Turbulence, Heat and Mass Transfer 10

On the evolution of Poisson solvers for extreme scale simulations

A. Alsalti-Baldellou[1,2], F. X. Trias[1] and A. Oliva[1]
[1] *Heat and Mass Transfer Technological Center, Technical University of Catalonia, Carrer de Colom 11, 08222 Terrassa (Barcelona), adel.alsalti@upc.edu*
[2] *Termo Fluids S.L., Carrer de Magí Colet 8, 08204 Sabadell (Barcelona), Spain*

Abstract – In this work, we aim to shed light to the following research question: *is the complexity of numerically solving Poisson's equation increasing or decreasing for very large DNS and LES simulations of incompressible flows?* Physical and numerical arguments are combined to derive power-law scalings at very high Reynolds numbers. Preliminary results of forced homogeneous isotropic turbulence seems to confirm these scalings. Theoretical convergence analysis for both Jacobi and multigrid solvers defines a two-dimensional phase space divided into two regions depending whether the number of solver iterations tend to decrease or increase with the Reynolds number. Preliminary results seem to confirm that we are indeed in the second region.

1. Introduction

We consider the simulation of turbulent incompressible flows of Newtonian fluids. Under these assumptions, the governing equations read

$$\partial_t u + (u\cdot\nabla)u = \rho^{-1}\nabla\cdot(2\mu\mathsf{S}(u)) - \nabla p, \qquad \nabla\cdot u = 0, \tag{1}$$

where $u(x,t)$ and $p(x,t)$ denote the velocity and pressure fields, and $\mathsf{S} = 1/2(\nabla u + \nabla u^T)$ is the rate-of-strain tensor. The density, ρ, is constant whereas the dynamic viscosity, $\mu(x,t)$, may depend on space and time. Notice that for (spatially) constant viscosity the diffusive term simplifies to $\nu\nabla^2 u$ where $\nu = \mu/\rho$ is the kinematic viscosity. Then, these equations have to be discretized both in space and time.

The basic physical properties of the Navier–Stokes (NS) equations (1) are deduced from the symmetries of the differential operators (see Ref.[1], for example). In a discrete setting, such operator symmetries must be retained to preserve the analogous (invariant) properties of the continuous equations [2, 3]: namely, the convective operator is represented by a skew-symmetric matrix, the diffusive operator by a symmetric, negative-definite matrix and the divergence is minus the transpose of the gradient operator. Therefore, even for coarse grids, the energy of the resolved scales of motion is convected in a stable manner, *i.e.,* the discrete convective operator transports energy from a resolved scale of motion to other resolved scales without dissipating any energy, as it should be from a physical point-of-view. Furthermore, high-order symmetry-preserving discretizatons can be constructed for Cartesian staggered grids [2]. It is noteworthy to mention that in the last decade, many DNS reference results have been successfully generated using this type of discretization (see Figure 1, for example).

On the other hand, from the pioneering turbulent channel flow simulations in the mid-'80s [6], DNS and LES simulations of incompressible flows have usually been carried out by means of a fractional step method (FSM) together with explicit or semi-implicit time-integration

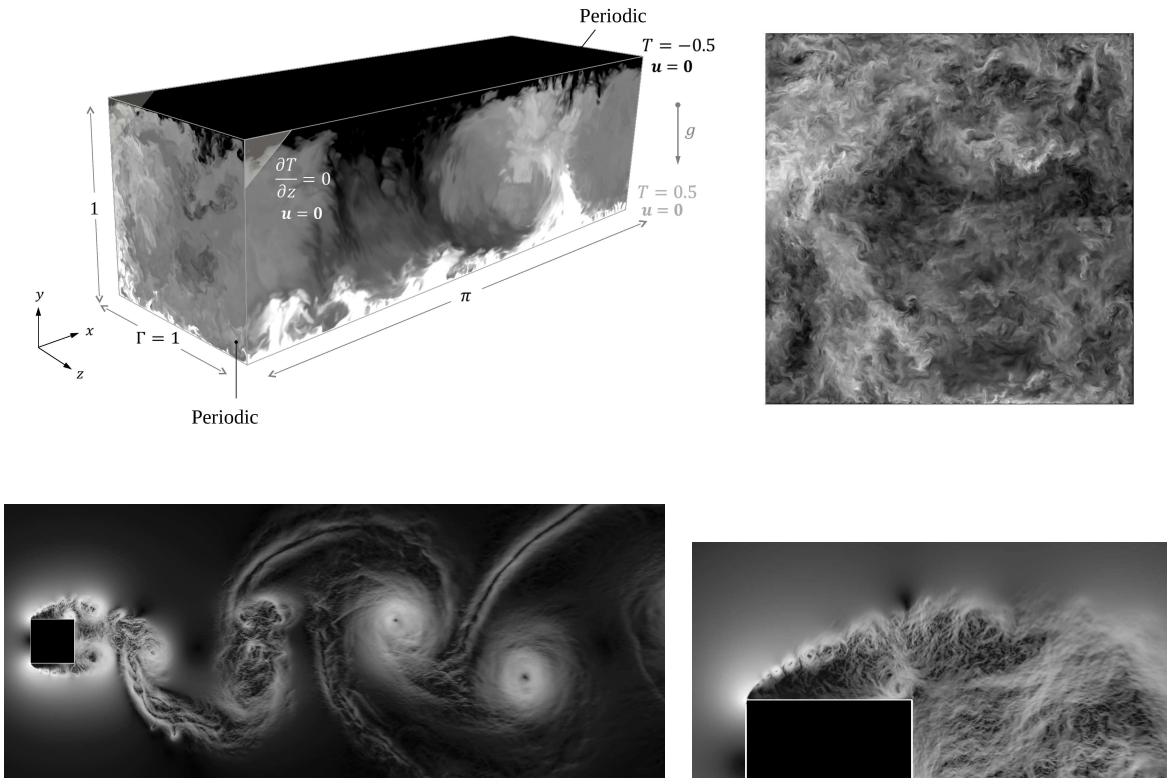

Figure 1: Examples of DNSs computed using symmetry-preserving discretizations. Top: air-filled ($Pr = 0.7$) Rayleigh-Bénard configuration studied in Ref. [4]. Instantaneous temperature field at $Ra = 10^{10}$ (left) and instantaneous velocity magnitude at $Ra = 10^{11}$ (right) for a spanwise cross section are shown. The latter was computed on 8192 CPU cores of the MareNostrum 4 supercomputer on a mesh of 5.7 billion grid points. Bottom: DNS of the turbulent flow around a square cylinder at $Re = 22000$ computed on 784 CPU cores of the MareNostrum 3 supercomputer on a mesh of 323 million grid points [5] .

methods for momentum, *i.e.,* the convective term is virtually always integrated with an explicit scheme whereas the diffusive term may be treated either implicitly or explicitly. This leads to the following CFL restrictions for the time-step

$$\Delta t^{\text{conv}} \sim \frac{\Delta x}{U} \qquad \Delta t^{\text{diff}} \sim \frac{\Delta x^2}{\nu}, \tag{2}$$

where Δx represents the grid spacing and U is a characteristic velocity of the large flow scales. Then, the pressure-velocity coupling is solved with a FSM leading to a Poisson equation for pressure, which is the most time-consuming part and the main bottleneck for extreme-scale CFD simulations. In this context, we aim to shed light to the following research question: *is the complexity of numerically solving Poisson's equation increasing or decreasing for very large DNS and LES simulations of incompressible flows?*

2. Two competing effects

2.1 Mesh size and time-step as a function of Reynolds

The never-ending increasing capacity of modern HPC system enables running DNS simulations at a higher and higher Reynolds number, Re. Estimations of how the number of grid points, N_x, and time-steps, N_t, grow with $Re = Ul/\nu$ can be easily obtained from the classical Kolmogorov theory (K41)

$$N_x^{K41} = \frac{L_x}{\Delta x} \sim \frac{l}{\eta} \sim Re^{3/4} \qquad N_t^{K41} = \frac{t_{\text{sim}}}{\Delta t} \sim \frac{t_l}{t_\eta} \sim \frac{l}{\eta}\frac{u}{U} \sim Re^{3/4}Re^{-1/4} = Re^{1/2}, \tag{3}$$

where L_x and t_{sim} are the domain size and the time integration period, which are assumed to be similar to the size of the largest scales, l, and its corresponding characteristic time, $t_l \sim l/U$, *i.e.,* $L_x \sim l$ and $t_{\text{sim}} \sim t_l$. For a DNS, we can also assume that $\Delta x \sim \eta$ and $\Delta t \sim t_\eta \sim \eta/u$, where $t_\eta \sim \eta/u$ and u are the characteristic time and velocity of the Kolmogorov length scales, η. Plugging this into Eq.(2) leads to the following estimations

$$N_t^{\text{conv}} \sim \frac{t_l}{\Delta t^{\text{conv}}} \sim \frac{l}{U}\frac{U}{lRe^{-3/4}} = Re^{3/4} \qquad N_t^{\text{diff}} \sim \frac{t_l}{\Delta t^{\text{diff}}} \sim \frac{l}{U}\frac{\nu}{l^2(Re^{-3/4})^2} = Re^{1/2}. \tag{4}$$

Therefore, we can conclude that

$$\frac{\Delta t}{t_l} \sim \frac{1}{N_t} \sim Re^{\alpha}, \tag{5}$$

where $\alpha = -1/2$ for the K41 theory (see Eq. 3) or diffusion dominated (see Eq. 4), and $\alpha = -3/4$ for convection dominated (see Eq. 4). Therefore, higher Re leads to (i) larger meshes and (ii) smaller time-steps, Δt. These are two competing effects on the convergence of Poisson's equation: namely, the former increases the condition number of the discrete Poisson equation whereas the latter leads to better initial guess.

2.2 Analysis of the residual of Poisson's equation

Next step is to analyze the residual of Poisson's equation as a function of the Reynolds numbers. There are two relevant aspects in this regard: the magnitude and the spectral distribution. To study this, let us consider a FSM where the so-called predictor velocity, u^p, results from a fully explicit forward Euler scheme

$$u^p = u^n + \Delta t R(u^n) \qquad \text{where} \quad R(u) = -(u\cdot\nabla)u + \rho^{-1}\nabla\cdot(2\mu\mathsf{S}(u)). \tag{6}$$

The forthcoming analysis could also be done for more appropriate (higher-order) time-integration schemes; however, for simplicity, we restrict ourselves to the first-order forward Euler. Then, imposing that the velocity at the next time-step must be divergence-free, $\nabla\cdot u^{n+1} = 0$, leads to a Poisson equation for pressure, p^{n+1},

$$u^{n+1} = u^p - \Delta t\nabla p^{n+1} \qquad \xRightarrow{\nabla\cdot} \qquad \nabla^2 p^{n+1} = \frac{1}{\Delta t}\nabla\cdot u^p. \tag{7}$$

Finally, assuming that $\nabla\cdot u^n = 0$ and taking p^n as initial guess, we obtain the following initial residual

$$r^0 = \nabla^2 p^n - \frac{1}{\Delta t}\nabla\cdot u^{p,n+1} \overset{(7)}{=} \frac{1}{\Delta t}\left(\nabla\cdot u^{p,n} - \nabla\cdot u^{p,n+1}\right) \approx \frac{\partial \nabla\cdot u^p}{\partial t}. \tag{8}$$

Here, for simplicity, we are assuming that the Δt remains constant. Alternatively, we can also consider $\tilde{r}^0 = \Delta t r^0$. In this case, the residual reads

$$\tilde{r}^0 = \nabla^2 \tilde{p}^n - \nabla \cdot u^{p,n+1} \overset{(7)}{=} \left(\nabla \cdot u^{p,n} - \nabla \cdot u^{p,n+1}\right) \approx \Delta t \frac{\partial \nabla \cdot u^p}{\partial t}, \tag{9}$$

where $\tilde{p} = \Delta t p$ is a pseudo-pressure. Notice that the second residual, $\tilde{r}^0$, is more meaningful from a physical point-of-view, since it directly translates on how accurately we are imposing the incompressibility constraint, $\nabla \cdot u = 0$.

At this point, we need to recall that $\nabla \cdot u^p$ can be expressed as follows [7]

$$\nabla \cdot u^p \approx \Delta t \nabla \cdot (u \cdot \nabla u) = 2\Delta t Q_{\mathsf{G}}, \tag{10}$$

where $Q_G = -1/2 tr(\mathsf{G}^2)$ is the second invariant of the velocity gradient tensor, $\mathsf{G} \equiv \nabla u$. Plugging this into Eqs. (8) and (9) leads to

$$r^0 \approx 2\Delta t \ \frac{\partial Q_{\mathsf{G}}}{\partial t}, \tag{11}$$

$$\tilde{r}^0 \approx 2\Delta t^2 \frac{\partial Q_{\mathsf{G}}}{\partial t}, \tag{12}$$

Therefore, smaller time-steps will decrease the magnitude of the residual leading to a better convergence of Poisson solver.

On the other hand, increasing *Re* also leads to finer meshes (see Eq. 3) and, therefore, to more ill-conditioned systems with a wider and wider range of scales to be resolved. In the forthcoming analysis, the spectral distribution of the initial residual, $\hat{r}_k^0$, plays a crucial role. In general, we can assume a power-law scaling within the inertial range

$$\frac{\partial Q_{\mathsf{G}}}{\partial t} \propto k^\beta \qquad \Longrightarrow \qquad \hat{r}_k^0 \propto \Delta t^p k^\beta, \tag{13}$$

where k is the wavenumber and $p \in \{1,2\}$ depends on the definition of the residual: $p = 1$ for Eq. (8) and $p = 2$ for Eq. (9). Then, a power-law scaling for Q_{G} can be derived from Eqs.(7) and (10)

$$2Q_{\mathsf{G}} = \nabla^2 p, \tag{14}$$

and the $k^{-7/3}$ scaling of the shell-summed squared pressure spectrum [8],

$$(\hat{Q}_{\mathsf{G}})_k \propto k^2 (k^{-7/3})^{1/2} = k^{5/6}. \tag{15}$$

Then, the value of β in Eq.(13) can be estimated from the dynamics of the invariants obtained from the so-called restricted Euler equations [9],

$$\frac{dQ_{\mathsf{G}}}{dt} = -3R_{\mathsf{G}} \qquad \Longrightarrow \qquad \frac{\partial Q_{\mathsf{G}}}{\partial t} = -(u \cdot \nabla) Q_{\mathsf{G}} - 3R_{\mathsf{G}}, \tag{16}$$

where $R_{\mathsf{G}} = det(\mathsf{G}) = 1/3 tr(\mathsf{G}^3)$ is the third invariant of G. Two terms in the right-hand-side of

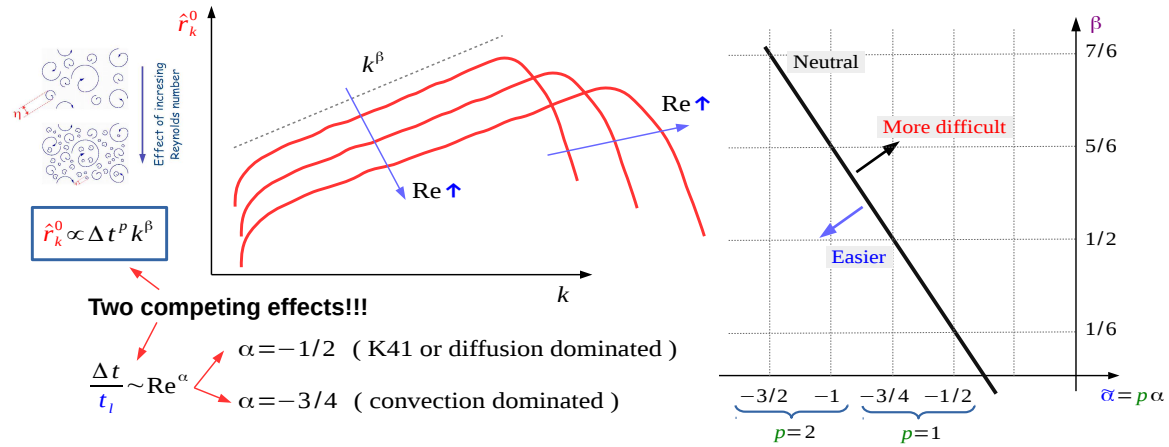

Figure 2: Left: illustrative explanation of the two competing effects on the solution of Poisson's equation when increasing Re number: time-step, Δt, decreases whereas the range of scales increases. Right: $\{\tilde{\alpha}, \beta\}$ phase space. Solid black line corresponds to $\propto Re^0$ in Eqs.(22) and (23), *i.e.,* neutral effect of Re-number in the total number of iterations.

Eq.(16) are expected to have different power-law scaling. Namely,

$$((u\cdot\widehat{\nabla})Q_\mathsf{G})_k \propto (\widehat{\nabla Q_\mathsf{G}})_k \propto k(k^{5/6}) = k^{11/6}, \tag{17}$$

$$(\hat{R}_\mathsf{G})_k \propto (k^{5/6})^{3/2} = k^{5/4}, \tag{18}$$

where the Taylor's frozen-turbulence hypothesis is applied to approximate the power-law scaling of the convective term, $(u\cdot\nabla)Q_\mathsf{G}$. Finally, combining the results obtained in Eqs.(13), (16) and (17) leads to

$$\boxed{\hat{r}_k^0 \propto \Delta t^p k^\beta \quad \text{with} \quad \beta = 11/6 \quad \text{and} \quad p = \begin{cases} 1 & \text{if } \hat{r} \text{ defined as Eq.(8)} \\ 2 & \text{if } \hat{r} \text{ defined as Eq.(9)} \end{cases}} \tag{19}$$

In summary, there are two competing effects (see Figure 2, left) when increasing Re number: time-step, Δt, decreases whereas the range of scales increases. The next step is to analyze how the solver convergence is affected.

3. Analysis of the solver convergence

We want to study whether the number of iterations inside the Poisson's solver increases or decreases with Re. To do so, we can relate the L2-norm of the residual with the integral of $\hat{r}_k$ for all the wavenumbers using the Parseval's theorem

$$||r||^2 = \int_\Omega r^2 dV = \int_1^{k_\mathrm{max}} \hat{r}_k^2 dk, \tag{20}$$

where $k_\mathrm{max} \approx 1/\eta \sim Re^{3/4}$. Then, the residual after n iterations can be computed as

$$||r^n||^2 = \int_1^{k_\mathrm{max}} \left(\hat{\omega}_k^n \hat{r}_k^0\right)^2 dk \overset{(5)(19)}{\approx} \int_1^{Re^{3/4}} \hat{\omega}_k^{2n} Re^{2\tilde{\alpha}} k^{2\beta} dk, \tag{21}$$

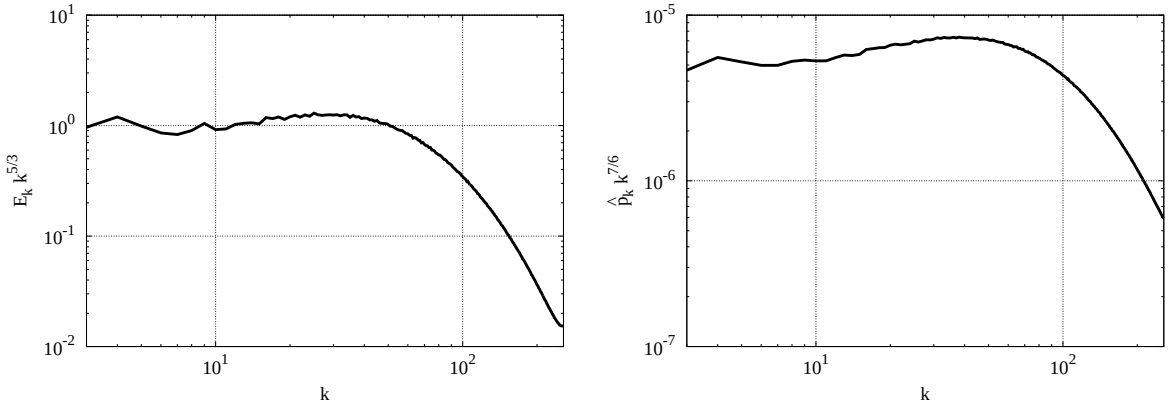

Figure 3: Compensated spectra for kinetic energy and pressure.

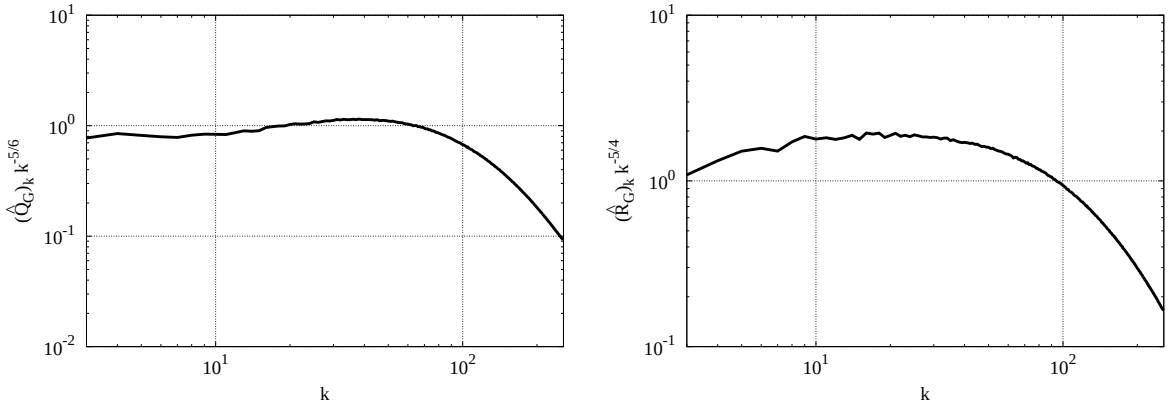

Figure 4: Compensated spectra for the second, Q_G, and third invariant, R_G, of the velocity gradient tensor.

where $\hat{\omega}_k = \hat{r}_k^{n+1}/\hat{r}_k^n$ is the convergence ratio of the solver and $\tilde{\alpha} = p\alpha$. For instance, for a Jacobi solver, $\hat{\omega}_k = \cos(\frac{\pi}{2}\rho)$ where $\rho \equiv k/k_{\max}$. In this case, using a quadratic approximation of $\cos(x) \approx 1 - 4x^2/\pi^2$ leads to

$$||r^n||^2 \approx \frac{Re^{2(\tilde{\alpha}+3/4(\beta+1/2))}}{2(2n+1)}. \tag{22}$$

This analysis can be extended for a multigrid solver (MG) with $l_{\max} \sim \log_2 N_x \sim (3/4)\log_2 Re$ levels and Jacobi as smoother

$$||r^n||^2 \approx \frac{Re^{2(\tilde{\alpha}+3/4(\beta+1/2))}}{2(2n+1)} \left\{ \left(\sum_{l=0}^{l_{\max}} \frac{(3/4)^{2n+1}}{2^{2l}} \right) + \frac{1}{2^{2l_{\max}+1}} \right\}. \tag{23}$$

Notice that in contrast to Eq.(22), MG's convergence is strongly accelerated by the term in brackets, which in the limit tends to $(3/4)^{2n}$. Nevertheless, the power law scaling with Re is the same; therefore, the regions defined in the $\{\tilde{\alpha}, \beta\}$ phase space remain unchanged (see Figure 2, right).

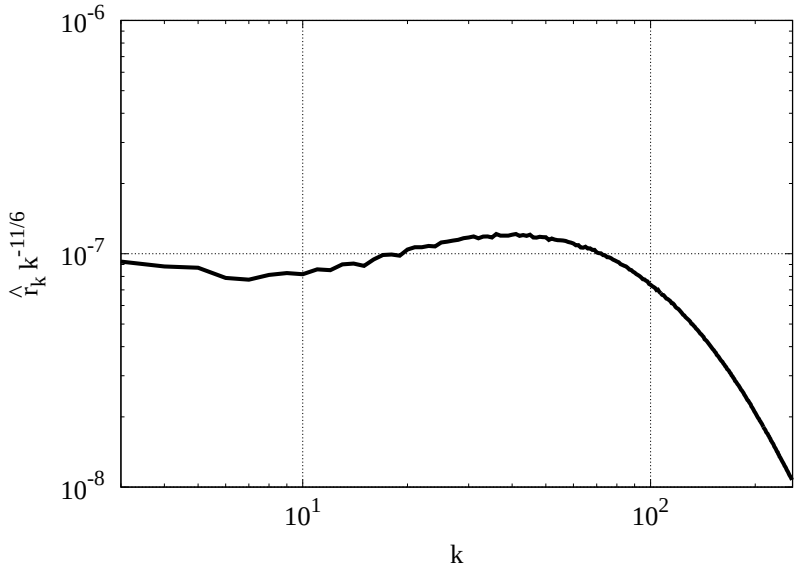

Figure 5: Compensated spectrum for the initial solver residual.

4. Numerical validation

The theory developed in previous sections relies on the validaty of the power-law scaling given in Eq.(19), which subsequently relies on a set of power-law scalings of the second, Q_{G}, and third invariant, R_{G}, of the velocity gradient tensor. First test-case to validate all these assumptions corresponds to a DNS of forced homogeneous isotropic turbulence at Taylor micro-scale Reynolds number $Re_\lambda = 325$ using a 512^3 mesh. Simulation has been carried out using the SpNS (pseudo-Spectral Navier–Stokes) code using the $3/2$ dealiasing rule. Energy is injected by fixing $E_k = 1$ for $||k|| = 1$. An explicit second-order Adams–Bashforth scheme is used for time-integration. The SpNS code is publicly available on GitHub[1].

Figure 3 displays the energy and pressure compensated spectra, showing a good agreement with the scaling predicted by the classical Kolmogorov theory [1]. Notice that there is a rather good agreement with the expected $-7/6$ power-law scaling for pressure [8]. Then, pressure and second invariant Q_{G} are directly linked via Eq.(14). This leads to the $5/6$ power-law scaling for Q_{G} which is observed in Figure 4 (left), and subsequently to the $5/4$ scaling of R_{G} which can be observed in Figure 4 (right). Finally, the combination of these scalings with the Taylor frozen-turbulence hypothesis (see Eq. 17) together with the restricted Euler equation (see Eq. 16) leads to the $11/6$ power-law scaling for the initial solver residual. The compensated spectrum for this residual displayed in Figure 5 seems to confirm our theory, *i.e.,* $\beta = 11/6$. Altogether leads to the preliminary conclusion that the number of iterations tends to increase with Re (see $\{\tilde{\alpha}, \beta\}$ phase space in Figure 2, right).

5. Concluding remarks

In this work, we aimed to shed light to the following research question: *is the complexity of numerically solving Poisson's equation increasing or decreasing for very large DNS and LES*

[1] Source code publicly available at `https://github.com/adalbal/SpNS`.

simulations of incompressible flows? Both physical and numerical arguments have been used to derive power law scalings at very high Reynolds numbers. Preliminary results of forced homogeneous isotropic turbulence have been carried out to confirm these scalings. Furthermore, theoretical convergence analysis for both Jacobi and multigrid solvers has defined a two-dimensional phase space divided into two regions depending whether the number of solver iterations tend to decrease or increase with the Reynolds number. Preliminary results seem to confirm that we are indeed in the second region, meaning that the complexity of solver Poisson's equation will increase. This prelimary finding may be useful to define roadmaps on the evolution of Poisson solvers for very large scale DNS and LES simulations.

Acknowledgements

This work was financially supported by two competitive R+D projects: RETOtwin (PDC2021-120970-I00), given by MCIN/AEI/ 10.13039/501100011033 and European Union Next GenerationEU, and FusionCAT (001-P-001722), given by *Generalitat de Catalunya* RIS3CAT-FEDER. Àdel Alsalti-Baldellou has also been supported by the predoctoral grants DIN2018-010061 and 2019-DI-90, given by MCIN/AEI/10.13039/501100011033 and the Catalan Agency for Management of University and Research Grants (AGAUR), respectively. Calculations have been carried out on the MareNostrum 4 supercomputer at BSC. We thankfully acknowledge these institutions.

References

1. Uriel Frisch. *Turbulence. The Legacy of A.N.Kolmogorov*. Cambridge University Press, 1995.
2. R. W. C. P. Verstappen and A. E. P. Veldman. Symmetry-Preserving Discretization of Turbulent Flow. *Journal of Computational Physics*, 187:343–368, 2003.
3. F. X. Trias, O. Lehmkuhl, A. Oliva, C.D. Pérez-Segarra, and R.W.C.P. Verstappen. Symmetry-preserving discretization of Navier-Stokes equations on collocated unstructured meshes. *Journal of Computational Physics*, 258:246–267, 2014.
4. F. Dabbagh, F. X. Trias, A. Gorobets, and A. Oliva. Flow topology dynamics in a three-dimensional phase space for turbulent Rayleigh-Bénard convection. *Physical Review Fluids*, 5:024603, 2020.
5. F. X. Trias, A. Gorobets, and A. Oliva. Turbulent flow around a square cylinder at Reynolds number 22000: a DNS study. *Computers & Fluids*, 123:87–98, 2015.
6. J. Kim and P. Moin. Application of a Fractional-Step Method to Incompressible Navier-Stokes Equations. *Journal of Computational Physics*, 123:308–323, 1985.
7. Stephen B. Pope. *Turbulent Flows*. Cambridge University Press, 2000.
8. D. I. Pullin. Pressure spectra for vortex models of fine-scale homogeneous turbulence. *Physics of Fluids*, 7:849, 1995.
9. B. J. Cantwell. Exact solution of a restricted Euler equation for the velocity gradient tensor. *Physics of Fluids A*, 4:782–793, 1992.