procedure. It will be shown that similar treatment is necessary to avoid prolonging largely corrupt features at the grid-scale.

2.3. Multigrid Reduction in Time Algorithm

The main concepts behind the MGRIT algorithm are presented here and reproduced from [20], where the algorithm is presented in extensive detail. The MGRIT implementation used for this work is provided by the XBraid software [21]. In an MGRIT algorithm, the mesh is defined in both space and time (see Fig. 2) and the solution state is stored in both space and time. For one-step integration methods such as fourth-order Runge-Kutta, the time discretization method can be represented as

$$\mathbf{U}_0 = g_0, \quad \mathbf{U}_i = \Phi_i(\mathbf{U}_{i-1}) + \boldsymbol{g}_i, \quad i = 1, 2, \ldots, N_t\,, \tag{8}$$

where $\boldsymbol{g}_i$ are solution-independent terms. Assembling the time-points into matrix form for a linear case results in,

$$\mathrm{A}\mathbf{U} \equiv \begin{pmatrix} \mathrm{I} & & & \\ -\Phi_1 & \mathrm{I} & & \\ & \ddots & \ddots & \\ & & -\Phi_{N_t} & \mathrm{I} \end{pmatrix} \begin{pmatrix} \mathbf{U}_0 \\ \mathbf{U}_1 \\ \vdots \\ \mathbf{U}_{N_t} \end{pmatrix} = \begin{pmatrix} \boldsymbol{g}_0 \\ \boldsymbol{g}_1 \\ \vdots \\ \boldsymbol{g}_{N_t} \end{pmatrix} \equiv \boldsymbol{g}\,. \tag{9}$$

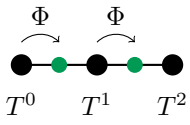

(a) F-relaxation to F-points.

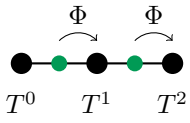

(b) C-relaxation to C-points.

Figure 4: Parallel application of Φ, the time integration operator, to F and C points on a fine multigrid level. Note that each application of Φ (C-relaxation or F-relaxation) integrates forward in time by δt, see Figure 3.

A forward block solve of this system corresponds to sequential time integration. MGRIT iteratively solves (9) with multigrid.

In the MGRIT algorithm, the temporal mesh of Fig. 2 is coarsened into a hierarchy of temporal grids using a coarsening factor. Either two- or three-level space-time grids with a coarsening factor of two are used for all results presented herein. Each grid is partitioned into F-points, which exist only on the finer time grids, and C-points which co-exist on both the finer and coarser time grids. A two-level hierarchy is shown in Fig. 3 where the fine grid is the composite of F-points (in green) and C-points (in black), and the coarse grid is just the C-points. For time-stepping methods such as explicit Runge-Kutta, the spatial mesh is coarsened at the same ratio as the temporal mesh to preserve the CFL number.

Relaxation, prolongation, and restriction operators, are described next. Relaxation is performed by application of the time integration operator, Φ, to the time points in two stages. The first stage is F-relaxation which applies the time integration operator to each block of F-points in parallel. F-relaxation is shown in Fig. 4a which concurrently performs the applications of Φ to update the F-points. The second stage is C-relaxation in Fig. 4b, which propagates the solution to the C-points in parallel. Successive applications of F-relaxation and C-relaxation, called FCF-relaxation, make up the relaxation strategy for MGRIT on a given temporal level. Note that changing the multigrid coarsening factor, m, to 4 would insert 3 F-points between the C-points and only affect F-relaxation [1].

The coarse-grid system of equations, defined at mesh C-points $i = jm$, $j = 0, 1, \ldots, N_t/m$, is

$$\mathrm{A}_\Delta \boldsymbol{e}_\Delta \equiv \begin{pmatrix} \mathrm{I} & & & \\ -\Phi_{\Delta,1} & \mathrm{I} & & \\ & \ddots & \ddots & \\ & & -\Phi_{\Delta,N_t/m} & \mathrm{I} \end{pmatrix} \begin{pmatrix} \mathbf{e}_{\Delta,0} \\ \mathbf{e}_{\Delta,1} \\ \vdots \\ \mathbf{e}_{\Delta,N_t/m} \end{pmatrix} \equiv \boldsymbol{r}_\Delta, \tag{10}$$

with the coarse-grid error approximation $\mathbf{e}_{\Delta,i}$, the residual $\boldsymbol{r}_\Delta = \mathbf{R_T}\left(\boldsymbol{g} - \mathrm{A}\left(\mathbf{U}\right)\right)$, and $\mathbf{R_T}$ the temporal restriction operator. Temporal restriction is performed in the standard MGRIT way with injection at the C-points [1] and, in this study, is accompanied by spatial coarsening, with $\mathbf{R_S}$ the spatial coarsening operator. Since the solution state is already represented by averages in the CFD algorithm, spatial coarsening is an exact operation. After restriction, the coarse grid residual equation $\mathrm{A}_\Delta \mathbf{e}_\Delta = \boldsymbol{r}_\Delta$ is solved on the coarse grid sequentially using the coarse grid time propagator $\Phi_{\Delta,j}$.

Complementary to restriction, the coarse grid error approximation is first spatially interpolated to the fine spatial resolution, denoted with $\hat{\boldsymbol{e}}_\Delta = \mathbf{P_S}(\boldsymbol{e}_\Delta)$. Spatial interpolation is performed with a fourth-order-accurate least squares method [19]. Next, interpolation in time, $\mathbf{P_T}$, injects the coarse grid error correction to the C-points on the fine grid, and the solution is updated with $\mathbf{U} = \mathbf{U} + \mathbf{P_T}(\hat{\boldsymbol{e}}_\Delta)$. This completes a two-grid MGRIT cycle of relaxation, coarse grid solve for $\boldsymbol{e}_\Delta$, followed by the error correction on the fine grid. This can be applied recursively for additional levels. For nonlinear problems, the full approximation scheme (FAS) technique, also known as full approximation storage, is used [22].

The multigrid solution process with AMR is shown in Fig. 5 for three levels. An FMG cycle is used for initialization, starting with a time-sequential solve of the coarsest space-time level. During this sequential solve, regions of the space-time mesh are tagged for refinement. A new space-time level is created through

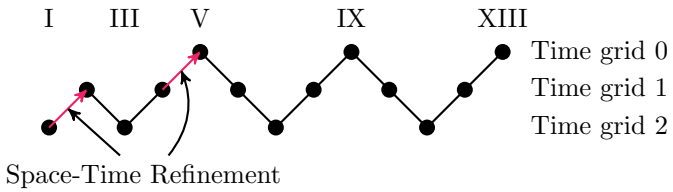

(a) Finer multigrid time grids are created by refinement in an FMG cycle

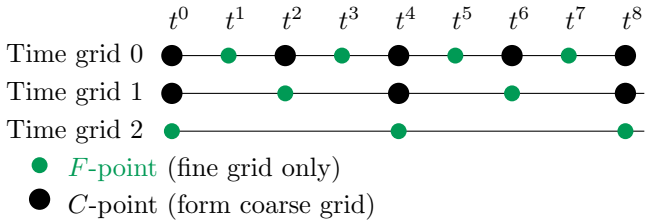

(b) After each temporal refinement, the coarser time grids are created by coarsening from the finest time grid with a fixed coarsening factor

Figure 5: Finer time grids are created in an FMG cycle.

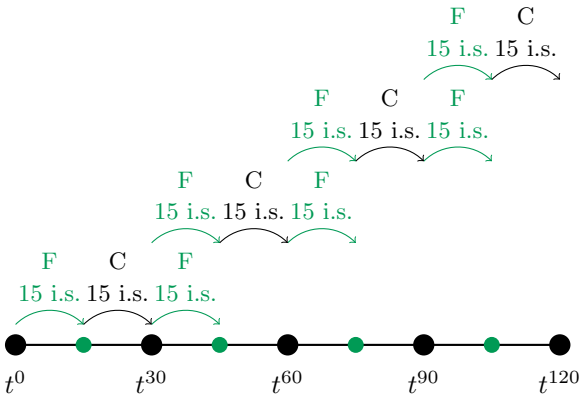

Figure 6: The use of internal steps (i.s.) in FCF-relaxation for turbulent flows. The time point numbers are representative of the number of steps taken by the explicit time-integration method. The vertical separation indicates operations that are performed in parallel.

refinement, as shown in Fig. 5a, and a two-level multigrid scheme is formed. Additional multigrid levels may be created by recursive applications of this FMG process. The time grid for three levels is shown in Fig. 5b, with the F-points of each level in green and the C-points in black. A temporal coarsening factor of two is used for each level in the time grid. As the time grid is coarsened, the spatial mesh is also coarsened by removing the finest AMR level. Due to the CFL condition, the coarsening ratio between the spatial meshes must match the coarsening ratio of the temporal grids. In this work, the spatial grids on each level cover the full domain, but this need not be the case. The finer grids can be localized around spatial features using adaptive-mesh refinement in a manner that couples with the MGRIT temporal grids [20].

For solving turbulent flows, some modifications to the classical MGRIT algorithm described above are necessary. Flexibility is required for tuning the amount of relaxation applied to fine turbulent scales with respect to error corrections from the coarse grid. If the fine time integrator is only applied twice between C-points, solutions on the coarse grid would have negligible change. Moreover, the coarse grid solves are sequential and, in that sense, expensive. To provide flexibility—and address the second hypothesis of separately converging scales—the multigrid F- and C-points are stretched further apart in the temporal domain and internal steps are taken in-between. One internal step here is equivalent to a single step made by the explicit Runge-Kutta method at the stability limit of the method, which could be tens or hundreds of times smaller than the time interval of the stretched-out F- and C-points. The concept is shown in Fig. 6 with 15

internal steps between multigrid time points, where each row of FCF-relaxation is performed concurrently. Although taking internal steps reduces the amount of temporal concurrency, there are similar reductions in communication and the algorithm's memory footprint. Consequently, taking internal steps often improves the overall performance of the algorithm. The right number of internal steps is currently determined experimentally. It is believed that a good number of internal steps balances relaxation to equilibrium in fine scales with feedback from their influence on coarse scales. Note that the number of internal steps is also a function of the time-stepping algorithm and the physics solved. For the compressible solver being used, the number of internal steps is quite high in order to satisfy restrictions related to acoustic time scales.

The number of C-points that a relaxation moves the problem forward in time is important for understanding the *exactness* property. Exactness means relaxation on the fine grid has advanced the solution sequentially all the way from the initial conditions (IC) to the end time. The solution is the same as sequential time stepping and one may as well have done that instead. Relying on the exactness property means a reduction in wall time will not be observed using PinT. For turbulence simulations, it means error in fine turbulent scales was propagated out to the end time, rather than being evolved to an equilibrium state. When spatial coarsening is used, MGRIT propagates the exact solution forward by one C-point at each iteration. The number of iterations required to converge the problem, compared to the number of iterations before exactness is realized, is key to performance. Ideally the former is a small fraction of the latter. In the results, this ratio is called the "fraction to exactness."

3. Test Cases

Two test cases are used to investigate the validity of the hypotheses. The first is a Taylor-Green problem solved at theoretically infinite Reynolds number. The second is a time-evolving double-shear mixing layer with physical viscosity providing a Reynolds number of 11650.

3.1. Taylor-Green Problem

The Taylor-Green problem is chosen to investigate the performance of the MGRIT algorithm applied to turbulence. The problem is initialized with a few large vortices which provide an energy store for decaying turbulence. The intent of this study is to investigate if MGRIT can solve inertial flows that are turbulent. Consequently, the viscosity for the flow is set to zero meaning the Reynolds number is theoretically infinite and the simulated physics are entirely inertial. In Fig. 1, the Kolmogorov scale would be 0 and the inertial cascade would continue as a flat line, never quite reaching $E = 0$. A drawback is that it is impossible to perform DNS or physical experiment to ascertain the true cascade. A consequence of zero viscosity is that numerical methods with a finite grid must provide some form of regularization to dissipate energy, either a turbulence model or other numerical stabilization mechanisms. For this research, the PPM limiter serves this purpose and also the role of modeling subgrid-scale terms. One can consider this as modeling energy transport to unrepresented turbulent scales (no physical diffusion and infinite Reynolds number) or mimicking physical diffusion in which case the Reynolds number is estimated as 1500 for a 64^3 mesh and 2300 for a 128^3 mesh, following the approach of Aspden et al. [23]. The true Reynolds number is somewhere between these values and infinity. The point of setting $\mu = \kappa = 0$ is to remove all physical diffusion and only have the minimum dissipation because the goal is to prove our hypothesis that the small scales are iterated to equilibrium and not simply damped. How well the simulations predict the true energy cascade is somewhat immaterial as long as the flow is representative of turbulence, because the same approach is used for both PinT and sequential solutions. For this particular study, PinT solutions are only compared against sequential solutions. Since the grid cutoffs are all in the inertial range of the energy cascade (if only barely for the resolution of 32^3), the simulations are representative of turbulence and the approach is appropriate for investigating whether MGRIT can be used to solve turbulent flows.

The Taylor-Green vortex flow is initialized in a fully-periodic cube of side-length D with a sinusoidal

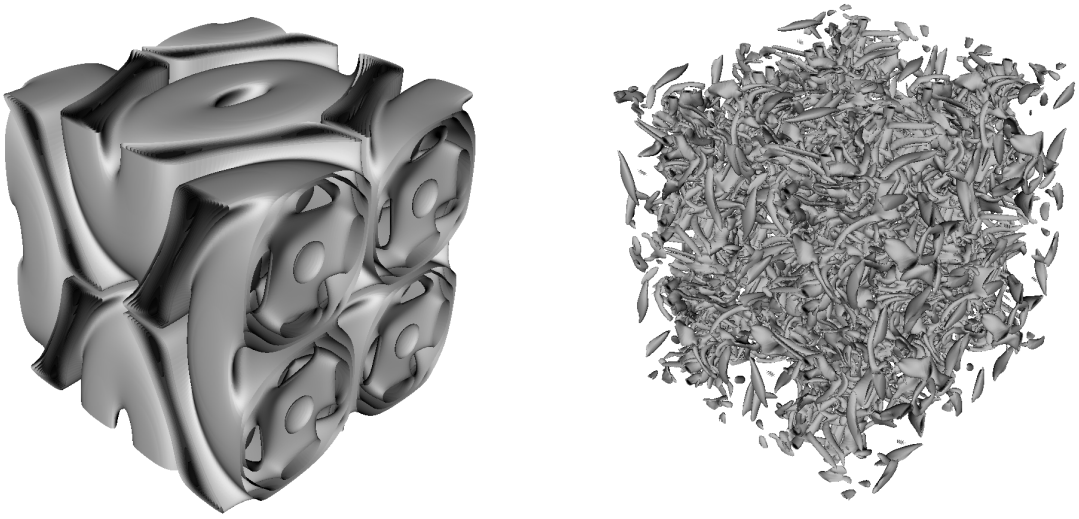

(a) Initial roll-up at $\tau = 2.1$.

(b) Fully-developed turbulence at $\tau = 18.7$.

Figure 7: Iso-surfaces of enstrophy in a Taylor-Green flow solved on a 128^3 grid. Eight large vortices decay into turbulence.

initial condition given by

$$u = -U_0 \sin\left(\frac{n\pi x}{D}\right) \cos\left(\frac{n\pi y}{D}\right) \sin\left(\frac{n\pi z}{D}\right) \tag{11}$$

$$v = U_0 \cos\left(\frac{n\pi x}{D}\right) \sin\left(\frac{n\pi y}{D}\right) \sin\left(\frac{n\pi z}{D}\right) \tag{12}$$

$$w = 0 \tag{13}$$

$$p = p_0 + \frac{\rho_0 U_0^2}{16}\left(\cos\left(\frac{2n\pi x}{D}\right) + \cos\left(\frac{2n\pi y}{D}\right)\right)\left(\cos\left(\frac{2n\pi z}{D}\right) + 2\right) \tag{14}$$

$$\rho = \frac{p}{RT_0} = \frac{p\rho_0}{p_0} \tag{15}$$

where U_0 is the velocity fluctuation magnitude and n is the number of vortices contained in the domain in each coordinate direction [14]. The flow has a Mach number based on U_0 of 0.1. Characteristic times based on the eddy turnover time are given by $\tau = tU_0/D$, where t is the simulation time. Decaying Taylor-Green flows transition to fully developed turbulence by $\tau \approx 10$. Enstrophy for a Taylor-Green flow is shown in Fig. 7 for initial roll-up at $\tau = 2.1$ and fully developed turbulence at $\tau = 18.7$.

This case is used in section 4 to investigate the behavior of applying MGRIT to turbulent flows. There, the mesh on the finest level is 64^3 and two temporal multigrid levels are use, separated by a refinement ratio of two. In section 5, the problem is solved on a 128^3 grid with three temporal multigrid levels, again using a refinement ratio of two.

3.2. Double-Shear Mixing Layer

The double-shear mixing layer differs from the Taylor-Green flow in that the turbulence is anisotropic. Furthermore, physical diffusion is included and the Reynolds number is 11650 based on the initial velocity U_∞ and the initial vorticity thickness, δ_w [14]. An overview of the flow problem is shown in Fig. 8. For the initial conditions, the velocities in each stream were sinusoidally perturbed and computed from a stream function in order to achieve an analytically divergence-free initial velocity field. A complete description of the initial conditions is given by Walters et al. [14]. The turbulence is decaying from the initial kinetic energy store of the opposite streams. Defining the characteristic eddy turn-over time as $\tau = tU_\infty/\delta_w$, the spectrum of kinetic energy is fully developed by $\tau \approx 20$ and the two layers begin interacting at $\tau \approx 35$.

This test case is explored in section 5 using a $128 \times 128 \times 64$ mesh and three temporal multigrid levels, with a refinement ratio of two between the levels. The temporal domain extends from $0 \leq \tau \leq 50$.

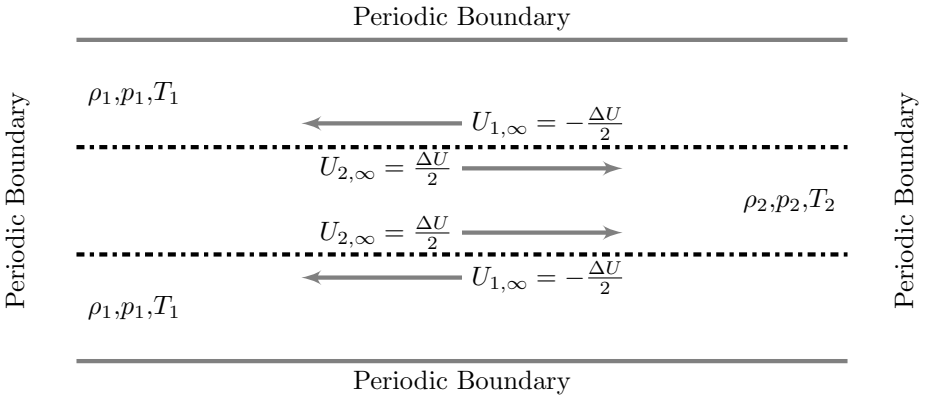

Figure 8: Diagram of the double-shear-layer configuration.

4. MGRIT Algorithm for Turbulence

In this section, strategies related to several algorithmic components are explored. Beginning with a demonstration of how an MGRIT algorithm, optimized for elliptic transport, fails to converge, strategies are then proposed for a viable algorithm. Pitfalls of applying MGRIT to turbulent flows are highlighted and remedies that were taken are presented. All studies involved solving a turbulent Taylor Green flow on a two-level multigrid mesh. The finer level has 64^3 spatial grid and the grid on the coarser level is 32^3. Early research activities explored a temporal domain from $20 \leq \tau \leq 21$ with the initial PinT solution on the coarse grid initialized from the sequential solutions at $\tau = 20$. The initial PinT solution on the fine grid is either initialized from the sequential solution or estimated from the coarse-grid solution. The latter approach is of course required in the general case where the sequential solution is not known. The final results cover the complete temporal domain from $0 \leq \tau \leq 20$. Solutions obtained by sequential time stepping were obtained on a 64^3 grid unless otherwise noted. In all cases, the PinT solution is compared to the sequential solution to assess correctness. This is done by comparing the turbulent energy spectrum for both cases. The spectrum is produced by a discrete Fourier transform of the velocity and summing the 3-D spectra into bins represented by spherical shells to produce plots of energy E versus wavenumber k. Supersampling is used to improve the distribution of the spectra into the shells—the grid in Fourier space is refined by a sampling factor and filled with piecewise-constant data before summing into shells of width Δk. All reported energy spectra are normalized by the total kinetic energy in the initial conditions.

4.1. Filtering

Using an MGRIT strategy optimized for diffusive flows does not converge to the sequential solution when applied to the Taylor-Green turbulent flow. This is the "Base" solution shown in Fig. 9. This figure shows both the residual and energy spectrum of the turbulent flow solved with XBraid. The residual is $\boldsymbol{g} - \mathrm{A}\mathbf{U}$ from (9) as computed on a spatial coarsening of the fine grid solution. In other words, fine scale features are not considered when evaluating the residual as they are likely chaotic. This residual is primarily used to assess convergence of the coarse scales. All grids were initialized to the sequential solution at $\tau = 20$, the IC, and evolved to $\tau = 21$. The asymptotic divergence in the residual manifests as the wrong energy spectrum compared to the sequential solution.

The divergence of the residual was thought to be a result of transferring invalid or corrupt scales, either arising from aliasing or discretization error. Spectral filtering was added separately to the restriction and prolongation operators in space. Weak convergence was only observed when the prolonged coarse-grid correction was filtered to a resolution of 16^3 before being added to the fine-grid solution. Although the convergence is weak, very good predictions are made of energy in small wave-numbers. The large-wave numbers (small scales of turbulence) have barely evolved and still follow the IC. Note that the small scales at $\tau = 20$ are nearly identical to those at $\tau = 21$. Thus, the influence of the small scales on the large scales is nearly identical at both times. The results indicate that discretization error on the coarse grid corrupts features from $8 < k < 16$. This is the action of the PPM limiter on the solution at the grid scale where

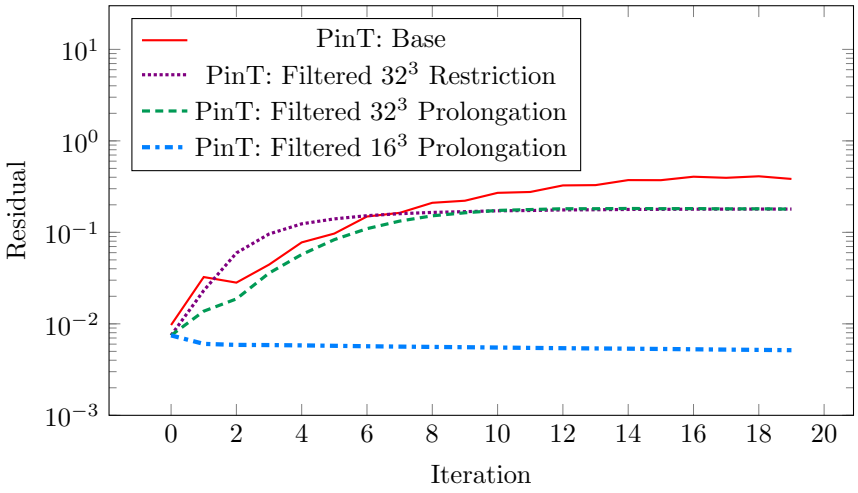

(a) Residuals.

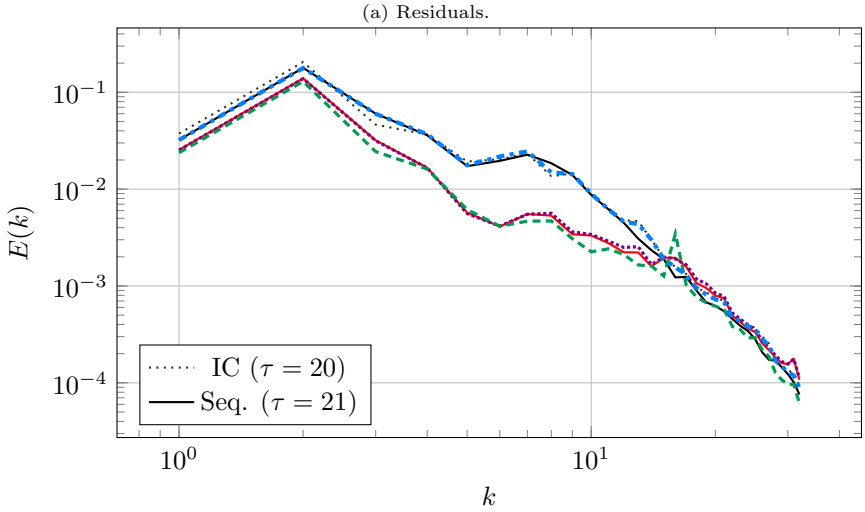

(b) Energy spectrum at iteration 12.

Figure 9: The base algorithm (optimized for diffusive flows) and the effect of filtering scales during restriction and prolongation. The two-level MGRIT spatial grids have resolutions 64^3 on the fine level and 32^3 on the coarse level. Weak convergence is only observed when filtering to a resolution of 16^3 during prolongation. The filtering prevents corrupted data, a result of grid-scale dissipation on the coarse grid, from being sent to the fine grid.

turbulence is recognized as discontinuities, causing increased dissipation. It is *not* an option to simply turn off the limiter because some form of grid-scale dissipation is required for this inviscid flow. Filtering the prolongation only transfers corrections from well-resolved scales on the coarse grid. This appears to be a necessary component for solving turbulent flows with MGRIT.

Two factors are relevant for the filtering, the refinement ratio between grids on different levels, and the coarsening factor for filtering (the ratio between a coarse grid and a grid representing the filtered solution). In this work, both factors are two. When prolonging from a coarse multigrid level with spatial cell spacing $\Delta\vec{x}$ to a finer multigrid level with spacing $\Delta\vec{x}/2$, the technique used here is to first interpolate in space to the finer level. Next, a low-pass spectral filter, implemented using FFTW, is applied that removes all solution content smaller than $2\Delta\vec{x}$.

4.2. Deconvolution

In the previous subsection, the fine-grid solution was initialized to the sequential solution at $\tau = 20$. For general problems, this is not possible and the fine-grid solution will have to be estimated from the first coarse-grid solution in an full multigrid (FMG) startup procedure. A simple approach is to interpolate from the coarse grid, but this will not populate fine scales. In this section, deconvolution follows the spatial interpolation to populate the fine scales. For non-projective filters, deconvolution is the inverse of filtering. Structural subgrid-scale models based on approximate deconvolution have been constructed using ad hoc mathematical procedure [24] and demonstrated for applications to incompressible and compressible turbulence [25, 26, 27, 28]. Figure 10 shows the averaging, interpolation, and deconvolution operators used in the work applied to the sample curve,

$$f(x) = \sin((8)2\pi x) + 0.5\sin((12)2\pi x) + 5\sin((32)2\pi x) + 7.3\sin((100)2\pi x)\,, \tag{16}$$

which is illustrated in the first row. Averaging from the reference to coarse is shown in the second row; interpolation back to the fine grid is shown in the third row; and the effect of deconvolution is shown in the fourth row. The differential deconvolution algorithm is given by

$$\phi = \langle\phi\rangle - \frac{\Delta_f^2}{24}\left.\frac{\partial^2\langle\phi\rangle}{\partial x_i^2}\right|_{(2)}\,, \tag{17}$$

where Δ_f is the filter width, $\partial^2/\partial x_i^2|_{(2)}$ is a second-order-accurate Laplacian, and $\langle\phi\rangle$ is the original solution state on the grid. Figure 11 shows the energy spectrum of the deconvolution (fourth row of Fig. 10). Note the additional waves beyond the original content in the reference solution ($k = 8, 12, 32,$ and 100) that arise beyond $k = 16$ due to aliasing errors.

Initial fine-grid solution estimates, interpolated and deconvolved from the coarse-grid solve as part of an FMG startup procedure, are shown in Fig. 12. Several different filter widths are tested. While $\Delta_f = 2\Delta x_{\text{coarse}}$ appears to best estimate the correct spectrum for this problem, a preference is made for the more conservative choice of $\Delta_f = \Delta x_{\text{coarse}}$. Adding in excessive energy content carries some risk and preliminary experimentation suggests there is not much performance benefit with MGRIT from using $\Delta_f = 2\Delta x_{\text{coarse}}$. A more rigorous study covering several different turbulent flow types would be needed to provide a recommendation for the best value of Δ_f. In the remaining results, $\Delta_f = \Delta x_{\text{coarse}}$ is used.

Still working in the restricted temporal domain of $20 \leq \tau \leq 21$, the effect of the coarse-grid correction on the fine-grid solution is shown in Fig. 13. The C-points are still using a spacing more optimal for diffusive flows, at every second time-step on the fine grid. Although there is little evolution of the fine scales, the improved initialization is sufficient to achieve convergence of coarse scales. A rapid convergence of coarse scales, influenced by the fine-scale initialization is observed.

4.3. Number of C-Points

In this section, attempts are made to converge the fine scales of the turbulent flow. In previous sections, $\Delta\tau = 1$, and initializing the fine scales of turbulence by deconvolution was sufficient for influencing the coarse scales. In this section, full simulations of the Taylor-Green problem from $0 \leq \tau \leq 20$ are performed and the fine scales must be evolved to equilibrium. Consider two possible cases representing extremes. The first was the previous case where the fine scales were only evolved by two time-steps per multigrid iteration. The

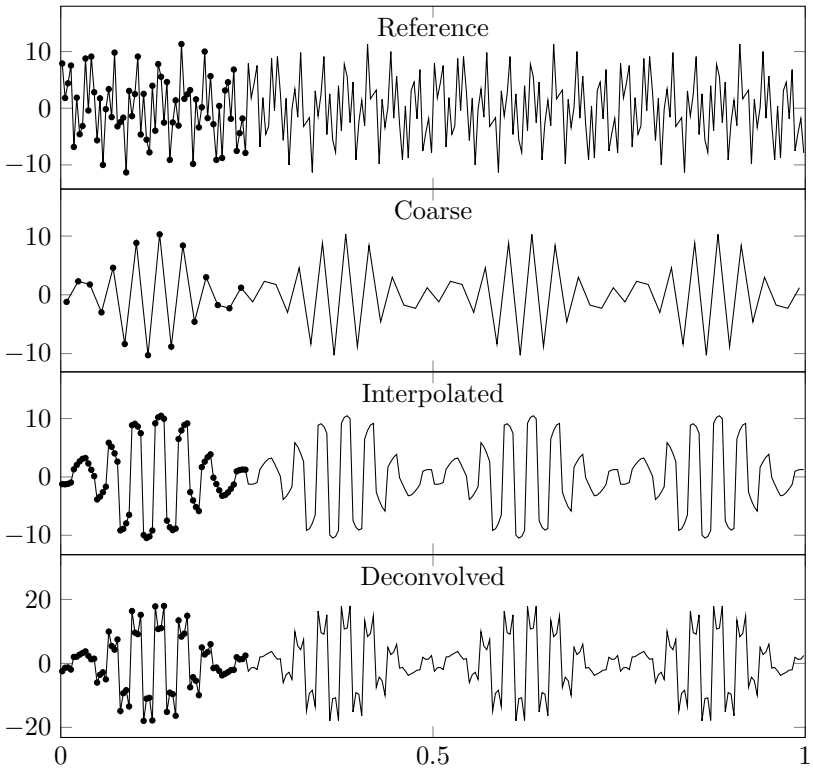

Figure 10: One-dimensional coarsening, interpolation, and deconvolution applied to a test function given by (16). Grid points are shown in the first quarter.

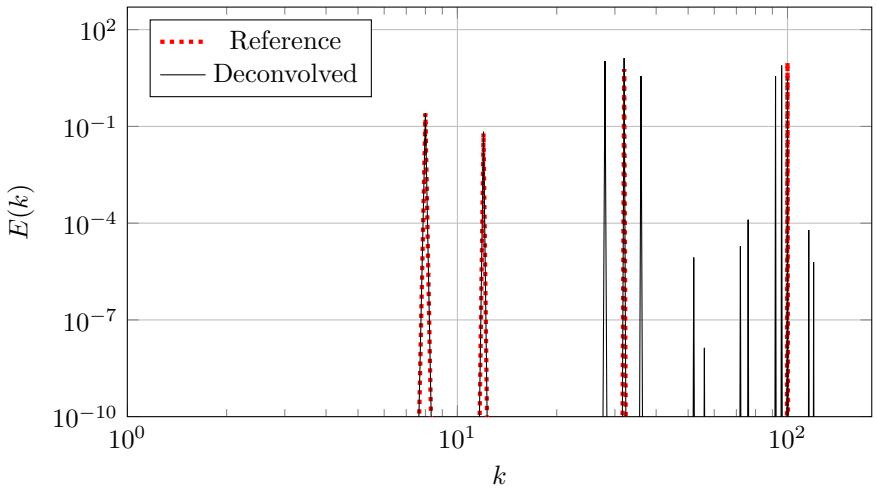

Figure 11: Energy spectrum of the deconvolution from Fig. 10. Only wavenumbers 8, 12, 32, and 100 were represented in the original reference solution.

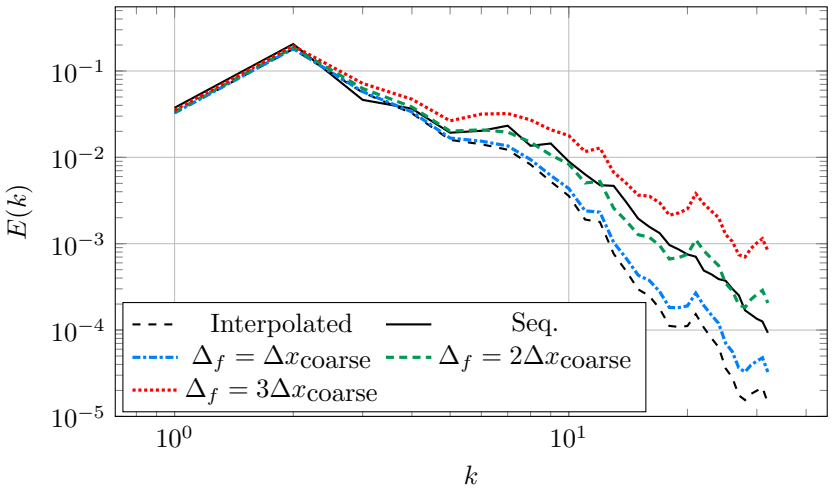

Figure 12: Initialization of the fine-grid solution following a coarse-grid solve as part of the FMG startup procedure. The effect of different filter widths for the deconvolution are shown.

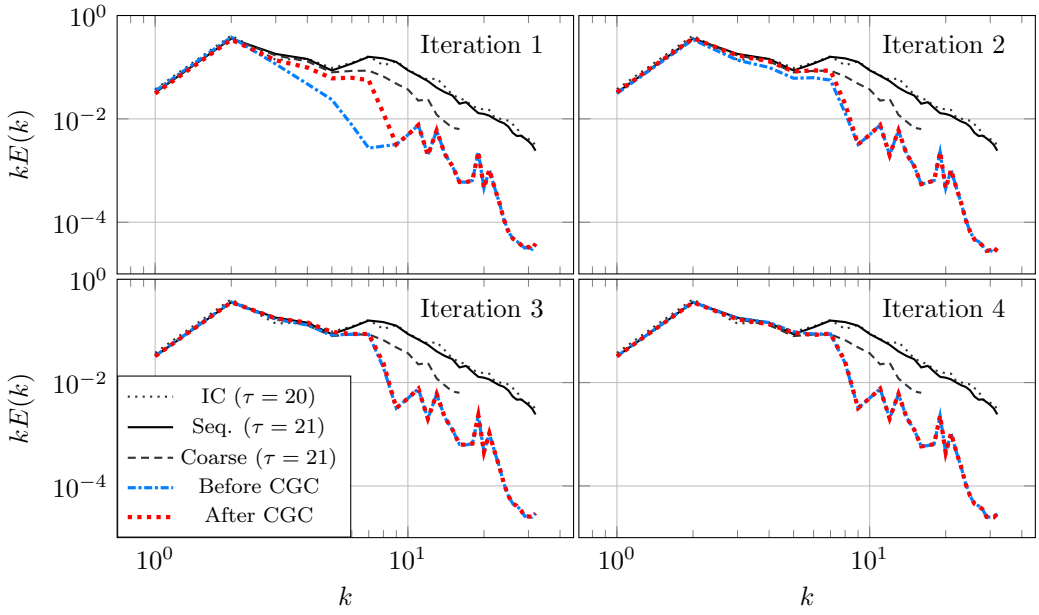

Figure 13: The effect of coarse grid corrections from MGRIT iterations at time $\tau = 21$. At each iteration, the solution on the fine grid is shown both before and after the coarse grid correction (CGC). The coarse scales rapidly converge under the influence of initialized fine-scales. For reference, the coarse sequential solution is also shown. Note that the y-axis is $kE(k)$ which amplifies differences in low wavenumbers.

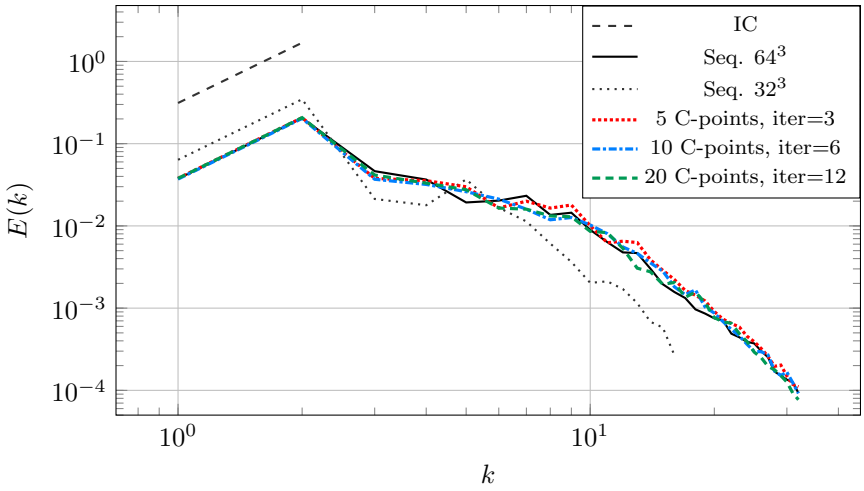

Figure 14: Energy spectrum at $\tau = 20$ of a Taylor-Green flow solved from $\tau = 0$ to $\tau = 20$ using a two-level MGRIT algorithm. Comparison of cases with different numbers of C-points.

opposite extreme is iterating the fine scales to equilibrium for each multigrid iteration. A preferred balance is that the fine scales reach equilibrium as changes in the coarse scales, due to the influence of fine scales, approach zero. Recall that the coarse scales are iterated sequentially for each multigrid iteration—all the error is removed by advecting to the end time. A balanced approach features a few multigrid iterations which is dictated by the number of C-points in the temporal domain as follows. The overall number of time steps on the fine grid remains constant; as the number of C-points is reduced, the number of internal steps increases (see Fig. 6). The maximum number of multigrid iterations is equal to the number of C-points in the domain. If this limit is reached, the solution has been sequentially stepped on the fine grid to reach the exactness property, obviating the point of MGRIT. Ideally convergence is achieved in fewer iterations, providing an opportunity for a parallel speedup. Figure 14 shows a converged iteration for solutions achieved with total C-points equaling 20, 10, and 5. This corresponds to temporal separations of 1, 2, and 4 τ between C-points, respectively. In terms of performance, fewer C-points are preferred because this minimizes the number of multigrid iterations and reduces the number of time points that must be stored. Although this reduces the temporal parallelization, it also minimizes the number of relaxations on the coarse grid which are fully sequential. Internal time-stepping only requires communication between spatial processors and not between sets of temporal processors.

In Fig. 14 solutions obtained with sequential time-stepping on a 32^3 grid and a 64^3 grid are also shown for comparison. The reason for showing the coarse-grid 32^3 sequential solution in Fig. 14 is that if it was indeed the same as the fine-grid 64^3 solution, then little effort would have been required by the MGRIT algorithm; the first iteration on the coarse grid would be nearly correct. The initial condition at $\tau = 0$, which only has two wavenumbers (see Eqs. (11) to (15)), is also shown.

For all cases, the problem is "reasonably converged" at 60% of the time-steps required before reaching exactness property (3, 6 and 12 multigrid iterations, respectively for the 5, 10, and 20 C-point cases). In this description, "reasonably converged" is assessed qualitatively by comparing against the sequential solution. This information is summarized in Table 1. Independent of the number of C-points, the fine scales require 60% of the sequential time steps to reach equilibrium.

4.4. MGRIT Convergence of Taylor-Green Flow

From the previous sub-section, convergence of the energy spectrum is shown at several iterations in Figs. 15 and 16 for the 20 and 5 C-point cases, respectively. In both cases, the coarsest scales converge quickly within a few iterations. The remaining iterations are spent converging the fine scales (taking 60% of the sequential steps as described previously). This suggests that the 5 C-point case has the closest balance between converging coarse and fine scales at the same rate. In Fig. 14, larger discrepancies remain in the medium scales ($4 < k < 10$). This issue is investigated more thoroughly in the next section.

Table 1: Summary of convergence properties for the results of Fig. 14. The exactness property is reached when the total number of time steps on the fine grid in the MGRIT algorithm equals that of the sequential algorithm. The "fraction to exactness" is the total number of fine-grid time steps taken by MGRIT to reach convergence at the indicated iteration divided by the number of time steps taken by the sequential algorithm.

C-points	MGRIT Iterations	Fraction to Exactness
5	3	0.6
10	6	0.6
20	12	0.6

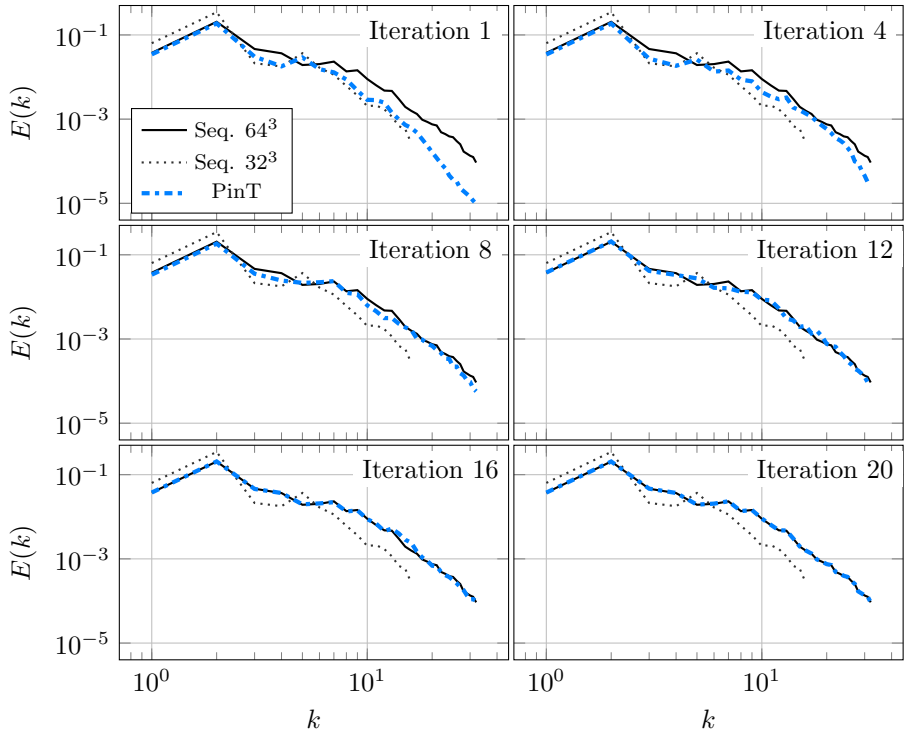

Figure 15: Energy spectrum at $\tau = 20$ of a Taylor-Green flow solved from $\tau = 0$ to $\tau = 20$ using a two-level MGRIT algorithm with 20 C-points.

The residual of coarse scales on the fine mesh is shown in Fig. 17. Note the residuals dropping to zero as the exactness property is realized. Convergence of all scales before exactness validates the hypothesis that the physics of turbulence can be used to enable temporal parallelization of flows dominated by inertial physics in an MGRIT context.

5. Results

The strategies from the previous section are applied to both test cases using three multigrid levels. The Taylor-Green problem is run with the finest mesh at 128^3. The finest mesh for the double-shear problem is $128 \times 128 \times 64$. For both cases, all parameters are held constant. During prolongation, each level is filtered by a factor of 2 with respect to the coarser level. For example, the Taylor-Green problem has three levels with spatial grids 32^3, 64^3, and 128^3. The prolonged scales are represented on spatial grids sized 16^3 and 32^3 before being added to the solution on the finer grid. An FMG startup procedure is used and initialization of finer levels is performed with a deconvolution filter width of $1\Delta x$. The spacing between the C-points is conservatively set to $5/3\tau$. As before, the Taylor-Green flow is solved to 20τ, now using 12 C-points. The double-shear is solved to 50τ using 30 C-points.

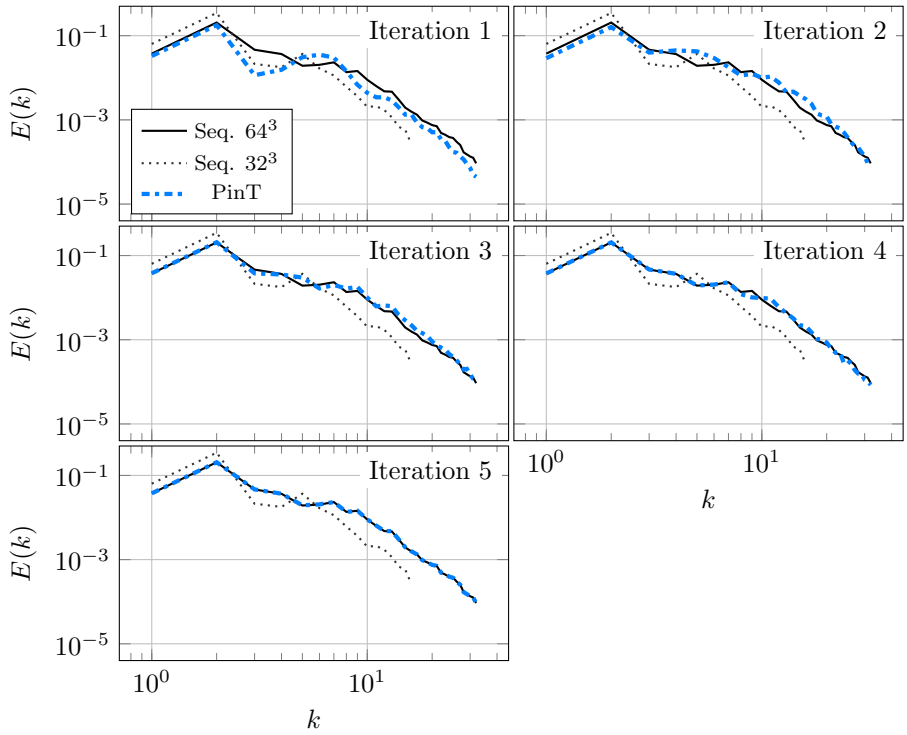

Figure 16: Energy spectrum at $\tau = 20$ of a Taylor-Green flow solved from $\tau = 0$ to $\tau = 20$ using a two-level MGRIT algorithm with 5 C-points.

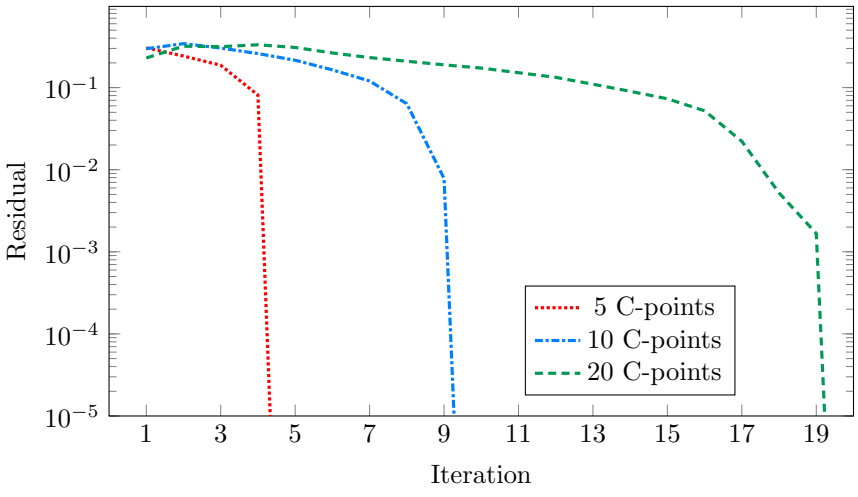

Figure 17: Residual on the coarse grid at each iteration for a two-level MGRIT solution of a Taylor-Green flow. The exactness property takes effect on iteration 5, 10, and 20 for the case with 5 C-points, 10 C-points, and 20 C-points, respectively.

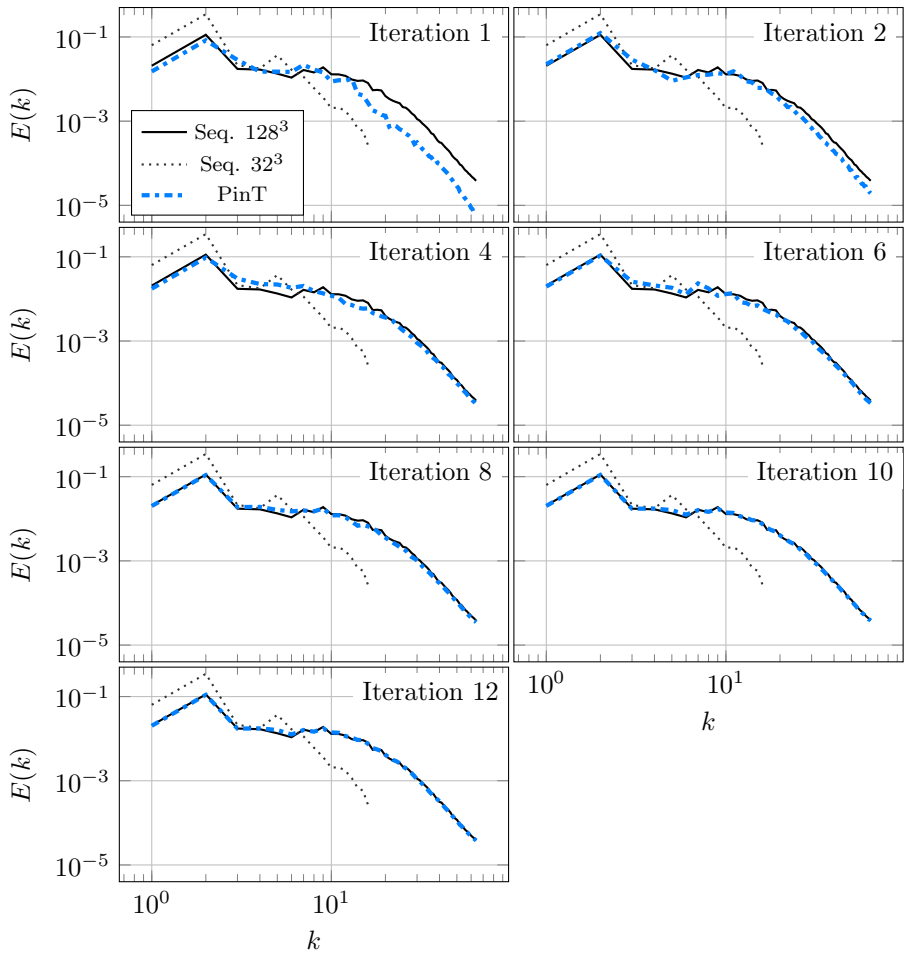

Figure 18: Energy spectrum at $\tau = 20$ of a Taylor-Green flow solved from $\tau = 0$ to $\tau = 20$ using a three-level MGRIT algorithm with 12 C-points. The finest grid has a resolution of 128^3.

5.1. Taylor-Green Flow

Energy spectrum at several iterations is shown in Fig. 18 for the Taylor-Green case solved on three levels. Convergence in large and small scales is achieved by iteration 4 but there are still discrepancies in the medium-sized scales. To understand this further, an ensemble of 25 sequential solutions were run with perturbations to the initial conditions in the wave numbers from $k = 13$ to $k = 64$. The magnitude of the perturbation was always less than 1% of the magnitude of the large vortices. The area covered by the ensemble is shown in Fig. 19 along with the original unperturbed sequential solution and iterations 2, 4, 8 of the MGRIT algorithm. MGRIT iteration 2 is outside the gray area, while iterations 4 and 8 fall within the ensemble. Similar data is shown in Fig. 20 but the shaded gray area now denotes a region of one-standard deviation away from the mean perturbed spectrum. The ensemble illustrates that the medium scales are chaotic and highly sensitive to the initial conditions. This is the reason MGRIT struggles to converge these scales before reaching the exactness property. Based on these plots, one can reasonably say that by iteration 4, the solution is converged within the ensemble. At some wave numbers, the solution is outside one standard deviation, but so is the unperturbed sequential solution.

There is one additional curiosity in Fig. 18, the exactness property is not perfectly satisfied. There are a few medium scales (at $k = 5, 6$) that are still slightly off and remain unchanged from iterations 8 to 12. This is thought to be a result of a small difference in floating point values exhibiting chaotic behavior. It may also be a result of some unknown and undesired interaction between levels in the three-level case. However, a similar effect is not observed in the double-shear test case.

The assessment of when convergence is achieved is expanded by considering changes to enstrophy. The

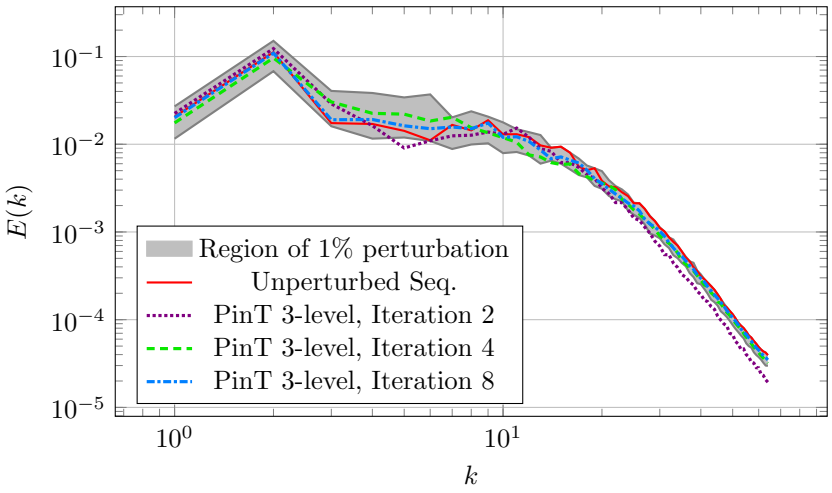

Figure 19: Time-sequential cases with a 1% perturbation of the initial conditions provides a region of valid energy values in gray. The original unperturbed time-sequential solution is shown in red. By the 4^{th} iteration, the time-parallel cases fall within the region of variance of the time-sequential cases.

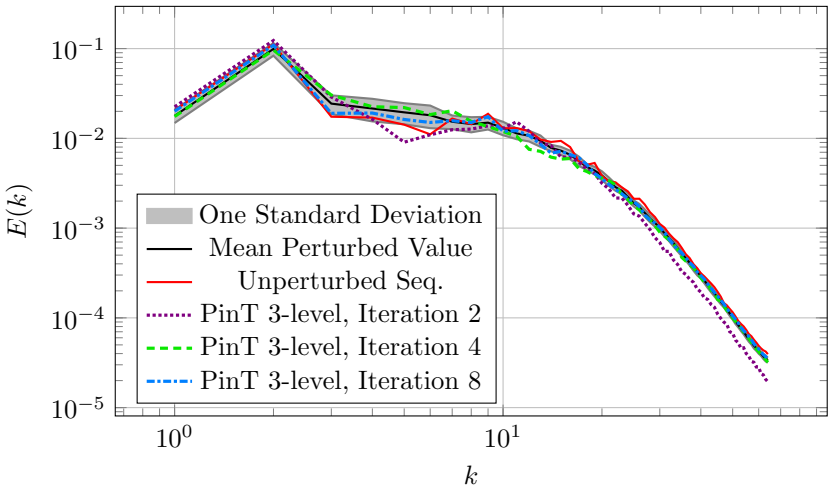

Figure 20: The mean perturbed solution is shown in black along with one standard deviation from the mean in the gray region. The original unperturbed time-sequential solution is shown in red. By the 4^{th} iteration, the time-parallel cases are as close to one standard deviation as the original unperturbed solution.

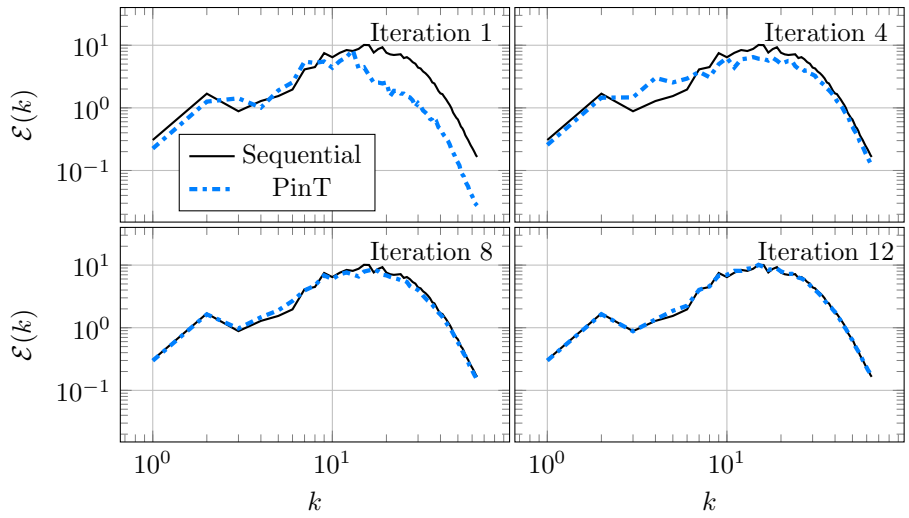

Figure 21: Enstrophy spectrum at $\tau = 20$ of a Taylor-Green flow solved from $\tau = 0$ to $\tau = 20$ using an MGRIT algorithm with 12 C-points. The finest grid has a resolution of 128^3.

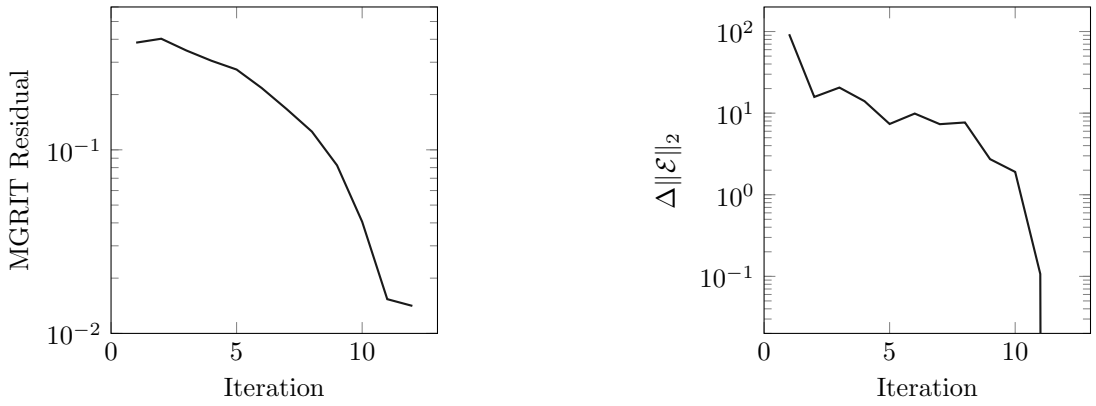

(a) MGRIT residual over iterations.

(b) Enstrophy residual over iterations.

Figure 22: Residuals of the two convergence measures for a 3-level Taylor-Green flow.

MGRIT residual, computed on a spatial coarsening of the fine grid solution, provides an assessment of the coarse scales. Something else is need to assess the fine scales. The enstrophy spectrum for several iterations of the 128^3 Taylor Green case is shown in Fig. 21. A characteristic of the enstrophy spectrum is that the largest magnitudes are for scales near the Taylor microscale. The Taylor microscale identifies the portion of the energy cascade just before dissipative effects become important. If a L_2-norm of the enstrophy is not changing, then eddies near the Taylor microscale are probably in equilibrium. Both the residual and change in enstrophy are plotted in Fig. 22. The residual behaves similar to that observed in Fig. 17 for the two level case, except there is no sharp drop off when exactness is reached due to the persistent perturbation in the medium scales. The enstrophy plateaus at iteration 5, roughly near where convergence was assessed from Fig. 18 to 20. At this point, the small scales are oscillating around some mean. In the last few iterations, as exactness is approached, $\Delta\mathcal{E}$ plummets to zero as expected.

5.2. Double-Shear Flow

Energy spectrum at several iterations is shown in Fig. 23 for the double-shear case solved on three levels. Convergence in large scales is achieved in 5 iterations and in small scales just after 10 iterations. Discrepancies in the medium scales are not evident as there were for the Taylor Green flow. Most likely, chaotic behavior in these scales is suppressed by the physical viscosity. Also, the exactness property is fully realized for this

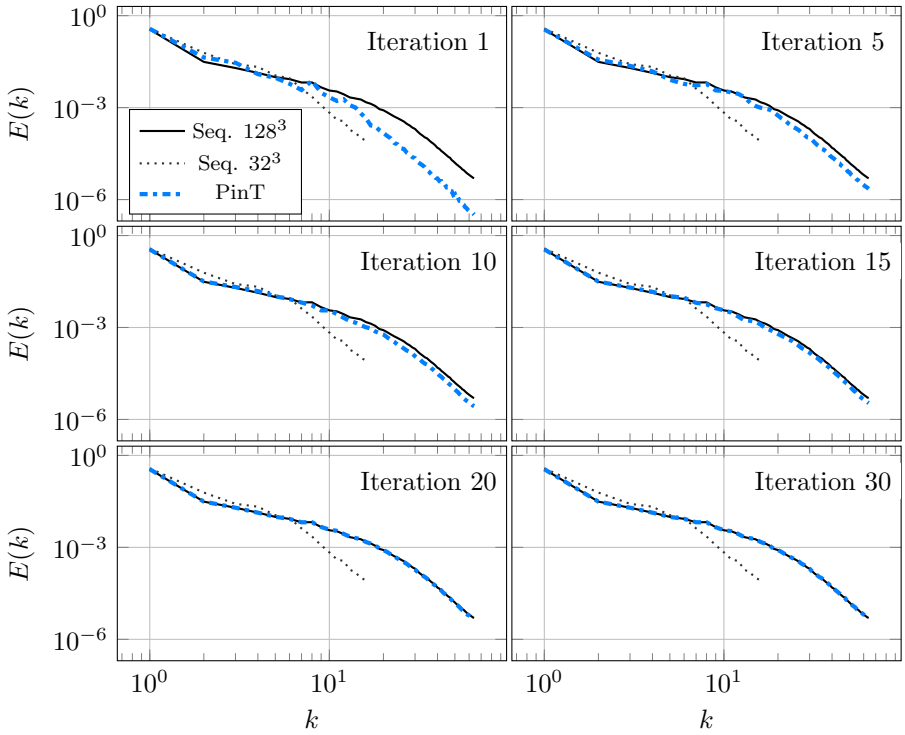

Figure 23: Energy spectrum at $\tau = 50$ of a double-shear-layer flow solved from $\tau = 0$ to $\tau = 50$ using a three-level MGRIT algorithm with 30 C-points. The finest grid has a resolution of $128 \times 128 \times 64$.

case as the residual drops to zero in Fig. 24a. Changes in enstrophy plateau at 10 iterations in Fig. 24b (again close to convergence identified from energy spectrum), and further decay from 20 iterations onwards.

The slow convergence of the small scales, with respect to the large scales, suggests that the C-points should be spaced much further apart for this problem. This indicates that optimal C-point spacing might be problem dependent.

5.3. Remarks on Assessing Convergence

In this work convergence was assessed by comparing the energy spectrum of the sequential and PinT solutions. For practical usage of the MGRIT algorithm, the sequential solution would not be known and other means are required to detect convergence. In the above results, iterations can be halted when the residual has converged by at least a small amount, and $\Delta\|\mathcal{E}\|_2$ has plateaued. Note that this is more easily detected when there are more C-points. These measures of convergence are not entirely satisfactory, being quite subjective, but may serve as a starting point for future improvements.

5.4. Remarks on Performance

It was considered premature to report on performance at this stage of the research. Nevertheless, readers should be aware that the authors do not consider the algorithm to be viable unless a roughly one-order reduction in wall-clock time can be realized for solving turbulent flows. This is a subject of future research. Note that both the Taylor-Green and double-shear cases converge when the finest level has taken only ≈ 0.35 of the steps taken by the sequential algorithm. This is much lower than fractions observed for the two-level Taylor-Green case and reported in Table 1. As this fraction reduces, the performance should increase. For high-fidelity solutions with many more multigrid levels and optimized time integrators, it is speculated that a one-order reduction in wall-clock time is possible.

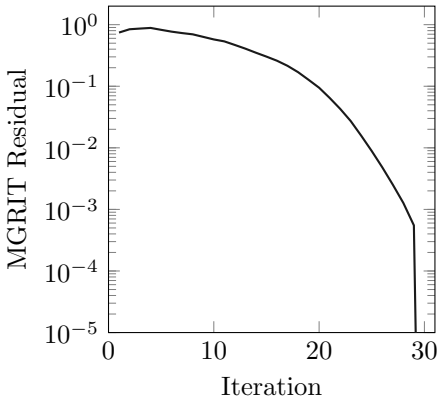

(a) MGRIT residual over iterations.

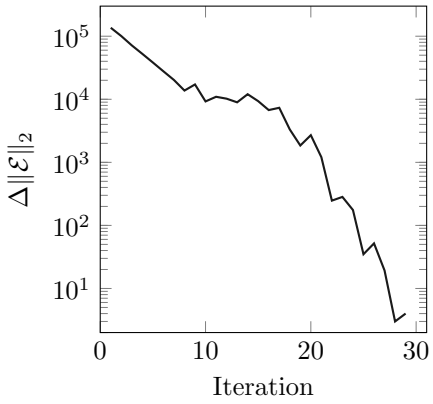

(b) Enstrophy residual over iterations.

Figure 24: Residuals of the two convergence measures for the double-shear-layer case. MGRIT's exactness property is reached at the 30^{th} iteration and the MGRIT residual goes to machine-zero.

6. Conclusion

The first hypothesis motivating this research was that the physics of turbulence could permit temporal parallelization using the MGRIT algorithm for a purely inertial flow described by hyperbolic PDEs. This was demonstrated by solving an infinite-Reynolds number Taylor-Green flow and a finite-Reynolds number double-shear flow and showing that the solution can be converged to have matching energy spectrum before reaching the exactness property (a recovery of sequential time stepping). Coarse-scale errors were advected out of the temporal domain while fine-scale errors were iterated to equilibrium by parallelizing the temporal domain. The approach assumes and takes advantage of the space-time locality of fine turbulent scales. The second hypothesis was that the fine and coarse scales could be separately converged, under each other's influence. In other words, the rates of convergence can be adjusted by the selection of multigrid-in-time C-points and internal time-stepping of the CFD algorithm. The ability to adjust convergence rates opens the door to adaptively switching between modeling and resolving fine turbulent scales, a topic for future research. An apparent necessity for MGRIT-based convergence of CFD algorithms that are not spectrally accurate is the filtering, during prolongation, of coarse-scale features that might have been corrupted by numerical dissipation at the coarse grid scale. The results indicate that the concept of applying MGRIT to turbulent flows has merit. Research is ongoing in applying the approach to a broader set of turbulence problems and exploring the performance of the algorithm.

7. Acknowledgments

Research at CSU was supported by sub-contracts B643713 and B647699 from Lawrence Livermore National Security, LLC. These subcontract were under prime contract No. DE-AC52-07NA27344 between LLNS and the United States Government.

References

[1] R. D. Falgout, S. Friedhoff, T. V. Kolev, S. P. MacLachlan, J. B. Schroder, Parallel time integration with multigrid, SIAM J. Sci. Comput. 36 (6) (2014) C635–C661.

[2] S. B. Pope, Ten questions concerning the large-eddy simulation of turbulent flows, New Journal of Physics 6 (2004) 1–24. doi:10.1088/1367-2630/6/1/035.

[3] A. Misra, D. I. Pullin, A vortex-based subgrid stress model for large-eddy simulation, Physics of Fluids 9 (8) (1997) 2443–2454.

[4] T. S. Lundgren, Strained spiral vortex model for turbulent fine structure, Phys. Fluids 25 (12) (1982) 2193–2203.

[5] T. H. Pulliam, D. W. Zingg, Fundamental Algorithms in Computational Fluid Dynamics, 1st Edition, Scientific Computation, Spinger International Publishing, New Delhi, India, 2014.

[6] A. J. Howse, H. De Sterck, R. D. Falgout, S. P. MacLachlan, J. Schroder, Parallel-in-time multigrid with adaptive spatial coarsening for the linear advection and inviscid Burgers equations, SIAM J. Sci. Comput. 41 (1) (2019) A538–A565.

[7] H. Hans De Sterck, R. D. Falgout, S. Friedhoff, O. A. Krzysik, S. P. MacLachlan, Optimizing multigrid reduction-in-time and Parareal coarse-grid operators for linear advection, Numer. Linear Algebra Appl. 28 (4) (2021) , https://doi.org/10.1002/nla.2367.

[8] F. P. Hamon, M. Schreiber, M. L. Minion, Parallel-in-time multi-level integration of the shallow-water equations on the rotating sphere, JCP 407 () (2020) , https://doi.org/10.1016/j.jcp.2019.109210.

[9] J. M. Reynolds-Barredo, D. E. Newman, R. Sanchez, D. Samaddar, L. A. Berry, W. R. Elwasif, Mechanisms for the convergence of time-parallelized, parareal turbulent plasma simulations, J. Comput. Phys. 231 (2012) 7851–7867.

[10] J. M. Reynolds-Barredo, D. E. Newman, R. Sanchez, An analytic model for the convergence of turbulent simulations time-parallelized via the parareal algorithm, JCP 255 (2013) 293–315.

[11] T. Lunet, J. Bodart, S. Gratton, X. Vasseur, Time-parallel simulation of the decay of homogeneous turbulence using Parareal with spatial coarsening, Comput. Vis. Sci. 19 (2018) 31–44.

[12] M. J. Gander, S. Vandewalle, Analysis of the parareal time-parallel time-integration method, SIAM J. Sci. Comput. 29 (2) (2007) 556–578.

[13] M. J. Gander, F. Kwok, H. Zhang, Multigrid interpretations of the parareal algorithm leading to an overlapping variant and MGRIT, Comput. Vis. Sci. 19 (3–4) (2018) 59–74.

[14] S. Walters, X. Gao, H. Johansen, S. Guzik, Assessing stretched-vortex subgrid-scale models in finite volume methods for unbounded turbulent flows, Flow, Turbulence and Combustion 106 (2020) 945–969, https://doi.org/10.1007/s10494-020-00206-1.

[15] D. Chung, D. I. Pullin, Direct numerical simulation and large-eddy simulation of stationary buoyancy-driven turbulence, J. Fluid Mech. 643 (2010) 279–308.

[16] T. W. Mattner, Large-eddy simulations of turbulent mixing layers using the stretched vortex model, J. Fluid Mech. 671 (2011) 507–534.

[17] S. M. Guzik, X. Gao, L. D. Owen, P. McCorquodale, P. Colella, A freestream-preserving fourth-order finite-volume method in mapped coordinates with adaptive-mesh refinement, Comput. Fluids 123 (2015) 202–217.

[18] X. Gao, L. D. Owen, S. M. J. Guzik, A parallel adaptive numerical method with generalized curvilinear coordinate transformation for compressible Navier-Stokes equations, Int. J. Numer. Meth. Fluids 82 (2016) 664–688.

[19] P. McCorquodale, P. Colella, A high-order finite-volume method for conservation laws on locally refined grids, Comm. App. Math. Comput. Sci. 6 (1) (2011) 1–25.

[20] J. Christoper, R. D. Falgout, S. J. B., S. M. Guzik, X. Gao, A space-time parallel algorithm with adaptive mesh refinement for computational fluid dynamics, Comput. Vis. Sci. 23 (13) (2020) , https://doi.org/10.1007/s00791-020-00334-1.

[21] XBraid: Parallel time integration with multigrid, https://computation.llnl.gov/projects/parallel-time-integration-multigrid.

[22] A. Brandt, Multi-level adaptive solutions to boundary-value problems, Math. Comp. 31 (138) (1977) 333–390.

[23] A. Aspden, N. Nikiforakis, S. Dalziel, J. B. Bell, Analysis of implicit LES methods, Comm. App. Math. and Comp. Sci. 3 (1) (2008) .

[24] J. A. Domaradzki, N. A. Adams, Direct modelling of subgrid scales of turbulence in large eddy simulations, Journal of Turbulence 3 (2002) , https://doi.org/10.1088/1468-5248/3/1/024.

[25] P. Sagaut, Large Eddy Simulation for Incompressible Flows: An Introduction, 3rd Edition, Springer-Verlag, 1998.

[26] E. Garnier, N. Adams, P. Sagaut, Large Eddy Simulation for Compressible Flows, 1st Edition, Springer, 2009.

[27] S. Stolz, A. Adams, An approximate deconvolution procedure for large-eddy simulation, Physics of Fluids 11 (7) (1999) 1699–1701.

[28] Q. Wang, M. Ihme, A regularized deconvolution method for turbulent closure modeling in implicitly filtered large-eddy simulation, Combustion and Flame 204 (2019) 341–355, https://doi.org/10.1016/j.combustflame.2019.03.009.