

Performance of high-order implicit large eddy simulations

Konstantinos Ritos*, Ioannis W. Kokkinakis, Dimitris Drikakis*

University of Strathclyde, Glasgow, G1 1XJ, UK



ARTICLE INFO

Article history:

Received 24 January 2017

Revised 30 October 2017

Accepted 23 January 2018

Available online 31 January 2018

Keywords:

iLES

High-Order methods

Turbulent flows

Parallel computing

ABSTRACT

The performance of parallel implicit Large Eddy Simulations (iLES) is investigated in conjunction with high-order weighted essentially non-oscillatory schemes up to 11th-order of accuracy. Simulations were performed for the Taylor Green Vortex and supersonic turbulent boundary layer flows on High Performance Computing (HPC) facilities. The present iLES are highly scalable achieving performance of approximately 93% and 68% on 1536 and 6144 cores, respectively, for simulations on a mesh of approximately 1.07 billion cells. The study also shows that high-order iLES attain accuracy similar to strict Direct Numerical Simulation (DNS) but at a reduced computational cost.

© 2018 The Authors. Published by Elsevier Ltd.

This is an open access article under the CC BY license. (<http://creativecommons.org/licenses/by/4.0/>)

1. Introduction

Implicit Large Eddy Simulations (iLES) originated from the observations made in [1] that the embedded dissipation of a certain class of numerical methods can be used in lieu of the explicit Sub-Grid Scale (SGS) models. Modified Equation Analysis (MEA) was developed [2] aiming at determining the stability of a difference equation by examining the truncation errors. Such an analysis was performed for the truncation error of certain schemes, e.g., [3–9]) leading to a better understanding of the implicit sub-grid dissipation. In iLES, the Navier–Stokes Equations (NSE) are discretised using high-resolution methods without involving a low-pass filtering operation, which leads to SGS terms that require additional modelling. Instead, only the (implicit) de facto filtering introduced through the finite volume integration of the NSE over the mesh cells is utilised in conjunction with non-linear numerical schemes that adhere to a number of principles; see [10,11], and reviews [9,12,13]. It has been shown [7] that iLES methods need to be carefully designed, optimised, and validated for the particular differential equation to be solved. Direct MEA of high-resolution schemes for NSE is difficult to be performed, thus understanding of the numerical properties of these methods to date still relies on performing computational experiments.

The use of iLES in free and wall-bounded flows has been justified by several authors [14,15], while a validation of the approach through theoretical analysis has been presented by Margolin et al. [8]. In incompressible flows, it is possible to develop an optimised

stencil with regards to numerical dissipation [16], however, in the case of compressible flows the numerical method should be monotonic with respect to the thermodynamic quantities. Poggie et al. [17] and Ritos et al. [18] applied compressible iLES to study Turbulent Boundary Layer (TBL) flows and showed that iLES can achieve close to strict DNS (see page 4 for definition of strict DNS) accuracy on significantly coarser meshes. Despite iLES (and similarly classical LES) being computationally less demanding than DNS, it still requires significant computational resources for simulating near-wall turbulence at high Reynolds numbers.

To date, there has been no systematic attempt to investigate the parallel scalability of different high-order compressible iLES methods in free and wall-bounded flows. The aim of this study is to present results regarding the accuracy, efficiency and parallel scalability of high-order iLES with reference to the Taylor Green Vortex (TGV) and supersonic TBL flows.

2. Numerical methods and flow cases

We have employed iLES in the framework of the CFD code CNS3D [12,15]. The Navier–Stokes equations are solved by using a finite volume Godunov-type method for the convective terms, which comprises the HLLC approximate Riemann solver [13,19] and two high-resolution schemes. The Monotone Upstream-centered Scheme for Conservation Laws (MUSCL) with second-order piecewise linear monotonicised central limiter [20] (labelled as M2), and the Weighted-Essentially Non-Oscillatory (WENO) ninth-order scheme [21] (labelled as W9). Furthermore, in order to examine the parallel scalability of high-order iLES, simulations were also performed using an eleventh-order WENO scheme (labelled as W11).

* Corresponding authors.

E-mail addresses: konstantinos.ritos@strath.ac.uk (K. Ritos), dimitris.drikakis@strath.ac.uk (D. Drikakis).

Table 1

Simulation parameters: u_∞ , T_∞ , M , P_∞ , ρ_∞ , and μ_∞ are the freestream velocity, temperature, Mach number, pressure, density and viscosity, respectively. T_w is the wall temperature, I is the turbulence intensity at the inlet and Re_L is the Reynolds number based on the freestream properties of air and the plate length, L .

L	u_∞	T_∞	M	P_∞
0.061 m	588 m/s	170 K	2.25	23.8 kPa
ρ_∞	T_w/T_∞	μ_∞	I	Re_L
0.488 kg/m ³	1.9	1.167×10^{-5} Pa s	3%	1.5×10^6

The viscous terms are discretised using a second-order central scheme. The solution is advanced in time using a five-stage (fourth-order accurate) optimal strong-stability-preserving Runge-Kutta method [22]. Further numerical details are provided in [15] and references therein.

The first flow case considered here is the TGV in a triple-periodic cubic domain of length 2π (m). A series of meshes was used: 32^3 , 64^3 , 128^3 , 256^3 and 512^3 evenly-spaced computational cells. The flow is initialised by solenoidal velocity profile,

$$\begin{aligned} u_0 &= U_0 \sin(kx) \cos(ky) \cos(kz), \\ v_0 &= -U_0 \cos(kx) \sin(ky) \cos(kz), \\ w_0 &= 0, \end{aligned} \quad (1)$$

and the pressure is obtained by solving the Poisson equation:

$$P_0 = P_\infty + \frac{1}{16} \rho_0 U_0^2 [2 + \cos(2kz)] \cdot [\cos(2kx) + \cos(2ky)], \quad (2)$$

where the wavenumber $k = 1$. An ideal gas equation of state is used and the Mach number, $U_0/\sqrt{\gamma P_0/\rho_0}$, is 0.08. The results are presented in terms of non-dimensional units; distance $x^* = kx$ and time $t^* = kU_0 t$.

The second flow case considered here is a supersonic turbulent flow over a flat plate at Mach number $M = 2.25$ and Reynolds number of 1.5×10^6 based on the freestream properties for air and the length of the plate, L ; see also Table 1.

Periodic boundary conditions are used in the spanwise direction (z). In the wall-normal direction (y) a no-slip isothermal wall at temperature $T_w = 323\text{K}$ is imposed. Supersonic outflow conditions are applied at the outlet, while far-field conditions are applied on the upper boundary opposite to the wall.

A synthetic turbulent inflow boundary condition is used to produce a freestream flow with a superimposed random turbulence. The synthetic turbulent inflow boundary condition is based on the digital filter (DF) method [18,23–25]. According to DF, instead of using a white-noise random perturbation at the inlet, energy modes within the Kolmogorov inertial range scaling with $k^{-5/3}$, where k is the wavenumber, are introduced into the turbulent boundary layer. A cutoff at the maximum frequency of 50 MHz is applied since the finest mesh would under-resolve higher frequency values. The turbulence intensity at the inlet (I) is set as $\pm 3\%$ of the intensity of the freestream velocity. This perturbation has been found to be sufficient to trigger bypass transition and turbulence downstream (see Fig. 1).

iLES have been performed on fine meshes but still coarser than DNS [17,26]. We employed four meshes with the coarsest and finest meshes containing ~ 4.5 million and ~ 100 million cells, respectively. For the calculation of the mesh spacing Δy the conventional inner variable scaling method $\Delta y^+ = u_\tau \Delta y / \nu_w$ is used, where $u_\tau = \sqrt{\tau_w / \rho_w}$ is the friction velocity; ν_w , τ_w and ρ_w are the wall viscosity, shear stress and density, respectively. Typical mesh resolution recommendations for LES lie in the range of $50 < \Delta x^+ < 150$ and $15 < \Delta z^+ < 40$, and for DNS in the range of $10 < \Delta x^+ < 20$ and $5 < \Delta z^+ < 10$ [17,27,28]. For wall-resolved LES and DNS the near-wall spacing should be $\Delta y^+ < 1$. A strict definition

Table 2

Boundary layer properties, including previous DNS and experimental studies. The compressible form of the momentum thickness (θ) has been used in the definition of Re_θ and Re_{δ_2} . Re_τ is the Reynolds number based on the friction velocity u_τ and the boundary layer thickness δ . Re_{δ_2} is based on θ and the near-wall viscosity μ_w . $H = \delta^* / \theta$ is the shape factor, where δ^* is the displacement thickness of compressible flow.

	Re_θ	Re_τ	Re_{δ_2}	H	M
W9	2170.0	414.0	1280.6	3.56	2.25
M2	1593.8	344.6	940.5	3.72	2.25
DNS [26]	2377.0	497.0	1516.0	2.98	2.0
strict DNS [17]	-	-	2000.0	-	2.25
Exp [29]	5100.0	1080.0	3100.0	2.00	2.28

for DNS mesh spacing requires $\Delta x^+ \leq 1$ and $\Delta y^+ \leq 1$. The mesh spacing used in this study is in the range of $9.06 < \Delta x^+ < 27.14$, $0.497 < \Delta y^+ < 1.22$ and $8.53 < \Delta z^+ < 24.95$, where the smallest values correspond to the finest mesh. Based on the above analysis, the present iLES on the finest mesh can be considered as an under-resolved DNS.

The TBL properties are presented in Table 2. To enable the comparison of the present results with other publications, various definitions of the Reynolds number have been employed based on the momentum thickness, the friction velocity, and the near-wall viscosity. The flow statistics are computed by averaging in time over three flow cycles and, spatially, in the spanwise direction. The statistical convergence of the simulations based on the standard error of the mean is less than 2%.

3. Results

3.1. TGV

Instantaneous visualisations of the TGV at $t^* = 15$ (Fig. 2) show the dominance of disorganised vortices in the decaying worm-vortex flow regime. The results were obtained using the ninth-order WENO scheme on 64^3 , 128^3 , 256^3 and 512^3 meshes. The snapshots of the flow are based on the Q -criterion, which defines a vortex as a continuous fluid region with a positive second invariant of the velocity gradient [30], i.e. $Q > 0$. All renderings are performed at the same level ($Q = 1$) and are coloured with the velocity magnitude.

The results on 256^3 and 512^3 meshes are very similar with respect to the turbulent structures resolved. The kinetic energy rate of dissipation, ε_1 , and pressure dilation-based dissipation rate, ε_2 , are shown in Fig. 3. The kinetic energy rate of dissipation is calculated by $\varepsilon_1 = -dE_k/dt$, where

$$E_k = \frac{1}{\rho_0 V} \int \frac{1}{2} \rho \mathbf{u} \cdot \mathbf{u} dV \quad (3)$$

is the volumetric-averaged kinetic energy. The simulations are nearly grid converged with respect to ε_1 and agree with other published results [31,32] (not shown here). The pressure dilatation-based dissipation rate is defined by

$$\varepsilon_2 = -\frac{1}{\rho_0 V} \int p \nabla \cdot \mathbf{u} dV. \quad (4)$$

ε_2 measures the effect of compressibility on the dissipation of turbulent energy and takes small values for low Mach flows.

A widely used performance metric for assessing parallel computations is the speedup:

$$S_n = \frac{T_{ref}}{T_n}, \quad (5)$$

where T_n is the execution time on n cores and T_{ref} is the execution time on a reference number of processors, usually equal to a single




Fig. 1. Iso-surfaces of Q -criterion, coloured by Mach number, for (a) M2 and (b) W9 iLES simulations; the computational domain has been truncated. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article).




Fig. 2. Iso-surfaces of Q -criterion ($Q = 1$) coloured by velocity magnitude at $t^* = 15$. The 32^3 mesh is not shown as no structure is visible at this level of Q . All shown TGV simulations are with W9. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article).

core or to the number of cores in a computational node of the HPC facility used. For the TGV simulations on meshes up to 512^3 cells, 12 cores were used as reference; one HPC node has two Intel Xeon X5650 processors with 6 cores each. The ideal speedup of parallel computations would be equal to n/n_{ref} , but this efficiency is not possible due the communication overhead between the computational cores and the idle time of computational nodes associated with load balancing. Fig. 4a shows the parallel speedup for the TGV case using the ninth-order iLES, achieving 77% speed-up using 480 cores. Furthermore, for scalability purposes the parallel performance of the eleventh-order WENO iLES on 6144 cores for the 1024^3 simulation is shown; a Cray HPC facility comprising two

Intel E5-2697v2 processors with 12 cores each was used. The reference execution time was obtained on 192 cores. The parallel performance of the 1024^3 simulation is approximately 93% and 68% on 1536 and 6144 cores, respectively. The parallel performance of the second-order iLES is not shown because it involves less calculations for the same mesh size and as a consequence the scalability will always be worse when comparing to higher order iLES.

3.2. TBL

The second flow case is a supersonic TBL for which DNS results and experimental data at similar Mach numbers are available.




Fig. 3. TGV case: (a) Kinetic energy dissipation rate (ϵ_1) and (b) pressure dilatation-based dissipation rate (ϵ_2). The results on the 512^3 mesh are in close agreement with previously published results [31,32] (not shown here). The y-axis in (b) is stretched by a factor of 10 compared to (a).




Fig. 4. Parallel scaling of (a) 9th- and (b) 11th-order iLES for the TGV case on 128^3 and 1024^3 meshes, respectively.




Fig. 5. Comparison of iLES with DNS and experimental data. (a) van Driest velocity profile (b) normal Reynolds stress. Strict DNS results [17] are included only in (b) because for the velocity profiles the results perfectly agree with the available DNS data of Pirozzoli et al. [26]. “LR” denotes a lower resolution mesh containing approximately 1/3 of the size of the fine mesh.

Comparisons with DNS and/or experiments are presented for the van Driest velocity profile, u_{VD} , and normal Reynolds stress, τ_{uu} (Fig. 5). The van Driest velocity profile is given by

$$u_{VD} = \int_0^{u^+} \sqrt{\frac{\rho}{\rho_w}} du^+, \quad (6)$$

where the superscript ‘+’ denotes wall scaling, $u^+ = u/u_\tau$. Previous publications [26,33] have shown that for adiabatic walls a satisfactory agreement of the velocity data is expected in near-wall region. Small variations are expected for different Reynolds number and the present iLES are in agreement with the DNS of Pirozzoli et al. [26]. The ninth-order iLES is also in excellent agreement with the experimental data [29]. The second-order iLES, conducted on the same mesh resolution, shows significant deviation from the

reference DNS and experiments. Performing the ninth-order iLES on 1/3 mesh resolution shows that mesh convergence is achieved, hence the high-order iLES reliably attain high accuracy on a relatively coarse mesh.

In respect of τ_{uu} , the second-order iLES significantly overestimate the Reynolds normal stress, especially in the peak region of the buffer zone. The ninth-order iLES show very good agreement with the DNS results up to about $y^+ \approx 20$, where the Reynolds similarity holds [34]. Further away from the wall it is typical to observe a strong dependence on Reynolds number for results presented in inner scaling coordinates. This explains the differences in the results in the logarithmic region due to the differences in the local Reynolds number.




Fig. 6. iLES speed-up for the 9th-order WENO scheme for the supersonic TBL case.

Table 3
Performance of second and ninth order iLES vs DNS [17,26].

Method	Computational cost	Error 1	Error 2
Strict DNS	~25 (years)	0.0%	0.0%
DNS	117 (days)	0.0%	6.5%
iLES W9	24 (days)	1.0%	6.3%
iLES M2	10 (days)	8.5%	23.7%
iLES W9 - LR	7 (days)	3.1%	8.1%

For the TBL case the speedup is calculated with reference to 36 cores (3 computational nodes with two Intel Xeon X5650 processors each node). The 36 cores reference was imposed due to the size of the fine mesh (~ 100 million cells). The parallel speedup is shown in Fig. 6. The ninth-order iLES provide a computational efficiency of 86% of the ideal efficiency, utilising 720 computational cores.

Table 3 shows the performance of low and high order iLES with reference to strict DNS [17]. For the DNS performance we have used the results of Pirozzoli et al. [26], where a mesh approximately 27 times coarser than the strict definition of DNS was utilised. The reported errors are averaged values calculated in the near-wall region, $y^+ \leq 30$, where “Error1” and “Error2” refer to relative difference from the reference van Driest velocity profile and normal wall stress, respectively. The computational cost has been calculated based on the assumption that simulations are performed on 240 cores. The computational cost for DNS are estimations based on the mesh size and number of cores found in the relevant publications. The results show that high-order iLES can attain high accuracy at a reduced computational cost, cf. iLES W9-LR with the rest of the results.

4. Conclusions

The accuracy, parallel scalability and efficiency of iLES were examined for different turbulent flow cases. A mesh convergence study was presented for the TGV case achieving nearly mesh independent results for the two finest meshes. The present high-order iLES exhibit high parallel efficiencies for simulations performed up to 6144 cores on a Cray HPC facility and for meshes containing up to 1.07 billion cells.

The first and second order statistics obtained from high-order iLES of a supersonic TBL flow are in excellent agreement with previous numerical and experimental data. iLES can achieve high accuracy in the near-wall region that is directly comparable to the results of strict DNS at a reduced computational cost. A combination of high-order iLES with relatively coarse meshes provides a

more pragmatic approach than using a second order method on a significantly finer mesh.

Acknowledgements

Results were obtained using the EPSRC funded ARCHIE-WeSt High Performance Computer (www.archie-west.ac.uk) under EPSRC grant no. EP/K000586/1. The authors would also like to thank EPSRC for providing access to computational resources on the National HPC facility ARCHER through the UK Applied Aerodynamics Consortium Leadership Project “e529”.

References

- [1] Boris J, Grinstein FF, Oran E, Kolbe R. New insights into large eddy simulation. *Fluid Dyn Res* 1992;10(4–6):199–228.
- [2] Hirt C. Heuristic stability theory for finite-difference equations. *J Comput Phys* 1968;2(4):339–55.
- [3] Margolin LG, Rider WJ. A rationale for implicit turbulence modelling. *Int J Numer Methods* 2002;39(9):821–41.
- [4] Rider WJ, Margolin LG. From numerical analysis to implicit subgrid turbulence modeling. In: 16th AIAA computational fluid dynamics conference; 2003. p. 1–11.
- [5] Drikakis D, Rider WJ. High-resolution methods for incompressible and low-speed flows, 1. Springer-Verlag; 2004.
- [6] Margolin LG, Rider WJ. The design and construction of implicit LES models. *Int J Numer Methods* 2005;47(10–11):1173–9.
- [7] Domaradzki JA, Radhakrishnan S. Effective eddy viscosities in implicit modeling of decaying high Reynolds number turbulence with and without rotation. *Fluid Dyn Res* 2005;36(4–6):385–406.
- [8] Margolin LG, Rider WJ, Grinstein FF. Modeling turbulent flow with implicit LES. *J Turbul* 2006;7(15).
- [9] Grinstein FF, Margolin LG, Rider WJ. *Implicit large eddy simulation: computing turbulent fluid dynamics*. Cambridge University Press; 2007.
- [10] Harten A. High resolution schemes for hyperbolic conservation laws. *J Comput Phys* 1983;49(3):357–93.
- [11] Harten A. High resolution schemes for hyperbolic conservation laws. *J Comput Phys* 1997;135(2):260–78.
- [12] Drikakis D. Advances in turbulent flow computations using high-resolution methods. *Prog Aerosp Sci* 2003;39(6–7):405–24.
- [13] Toro EF. *Riemann solvers and numerical methods for fluid dynamics*. 3rd. Springer; 2009.
- [14] Fureby C, Grinstein FF. Large eddy simulation of high-Reynolds-number free and wall-bounded flows. *J Comput Phys* 2002;181(1):68–97.
- [15] Drikakis D, Hahn M, Mosedale A, Thornber B. Large eddy simulation using high resolution and high order methods. *Philos Trans Royal Soc A* 2009;367:2985–97.
- [16] Hickel S, Adams NA, Domaradzki JA. An adaptive local deconvolution method for implicit LES. *J Comput Phys* 2006;213(1):413–36.
- [17] Poggie J, Bisek NJ, Gosse R. Resolution effects in compressible, turbulent boundary layer simulations. *Comput Fluids* 2015;120:57–69.
- [18] Ritos K, Kokkinakis IW, Drikakis D, Spottswood SM. Implicit large eddy simulation of acoustic loading in supersonic turbulent boundary layers. *Phys Fluids* 2017;29(4):1–11.
- [19] Toro EF, Spruce M, Speares W. Restoration of the contact surface in the HLL-Riemann solver. *Shock Waves* 1994;4(1):25–34.
- [20] van Leer B. Towards the ultimate conservative difference scheme III. Upstream-centered finite-difference schemes for ideal compressible flow. *J Comput Phys* 1977;23(3):263–75.
- [21] Balsara DS, Shu CW. Monotonicity preserving weighted essentially non-oscillatory schemes with increasingly high order of accuracy. *J Comput Phys* 2000;160(2):405–52.
- [22] Spiteri R, Ruuth SJ. New class of optimal high-order strong-stability-preserving time discretization methods. *SIAM J Numer Anal* 2002;40(2):469–91.
- [23] Lund TS, Wu X, Squires KD. Generation of turbulent inflow data for spatially-developing boundary layer simulations. *J Comput Phys* 1998;140(2):233–58.
- [24] Klein M, Sadiki A, Janicka J. A digital filter based generation of inflow data for spatially developing direct numerical of large eddy simulations. *J Comput Phys* 2003;186(2):652–65.
- [25] Toubert E, Sandham ND. Large-eddy simulation of low-frequency unsteadiness in a turbulent shock-induced separation bubble. *Theor Comput Fluid Dyn* 2009;23(2):79–107.
- [26] Pirozzoli S, Bernardini M. Turbulence in supersonic boundary layers at moderate Reynolds number. *J Fluid Mech* 2011;688:120–68.
- [27] Georgiadis NJ, Rizzetta DP, Fureby C. Large-eddy simulation: current capabilities, recommended practices, and future research. *AIAA J* 2010;48(8):1772–84.
- [28] Choi H, Moin P. Grid-point requirements for large eddy simulation: Chapman’s estimates revised. *Phys Fluids* 2012;24(1):011702.
- [29] Piponniau S, Dussauge JP, Débiève JF, Dupont P. A simple model for low-frequency unsteadiness in shock-induced separation. *J Fluid Mech* 2009;629:87–108.
- [30] Kolář V. Vortex identification: new requirements and limitations. *Int J Heat Fluid Flow* 2007;28:638–52.

- [31] Debonis J. Solutions of Taylor–Green vortex problem using high-resolution explicit finite difference methods. In: 51st AIAA Aerospace Sciences Meeting; 2013. p. 1–9.
- [32] Bull JR, Jameson A. Simulation of the Taylor–Green vortex using high-order flux reconstruction schemes. *AIAA J* 2015;53(9):2750–61.
- [33] Smits AJ, Dussauge JP. Turbulent shear layers in supersonic flow. 2nd. American Institute of Physics; 2006.
- [34] Fernholz HH, Finley PJ. The incompressible zero-pressure-gradient turbulent boundary layer: and assessment of the data. *Prog Aerosp Sci* 1996;32:245–311.