FLOW CAMBRIDGE
UNIVERSITY PRESS

RESEARCH ARTICLE

Large eddy simulation of aircraft at affordable cost: a milestone in computational fluid dynamics

Konrad A. Goc[1] (iD), Oriol Lehmkuhl[2], George Ilhwan Park[3,*] (iD), Sanjeeb T. Bose[4,5] and Parviz Moin[1] (iD)

[1]Center for Turbulence Research, Stanford University, Stanford, CA 94305, USA
[2]Barcelona Supercomputing Center, Barcelona 08034, Spain
[3]Department of Mechanical Engineering and Applied Mechanics, University of Pennsylvania, Philadelphia, PA 19104, USA
[4]Cascade Technologies, Inc., Palo Alto, CA 94303, USA
[5]Institute for Computational & Mathematical Engineering, Stanford University, Stanford, CA 94305, USA
*Corresponding author. E-mail: gipark@seas.upenn.edu

Abstract

While there have been numerous applications of large eddy simulations (LES) to complex flows, their application to practical engineering configurations, such as full aircraft models, have been limited to date. Recently, however, advances in rapid, high quality mesh generation, low-dissipation numerical schemes and physics-based subgrid-scale and wall models have led to, for the first time, accurate simulations of a realistic aircraft in landing configuration (the Japanese Aerospace Exploration Agency Standard Model) in less than a day of turnaround time with modest resource requirements. In this paper, a systematic study of the predictive capability of LES across a range of angles of attack (including maximum lift and post-stall regimes), the robustness of the predictions to grid resolution and the incorporation of wind tunnel effects is carried out. Integrated engineering quantities of interest, such as lift, drag and pitching moment will be compared with experimental data, while sectional pressure forces will be used to corroborate the accuracy of the integrated quantities. Good agreement with experimental C_L data is obtained across the lift curve with the coefficient of lift at maximum lift, $C_{L,max}$, consistently being predicted to within five lift counts of the experimental value. The grid point requirements to achieve this level of accuracy are reduced compared with recent estimates (even for wall modelled LES), with the solutions showing systematic improvement upon grid refinement, with the exception of the solution at the lowest angles of attack, which will be discussed later in the text. Simulations that include the wind tunnel walls and aircraft body mounting system are able to replicate important features of the flow field noted in the experiment that are absent from free air calculations of the same geometry, namely, the onset of inboard flow separation in the post-stall regime. Turnaround times of the order of a day are made possible in part by algorithmic advances made to leverage graphical processing units. The results presented herein suggest that this combined approach (meshing, numerical algorithms, modelling, efficient computer implementation) is on the threshold of readiness for industrial use in aeronautical design.

Impact Statement

The use of computational fluid dynamics for external aerodynamic applications has been a key tool for aircraft design in the modern aerospace industry. In take-off and landing configurations, predicting the maximum lift an aircraft can produce and the associated onset of boundary layer separation encountered at high angles

CrossMark

of attack is critically important. Flow solutions from state-of-the-art solvers are unable to routinely comply with the stringent accuracy and computational efficiency requirements demanded by industry. Leveraging large eddy simulation with appropriate wall/subgrid-scale models and low dissipation numerical methods suitable for complex geometries on modern computer architectures offers a tractable path towards meeting these accuracy and affordability requirements.

1. Introduction

The present investigation was in part motivated by one of the Grand Challenge Problems in computational aerosciences identified in the NASA 2030 Computational Fluid Dynamics (CFD) Vision Report (Slotnick et al., 2014): 'LES of a powered aircraft configuration across the full flight envelope'. Present use of CFD in the engineering design process has had considerable success in predicting flows at cruise conditions. These operating points are largely characterized by attached boundary layers where the turbulence models used in traditional CFD approaches (often, but not exclusively based on Reynolds-averaged Navier–Stokes (RANS) closures) are relatively accurate. Other operating conditions, such as take-off or landing, experience more complex and unsteady phenomena including boundary layer separation. Notably, the NASA 2030 CFD Vision Report lists the prediction of separated flows as a present deficiency of lower fidelity (e.g. RANS-based) solution approaches and suggests that unsteady simulation techniques such as large eddy simulations (LES) may provide sufficient accuracy. The target date for an LES-based 'technology demonstration' was estimated to be in the late 2020s.

Computational cost is the principal uncertainty of when higher fidelity simulation techniques can be adopted. There is little doubt that first-principles-based, direct numerical simulation would provide accurate predictions, but the resolution requirements at high Reynolds numbers would be prohibitive. Wall modelled approaches for LES in which the ensemble effect of the unresolved near-wall structures are modelled while eddies that scale with the boundary layer thickness are directly resolved (Bose & Park, 2018; Cabot & Moin, 2000; Larsson, Kawai, Bodart, & Bermejo-Moreno, 2016; Piomelli & Balaras, 2002; Wang & Moin, 2002) can ameliorate these resolution requirements. While LES wall models have been constructed from the equilibrium boundary layer assuming a constant stress layer, analysis of the von Karman momentum integral suggests that these approximations should be valid if sufficient resolution of the boundary layer thickness, δ, is provided (specifically the displacement and momentum thicknesses) (Bose & Park, 2018). The error in salient quantities of interest (e.g. the overall lift coefficient) as a function of boundary layer thickness resolution for a complex flow, such as an aircraft is, however, not well established. The relationship between error and boundary layer thickness resolution is likely to be a strong function of the numerical methods and subgrid-scale models utilized when marginal resolution of the boundary layer thickness is provided (1–10 grid points per δ) (Lozano-Durán, Bose, & Moin, 2021).

This paper presents a summary of simulations performed using a realistic aircraft model in landing configuration. The selected model problem is the Japanese Exploration Agency (JAXA) Standard Model (JSM) with the objective of demonstrating the desired level of capability and cost effectiveness of LES for the prediction of aeronautical flows of engineering significance. This configuration was selected due to the recent interest garnered by its featuring in the AIAA Third High-Lift Prediction Workshop (HLPW3) where 35 participants submitted a total of 79 data sets of CFD results predicting the integrated forces and moments across the lift curve. Except for a few participants who used unsteady techniques (unsteady RANS, Coder, Pulliam, and Jensen (2019); lattice Boltzmann very large-eddy simulations, Konig, Fares, Murayama, and Ito (2018); delayed detatched-eddy simulation, Cary, Yousuf, Li, and Mani (2018) and Balin and Jansen (2018)), most calculations presented in the workshop were steady RANS simulations deploying variants of the Spalart–Allmaras or Shear Stress Transport models for the Reynolds stress closure (Rumsey, Slotnick, & Sclafani, 2019). Figure 1 shows these results compiled onto a single lift versus angle of attack curve. A critical observation from this plot is that while good clustering of results is found in the linear lift range, a significant data scatter is observed near maximum
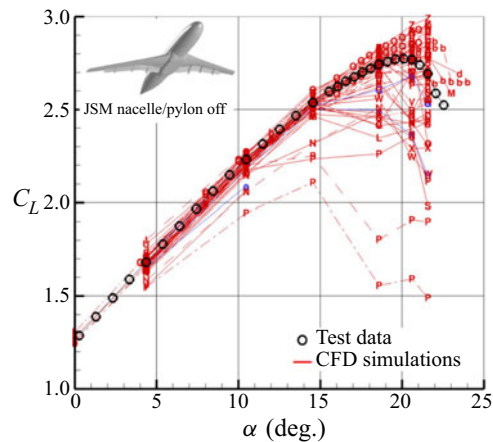
Figure 1. *Collection of the results of 30+ participants' data submissions to the HLPW3. A variety of solvers and gridding strategies are represented, however, the majority of the participants used steady RANS tools to generate their results. As a whole, the predictive capability is better in the linear range of the lift curve than near the point of maximum lift, where the results show appreciable scatter. A detailed legend for the data series is provided in* Rumsey et al. (2019).

lift, tempering confidence in the predictive capability of steady RANS when flow separation becomes significant.

Calculations described herein demonstrate accurate prediction of engineering quantities of interest at affordable cost well ahead of the 'technology demonstration' target date of the late 2020s set by the CFD Vision report in the 'technology development roadmap' it proposes. These calculations are an important step towards completing the grand challenge problem of a wall-resolved (large eddy) simulation of a powered aircraft at flight Reynolds number. The computational cost and predictive accuracy of these calculations (e.g. $\Delta C_L \leq 0.03$ and $\Delta C_D \leq 0.0010$ at maximal lift) are on the threshold of industrial readiness (Clark, Slotnick, Taylor, & Rumsey, 2020), particularly for our solution on the finest grid in the wind tunnel. Experimental repeatability values for the JAXA measurements at maximum lift are $C_L = 0.03$ and $C_D = 0.0015$. The preliminary calculations conducted in 2018 at the Center for Turbulence Research (CTR) Summer Program (Lehmkuhl, Park, Bose, & Moin, 2018) were targeted near the maximum lift conditions and indicated that acceptable levels of accuracy were attainable at resolutions well below the expectations of the CFD community (Choi & Moin, 2012; Slotnick et al., 2014; Spalart, 1997). We reinforce this view in this paper by demonstrating an expanded set of calculations spanning a range of angles of attack pertaining to both the linear lift and the stalled regimes. These simulations demonstrate that the agreement with experiment achieved by Lehmkuhl et al. (2018) was not fortuitous. A flurry of LES applications to complex three-dimensional aircraft flows have opened up in the wake of the 2018 CTR Summer Program, including the works by Ghate et al. (2021), Angelino, Fernández-Yáñez, Xia, and Page (2020), Lozano-Duran, Bose, and Moin (2020) and Iyer and Malik (2020), which have continued to highlight the promise of the predictive capabilities of LES technology in realistic settings.

The following sections remark on several potential reasons as to why these calculations have come to fruition earlier than originally anticipated. First, prior resolution estimates were not sensitized to particular quantities of interest (e.g. mean forces and moments). For instance, the requirements to accurately predict Reynolds stress distributions or wall pressure fluctuations are more stringent than the requirements to accurately predict integrated forces or average surface pressures (Park & Moin, 2016). In fact, there is some evidence that compulsory grid resolution is in part driven by capturing inviscid phenomena such as strong acceleration experienced at the airfoil leading edges at high angles of attack. Severe resolution requirements from the need to resolve thin, laminar boundary layers and

turbulent transition mechanisms also appear to be reduced, particularly at high angles of attack, where abrupt transition is observed near leading edges. Second, the deflation of computational cost due to the advent of accelerated computing architectures (e.g. graphical processing units (GPUs)) may bridge the remaining gap in computational cost associated with higher fidelity simulations. From the cases presented herein, algorithmic advances on modern computer hardware can deliver tractable turnaround times (within a day) and at costs (in a monetary sense) that can be an order of magnitude less than current central processing unit (CPU)-based supercomputing platforms.

The remainder of the manuscript is organized as follows. Section 2 provides a general description of the LES formulation (governing equations and wall models) and the flow solvers used. Section 3 details the results from the LES calculations, primarily in terms of the global force and surface pressure predictions compared with the available experimental measurements. A detailed discussion on the computational costs is provided in § 4, followed by concluding remarks in § 5.

2. Governing equations, physical models and numerical methods

2.1. *LES governing equations*

The governing equations for the low-pass filtered, compressible Navier–Stokes for mass, momentum and total energy are given by

$$\frac{\partial \bar{\rho}}{\partial t} + \frac{\partial \bar{\rho}\tilde{u}_j}{\partial x_j} = 0, \tag{2.1}$$

$$\frac{\partial \bar{\rho}\tilde{u}_i}{\partial t} + \frac{\partial \bar{\rho}\tilde{u}_i\tilde{u}_j}{\partial x_j} + \frac{\partial \bar{P}}{\partial x_i} = \frac{\partial \tilde{\sigma}_{ij}}{\partial x_j} - \frac{\partial \tau_{ij}^{sgs}}{\partial x_j}, \tag{2.2}$$

$$\frac{\partial \bar{E}}{\partial t} + \frac{\partial [(\bar{E}+\bar{P})\tilde{u}_j]}{\partial x_j} = \frac{\partial}{\partial x_j}\left(k\frac{\partial \tilde{T}}{\partial x_j}\right) + \frac{\partial(\tilde{u}_i\tilde{\sigma}_{ij})}{\partial x_j} - \frac{\partial Q_j^{sgs}}{\partial x_j} - \frac{\partial(\tilde{u}_i\tau_{ij}^{sgs})}{\partial x_j}, \tag{2.3}$$

where ρ, P, T and u_i refer to the fluid density, pressure, temperature and velocity vector, respectively. The total energy is defined as $E = \bar{\rho}\tilde{e} + \bar{\rho}\tilde{u}_k\tilde{u}_k/2$, $\tilde{\sigma}_{ij} = \mu(\tilde{T})(\tilde{S}_{ij} - 1/3\tilde{S}_{kk}\delta_{ij})$ is the deviatoric part of the resolved stress tensor, and $\tilde{S}_{ij} = 1/2(\partial \tilde{u}_i/\partial x_j + \partial \tilde{u}_j/\partial x_i)$ is the resolved strain rate tensor. The $(\bar{\cdot})$ and $(\tilde{\cdot})$ symbols refer to the Reynolds and Favre averaging operators, respectively. The subgrid-scale terms, indicated by the superscript *sgs*, account for the stress and heat flux of unresolved eddies on the scales resolved by the grid and are defined, respectively, as

$$\tau_{ij}^{sgs} = \bar{\rho}(\widetilde{u_iu_j} - \tilde{u}_i\tilde{u}_j), \tag{2.4}$$

$$Q_j^{sgs} = \bar{\rho}(\widetilde{eu_j} - \tilde{e}\tilde{u}_j). \tag{2.5}$$

In the present calculations, the static coefficient Vreman eddy viscosity model (Vreman, 2004) is used for the JSM free air cases, while the dynamic Smagorinsky model (Germano, Piomelli, Moin, & Cabot, 1991; Moin, Squires, Cabot, & Lee, 1991) is employed for the cases with the wind tunnel and nacelle. The JSM experiment is conducted with a free stream Mach number of 0.172 and a $Re_c = 1.96 \times 10^6$ (based on the mean aerodynamic chord and the free stream velocity). At these conditions, a calorically perfect gas equation of state is adopted for air ($Pr = 0.70$, $\gamma = 1.4$). The flow exhibits strong acceleration at the leading edges (particularly on the slat at higher angles of attack near maximum lift, where peak Mach numbers of $M \approx 0.75$ are observed), but overall the flow regime is characterized by low Mach numbers. As a result, some calculations shown below are obtained based on low Mach (incompressible) approximations of the governing equations ((2.1)–(2.3)) and will be noted as such. Although the flow solver 'Alya' is outside of its theoretical range of applicability in some restricted areas of the flow field (i.e. leading edges of the slat, main element and flap), we have found only minor differences between the compressible and incompressible solvers in this study.

2.2. Near-wall LES modelling

Viscous length scales near a no-slip wall introduce resolution requirements that scale with $l_\nu = \nu/u_\tau = \nu\sqrt{\rho}/\sqrt{\tau_w}$, where $u_\tau = \sqrt{\tau_w}/\sqrt{\rho}$, ν, ρ and τ_w denote the friction velocity, kinematic viscosity, fluid density and stress at the wall, respectively. As the Reynolds number (based on the mean aerodynamic chord length of the wing) increases, scale separation between the boundary layer thickness, δ, and l_ν increases, which can be characterized by a friction Reynolds number, $Re_\tau = \delta/l_\nu$. The high Reynolds numbers observed for flight vehicles ($Re_\tau \approx 5000$ for the JSM based on peak skin friction and 99 % boundary layer thickness predicted by the LES simulations) imposes prohibitive resolution requirements if l_ν is directly resolved. This resolution requirement is exacerbated in LES approaches (compared with RANS) as this viscous length scale must be resolved along each dimension as energetic near-wall turbulent structures also scale with respect to the viscous unit. Several estimates of the required number of grid points have been made for realistic aircraft based on experimental correlations for the skin friction (Choi & Moin, 2012) or RANS solutions (Spalart, 1997) suggesting that direct resolution would be untenable using available supercomputers in the foreseeable future. These grid point requirements are greatly reduced if only the eddies that scale with the boundary layer thickness are resolved and the aggregate effect of the near-wall viscous region is modelled, referred to as a wall modelled LES (WMLES). In this limit, the grid point count for the turbulent flow regime would scale linearly with respect to the chord-based Reynolds number, $N_{grid} \propto Re_c$. Further reduction in the overall computational cost can be achieved by efficient clustering of grid points (e.g. boundary-layer-confirming grids (Lozano-Durán et al., 2021)) and significant reduction in the time step size (Yang & Griffin, 2021), in particular for compressible flow solvers deploying explicit time-marching schemes that are limited by the acoustic CFL condition. Overall, WMLES of high-Reynolds-number applications would be potentially feasible using present computing capacity (e.g. Summit at the US Department of Energy's Oak Ridge National Laboratory).

The most common approach for the near-wall modelling is to assume that the thin boundary layer equations are valid. Assuming a local balance between the pressure gradient/streamwise convection and that the spatiotemporal resolution of the LES is large compared with viscous length and time scales such that the near-wall cell can be regarded as statistically stationary, then the momentum equations simplify to a constant stress layer in the wall normal direction,

$$\frac{d\tau}{dn} = \frac{d}{dn}\left(\mu\frac{dU}{dn} - \overline{\rho u'v'}\right) = 0, \tag{2.6}$$

subject to a no-slip condition ($U(n = 0) = 0$) at the wall and coupled to the LES solution at some distance from the wall ($U(n = n_{match}) = U_{LES}$). Admitting a closure for the Reynolds shear stress (a turbulent eddy viscosity based on a Prandtl mixing length), (2.6) can be solved for the wall stress. This stress is then provided as the wall boundary condition for the LES. Similar approximations can be made of the total energy equation to produce an approximation for the wall heat flux, q_w. This approximation is referred to as the equilibrium wall model and is adopted for the present studies. A schematic of the wall model/LES solution coupling methodology is sketched in figure 2. For additional details regarding the construction and implementation of wall models for LES, the reader is referred to recent reviews (Bose & Park, 2018; Larsson et al., 2016; Piomelli & Balaras, 2002). Details of the particular wall modelling methodologies employed in the present study can be found in Lehmkuhl et al. (2018).

While this investigation does not seek to improve state-of-the-art wall models for LES, it does provide important data regarding outstanding questions with respect to the analysis above. First and foremost, grid point estimates were based on the resolution requirements for the energetics of near-wall turbulence. To leading order, the prediction of aerodynamic loading (lift, drag, moments) are of paramount significance for simulations of a full scale flight vehicle. These quantities of interest can potentially have different resolution requirements. There are additional effects that carry resolution requirements: compulsory resolution of geometric features (e.g. slat brackets, flap support fairings, slat gap); inviscid effects (leading edge acceleration); and finite span phenomena (e.g. tip vortices). Second,
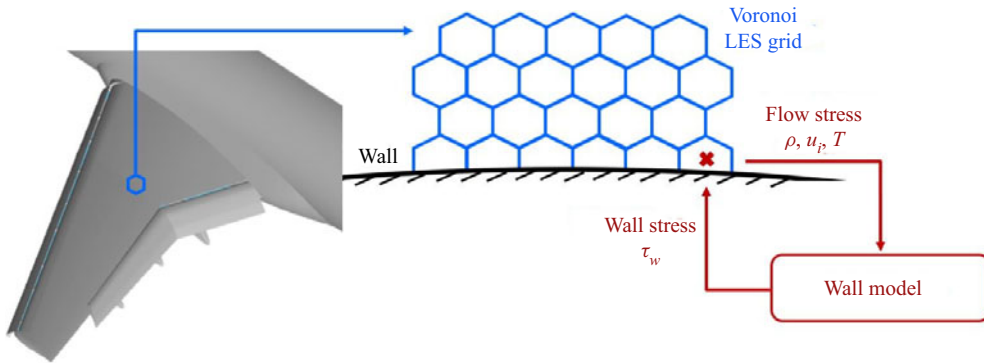
Figure 2. *Schematic showing the procedure for wall flux modelling in LES. The LES solution (blue) supplies the top boundary condition to an auxiliary wall model (red), whose role is to deliver wall fluxes to the LES in the form of a Neumann boundary condition coming from a solution to the wall model equations.*

many assumptions in the thin boundary layer equations are violated; there are strong pressure gradient effects and the conditions around maximal lift and post-stall regime include boundary layer separation. There is evidence from *a posteriori* wall modelled LES calculations that suggest that the equilibrium approximations are sufficient for flows over airfoils/aircraft models (Bodart, Larsson, & Moin, 2013; Park, 2017; Wang & Moin, 2002). It can be shown that the errors in the wall stress prediction by an equilibrium approximation can be bounded by the resolution of the LES matching location with respect to momentum/displacement thickness (y/δ^* or y/θ) even in the presence of pressure gradients (Bose & Park, 2018) supporting the *a posteriori* observations. This relies on the fact that the outer LES is capable of directly resolving the non-equilibrium effects in the outer regions of the boundary layer. Nevertheless, it will be seen below that extrapolating the grid resolutions utilized in these prior studies or strictly relying on the theoretical guidance would provide more stringent requirements than what was utilized to capture the present quantities of interest.

2.3. Flow solvers and numerical methods

Two different flow solvers are utilized for investigations of the JSM configuration: charLES and Alya. The flow solver charLES is a massively parallel, finite volume, compressible flow solver. It utilizes a low dissipation spatial discretion based on principles of discrete entropy preservation (Chandrashekar, 2013; Honein & Moin, 2004; Tadmor, 2003) where the fluxes are constructed to globally conserve entropy in an inviscid, shock-free flow and conserve kinetic energy in an inviscid, low Mach number regime. The numerical scheme has been shown to be suitable for coarsely resolved large eddy simulations of turbulent flows that are especially sensitive to numerical dissipation. The discretization is suitable for arbitrary unstructured, polyhedral meshes, and the solutions contained herein are computed using unstructured grids based on Voronoi diagrams. The use of Voronoi diagram-based meshes allows for the rapid generation of high quality grids with some guaranteed properties (for instance, the vector between two adjacent Voronoi sites is parallel to the normal of the face that they share). Time advancement is performed using a three stage explicit Runge–Kutta scheme (Gottlieb, Shu, & Tadmor, 2001), and the spatial discretization is formally second-order accurate. The code has been extensively validated, and recent applications of the code could be found in Bres et al. (2018), Lozano-Duran et al. (2020) and Fu, Karp, Bose, Moin, and Urzay (2021).

Alya is a parallel, multiscale simulation code developed at the Barcelona Supercomputing Centre (Vazquez et al., 2016). The formulation used in the calculations described herein is purely incompressible. The convective term is discretized using a Galerkin finite element scheme recently proposed

in Charnyi, Heister, Olshanskii, and Rebholz (2017), which conserves linear/angular momenta and kinetic energy at the discrete level. Both second- and third-order spatial discretizations are used. Neither upwinding nor any equivalent momentum stabilization is employed. In order to use equal-order elements, numerical dissipation is introduced only for the pressure stabilization via a fractional step scheme (Codina, 2001), which is similar to those approaches used for pressure–velocity coupling in unstructured, collocated finite-volume codes (Jofre, Lehmkuhl, Ventosa, Trias, & Oliva, 2014). The set of equations is integrated in time using a third-order Runge–Kutta explicit method combined with an eigenvalue-based time step estimator (Trias & Lehmkuhl, 2011). This approach is significantly less dissipative than the traditional stabilized finite element method approach (Lehmkuhl, Houzeaux, Owen, Chrysokentis, & Rodriguez, 2019).

One distinguishing feature of both numerical approaches is an emphasis on low-dissipation schemes. Numerical dissipation is known to be extremely detrimental to the simulation of turbulent flows (Mittal & Moin, 1997), particularly when the resolution is coarse with respect to integral scales of turbulence as it is in the present investigations. In complex geometries and on unstructured grids, the limitation of dissipation while maintaining stable numerical solutions is not trivial or commonplace. Many existing unsteady flow solvers use discretizations inherited from steady RANS approaches where solutions are not as sensitive to numerical dissipation or where dissipation can aid convergence to steady state solutions.

3. Computational results

We detail the set-up and results for the simulations of the JSM, which was the focus of the AIAA HLPW3. Computational results are compared with the experimental data originally collected by Yokokawa, Murayama, Ito, and Yamamoto (2006) and later expanded upon by Yokokawa et al. (2008, 2010). Two different configurations are considered: a JSM model in free air and a case where the JSM model is within a domain that represents the geometry of the JAXA experimental campaign (i.e. includes wind tunnel walls and sidewall offset). The free air simulation results for the integrated lift, drag and moments are compared with the experimental measurements that have been corrected to account for the blockage and boundary layer effects due to the wind tunnel walls. Comparisons are also made with experimental oil film visualizations and surface pressure measurements, but no wind tunnel corrections are available for the surface data. Simulations including the wind tunnel geometry facilitate more direct comparisons with the raw experimental measurements. The computational geometry and domain is shown alongside the experimental model in figure 3, a slice depicting the mesh utilized from the charLES and Alya flow solvers are shown in figure 4. The grid shown for the charLES solution corresponds to a grid count of 32 million control volumes, which was sized to fit approximately 10 points across the trailing edge boundary layer thickness (in all three directions) of the main airfoil element as predicted by the turbulent zero pressure gradient correlation based on the mean aerodynamic chord ($\delta_{99}/\Delta \approx 10$). The resolution does not vary in the streamwise direction, meaning that there are fewer points per boundary layer thickness towards the leading edge of the wing where the boundary layers are thinner. The majority of grid points for the baseline resolution for the charLES simulation are clustered near the leading edge of the slat to capture the flow acceleration, particularly at high angles of attack (it is worth noting that this effect is predominantly an inviscid phenomenon). In addition to better resolution of inviscid effects, laminar-to-turbulent transition effects may also be at play in the leading edge regions. The increased grid resolution improves the ability of the solution to capture transitional and turbulent flow structures in these leading edge regions where the boundary layers are thin and the pertinent viscous length scales are small. The sensitivity of the results to grid refinement are explored in this work.

3.1. JAXA free air configuration: baseline resolution

The lift coefficient versus the angle of attack for the free air configuration using the baseline resolution is shown in figure 5(a). Error bars in the experiment denote the quoted experimental repeatability at select angles of attack. The predictions of both solvers are reasonably consistent with one another
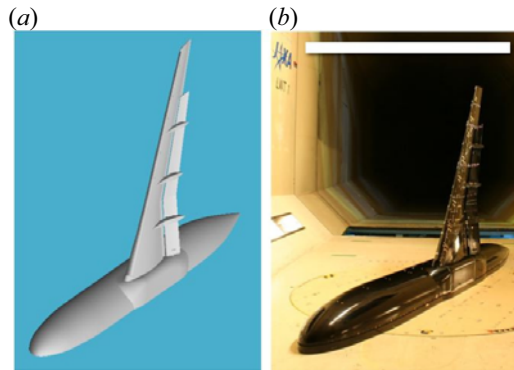
Figure 3. *The computational geometry (a) and the experimental wind tunnel model (b), with white scale bar in the upper right-hand corner of the experimental image corresponding to one wing semispan length of 2.3 m.*

ahead of stall (angles of attack less than 19 degrees). Both solvers show reasonable comparisons with the experimental measurements, although a shift of ($\Delta C_L \approx 0.05$–0.08) is seen for the lower angles in the charLES simulation results. The lift predictions from the charLES solutions at this resolution approximately predict both the maximum lift and the onset of the post-stall regime, while the Alya solution shows a deeper stall.

The drag polar is shown in figure 5(*b*). The shape of the predicted values prior to stall is consistent with the experimental measurements, although a shift of $\Delta C_D \approx 0.03$ versus the experiment is visible for the intermediate angles of attack becoming more pronounced near $C_{L,max}$. Results compiled from the AIAA HLPW3 showed a shift in the drag of similar magnitude from every simulation when comparing free air CFD with corrected wind tunnel data (Rumsey et al., 2019). A reason for this was not cited by the workshop committee, but could have to do with uncertainty associated with wind tunnel corrections for this configuration as our calculations that include the wind tunnel geometry do not exhibit such a shift when compared with uncorrected wind tunnel data. It is worth noting that the flow over a high lift aircraft can be considered similar to a bluff body flow. Friction drag contributes at most ≈10 % to the total drag at the lowest angle of attack, with the remaining ≈90 % being form drag. This breakdown is not a strong function of angle of attack (the breakdown between fiction and form drag is ≈5 %/95 % at the highest angle of attack).

Figure 5(*c*) shows the pitching moment coefficient for the JSM configuration, whose reference centre is at $x \approx 2.38$ m, $y = 0$, $z = 0$, or at approximately midchord of the inboard wing. The prediction of the pitching moment is of equal significance to the lift predictions as it describes the strength and direction of the response of the aircraft at a given angle of attack; a negative C_M would indicate a rotation of the body to push the nose down if the aircraft could move freely. The charLES results show that the moments are predicted to nearly within the experimental repeatability bounds up until the aircraft stall point. The key discrepancy noted here is the absence of a 'kink' observed in the experiment that produces a nearly constant C_M in the post-stall regime. The appearance of this kink in the experimental data is attributed to boundary layer separation on the inboard section near the wing-body juncture. Ito et al. (2019) suggest that this inboard separation is tied to the interaction of juncture flow with the tunnel floor boundary layer. As such, the ability of the simulations to capture this separation mechanism will be revisited in § 3.3 where the tunnel geometry is included in the calculations.

3.2. JAXA free air configuration: grid refinement near stall

The prediction of the maximum lift and the onset of boundary layer separation is one of largest deficiencies in existing simulation approaches for flight vehicles. The baseline resolutions suggested
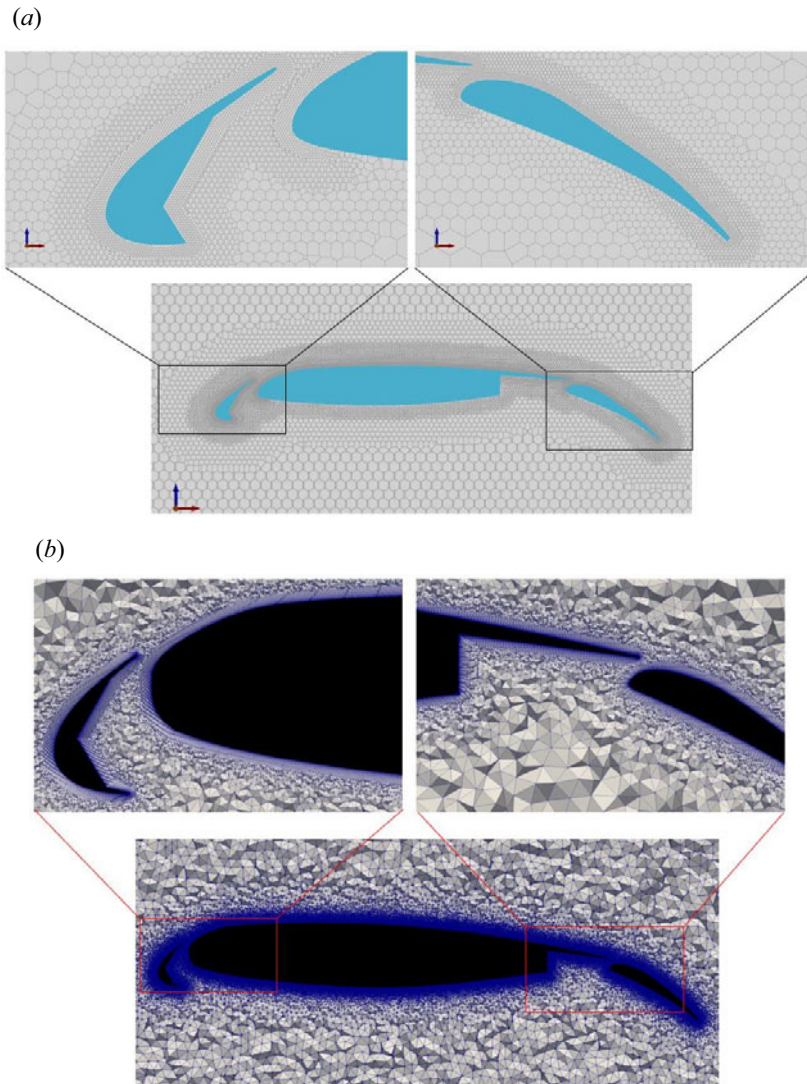
(*a*)

(*b*)

Figure 4. *Spanwise slices of the baseline grids used for JSM free air calculations including zoom views of the slat/main element leading edge and the main element trailing edge/flap. (a) Voronoi diagram based grid based on hexagonally close packed point arrangement for the charLES solver and (b) prismatic boundary layer elements with a tetrahedral far field used by the Alya flow solver.*

that the $C_{L,max}$ and post-stall regimes could be approximately captured with relatively coarse resolutions. To assess the robustness of these calculations, grid refinement studies are conducted.

Two different grid refinement studies have been carried out with Alya. In the first study, two different meshes were considered for the full angle of attack sweep. A mesh of 30M cells (with approximately five elements inside the boundary layer) and 180M cells (with approximately 10 elements inside the boundary layer), were chosen, the results of which are depicted in figure 6(*a*). A clear improvement in the lift prediction is observed with the increase of grid resolution; the 180M element mesh results lie inside the experimental repeatability bounds for all of the angles of attack considered.

In the second, targeted mesh refinement exercise a point after stall ($\alpha = 21.57°$) has been considered. A mesh refinement has been carried out up to 2B elements. The predicted C_L at this very high resolution
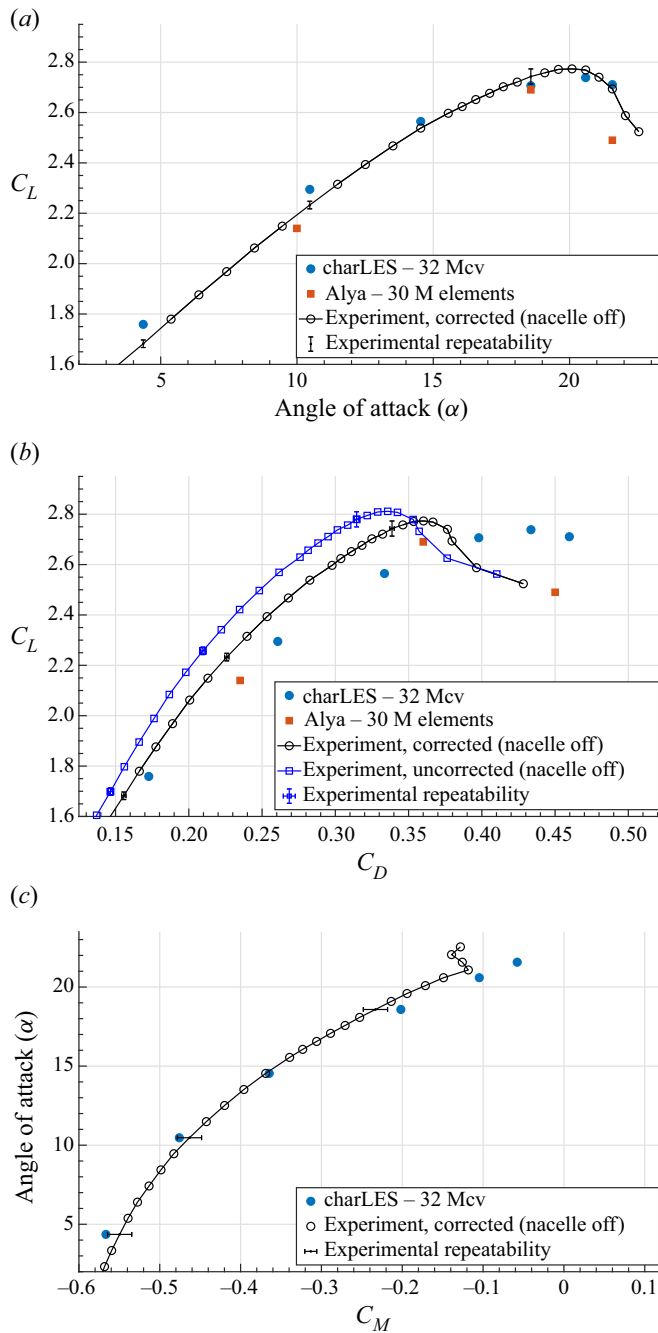
Figure 5. *(a) Lift curve (C_L versus α), (b) drag polar and (c) pitching moment coefficient (α versus C_M) from baseline resolution simulations compared with experimental measurements (charLES – 32 Mcv resolution is from Goc, Bose, & Moin, 2020). The uncorrected drag is shown in (b) as a measure of how large the tunnel to free air corrections are in the experiment.*
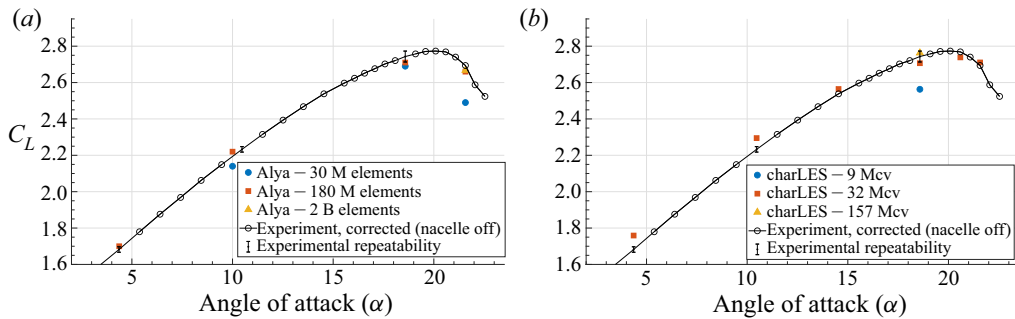
Figure 6. *The lift curve predicted by (a) Alya finite element code and (b) charLES finite volume code. A grid refinement sequence consisting of a coarse, medium and fine mesh is completed after the stall condition for Alya ($\alpha = 21.57°$), while the refinement sequence for charLES is near $C_{L,max}$ ($\alpha = 18.58°$) (charLES – 32 Mcv resolution is from* Goc et al. (2020)*).*
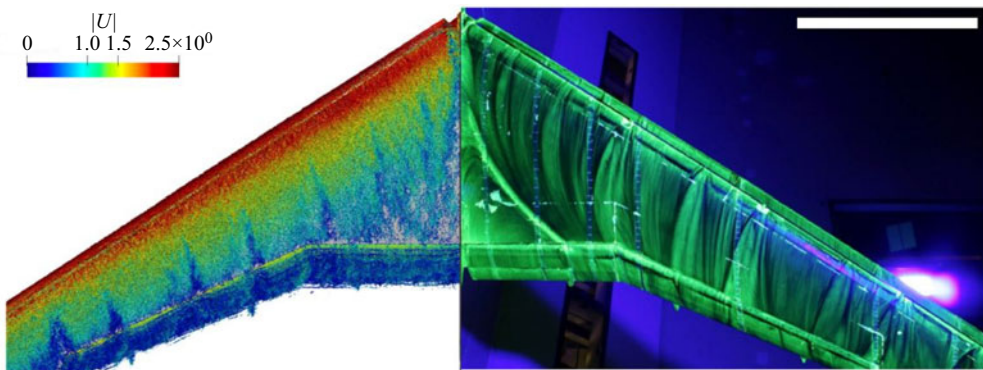
Figure 7. *Isocontours of the Q-criterion coloured by the non-dimensional velocity magnitude (red, 2.5; blue, 0) (as computed by Alya) versus oil visualizations from* Yokokawa et al. (2006)*, with white scale bar in the upper right-hand corner of the experimental image corresponding to one mean aerodynamic chord length of ≈ 530 mm. An isocontour value of one corresponds to the free stream velocity of ≈ 60 m s^{-1}.*

level is in agreement with the C_L value from the intermediate grid resolution considered. Figure 7, isocontours of the Q-criterion (Hunt, 1988) are depicted together with the experimental oil visualizations from Yokokawa et al. (2006). While the lift coefficient compares well with the corrected experimental data, the inboard separation observed is substantially weaker than is suggested by the experimental oil flow visualization. Given the level of grid convergence achieved, this discrepancy in inboard separation is consistent with the hypothesis of missing wind tunnel effects rather than poor numerical resolution. This will be further explored in § 3.3.

Additional grid refinement studies are also conducted using both solvers near angles of attack corresponding to $C_{L,max}$. Sectional pressure predictions at midspan and near the wing tip from both solvers at a finer resolution compare favourably with the experimental measurements at $\alpha = 18.6°$ (see figure 8). This suggests that the lift predictions are not due to fortuitous error cancellation along the wing. Sectional pressures from a series of grids at $\alpha = 18.6°$ at various spanwise locations are shown in figure 9. The predictions collapse relatively well at all resolutions, with the exception of the coarsest grid at the outboard wing station ($\eta = 0.77$). At the outboard section, the suction produced on the slat and main elements are notably reduced compared with the experimental measurements. The suction peak at the outboard section is strong ($-C_p \approx 10$) on the slat surface and is significantly higher than the inboard suction peaks. This suggests that the effect of under-resolution may be connected with the inability to
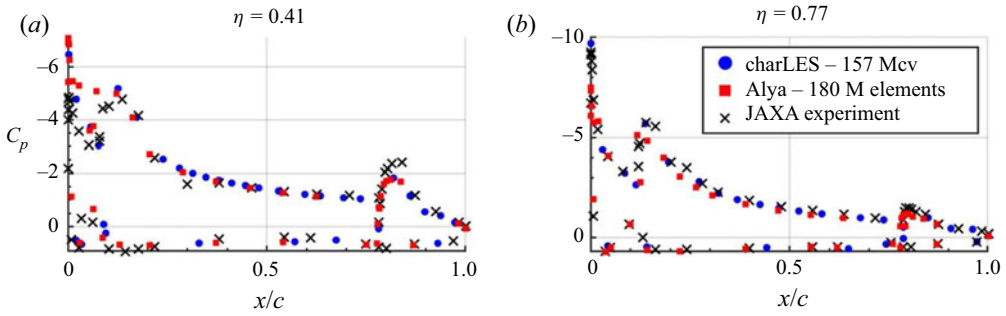
Figure 8. *Comparison of the sectional pressure predictions near the maximum lift condition, α = 18.58° at two stations along the wing, one at semispan fraction η = 0.41 and another near the wing tip at η = 0.77 compared with (uncorrected) experimental measurements.*
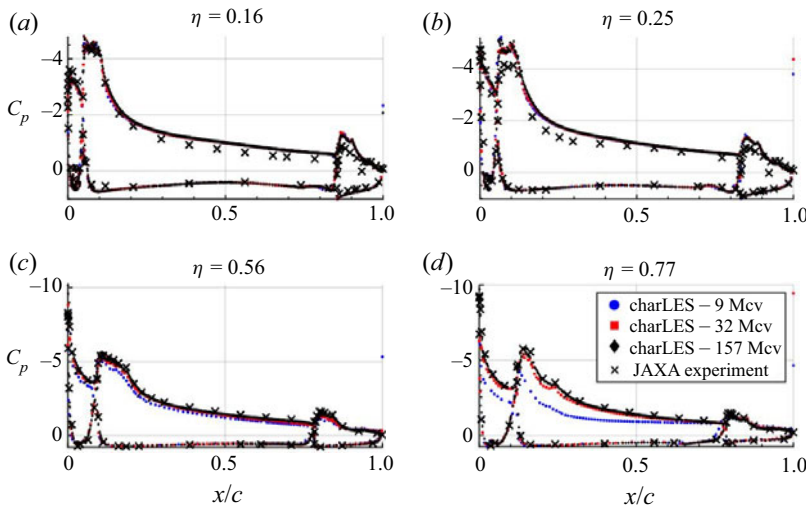

Figure 9. *Sectional pressure measurements from a grid refinement sequence using the charLES solver at α = 18.58° (charLES – 32 Mcv resolution is from Goc et al. (2020)).*

capture strong leading edge acceleration, which is predominantly an inviscid flow phenomenon. At the medium and fine grids, sectional pressures are well collapsed at the outboard sections and compare favourably with the experimental measurements. Figure 10 shows the effect of the grid resolution on the drag polar. Even at the finest resolution levels, errors in the prediction of drag persist using both solvers. We do not offer an explanation as to this except to point to the results in § 3.3 in which the inclusion of greater geometric fidelity in the representation of the test facility and comparison with uncorrected data led to significant improvements in the prediction of drag at comparable grid refinement levels.

3.3. Inclusion of wind tunnel effects in JSM calculations

The previous section showed encouraging comparisons with the experimental data, but there are two notable deficiencies: there is a systematic shift in the drag values obtained and a 'kickback' of the pitching moment associated with the appearance of an inboard stall mechanism is absent. It is possible that both of these issues are related to wind tunnel effects that are not captured in the free air setting. To test this hypothesis, simulations are conducted replicating the experimental facilities. All simulations in this section are performed using the charLES solver. The configuration now includes a flow-through nacelle
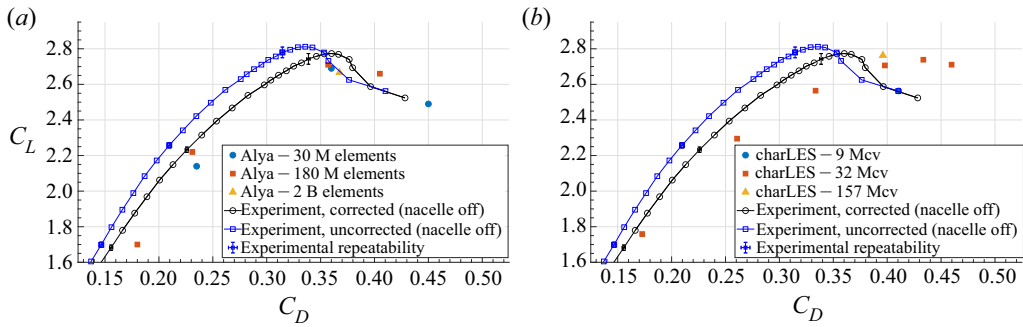
Figure 10. *The drag polar predicted by (a) Alya finite element code and (b) charLES finite volume code. A grid refinement sequence consisting of a coarse, medium and fine mesh is completed after the stall condition for Alya ($\alpha = 21.57°$), while the refinement sequence for charLES is near $C_{L,max}$ ($\alpha = 18.58°$) (charLES – 32 Mcv resolution is from* Goc et al. (2020)).

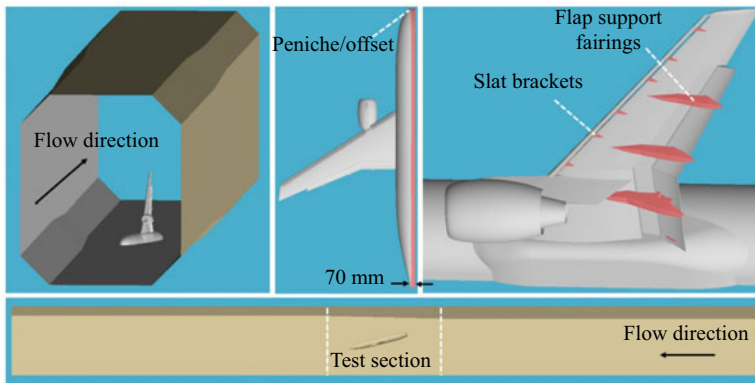**Figure 11.** *Schematic of the computational geometry for the JSM, nacelle on configuration, in the JAXA 6.5 m × 5.5 m Low-Speed Wind Tunnel.*
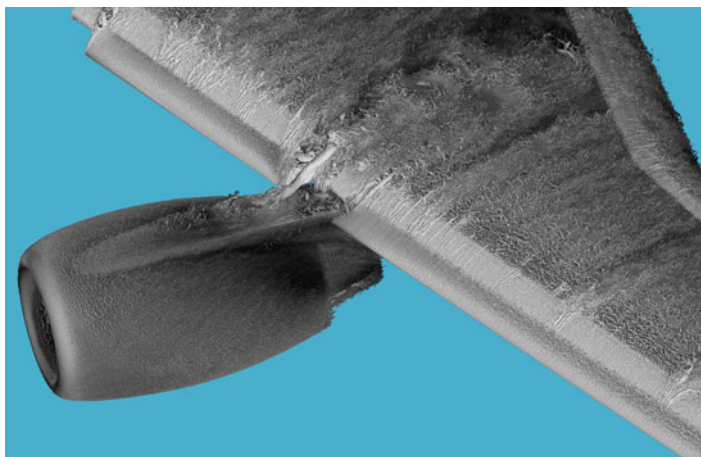

Figure 12. *Isosurfaces of Q-criterion coloured by velocity magnitude showing the structure of the turbulent flow field over an aircraft wing in landing configuration. Computed using the charLES solver at $\alpha = 18°$ including the engine nacelle and wind tunnel test section (not pictured) geometries.*
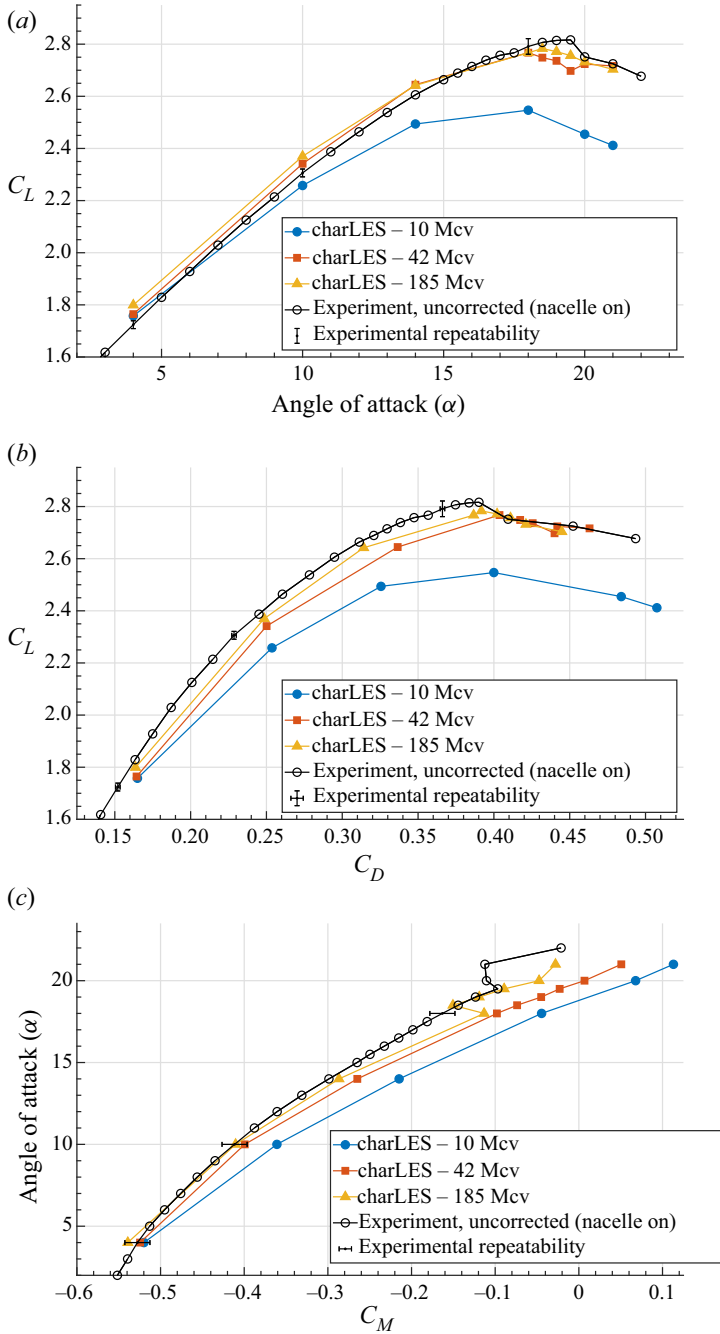
Figure 13. *(a) Lift curve (C_L versus α), (b) drag polar and (c) pitching moment coefficient (α versus C_M) plots for the JSM configuration including wind tunnel effects (42 Mcv resolution is from Goc et al., 2020).*

mounted on the underside of the wing and the test section walls (see figure 11). We note that our free air and wind tunnel calculations differ in that the wind tunnel geometry includes the engine nacelle while the free air calculations do not. This was done for expediency and in light of the experience in the context of RANS simulations by JAXA which revealed the important influence of the wind tunnel installation
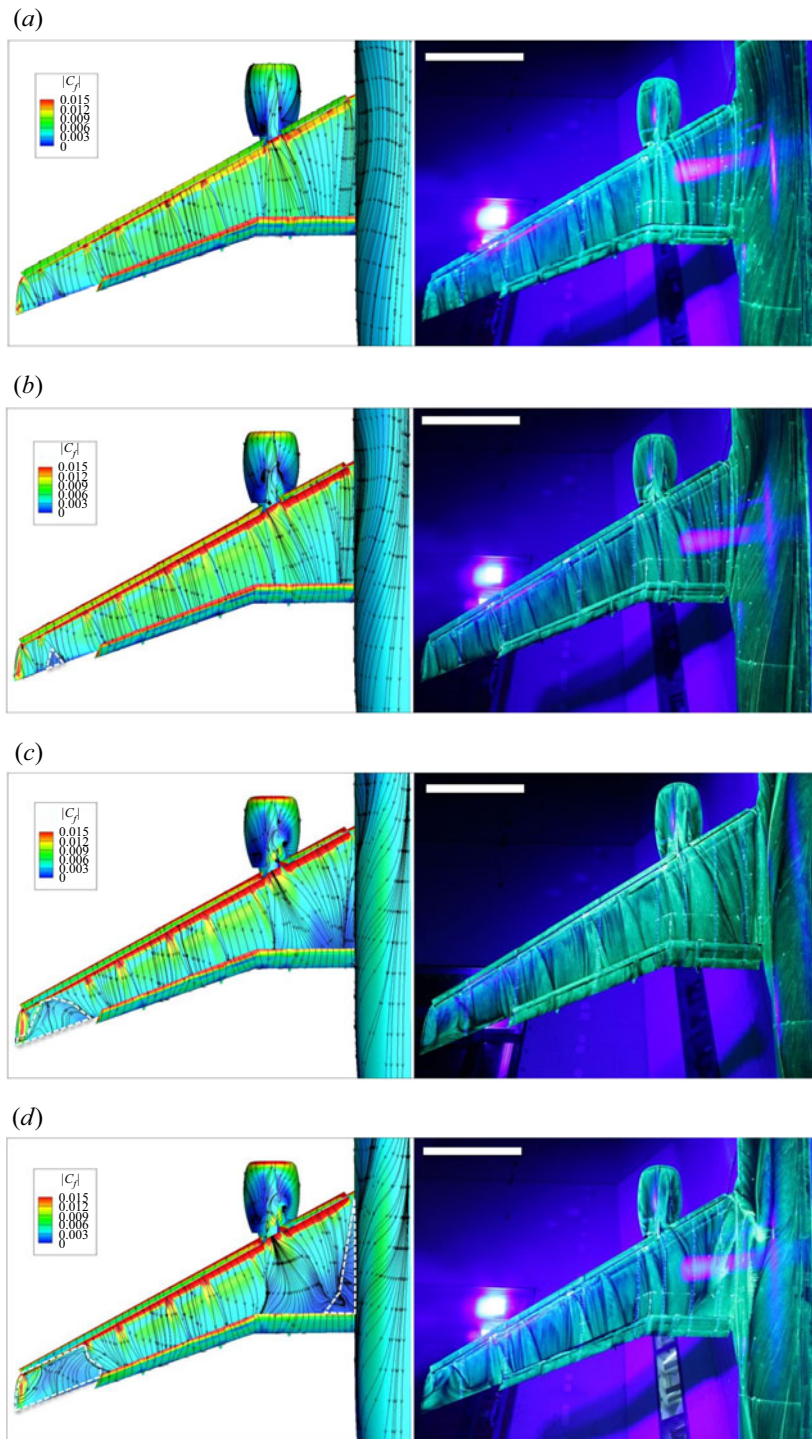
(*a*)

(*b*)

(*c*)

(*d*)

Figure 14. *Average skin friction contours with skin friction streamlines (left-hand column) from charLES fine grid simulations compared with experimental oil flow visualizations (right column), with white scale bar in the upper left-hand corner of the experimental images corresponding to one mean aerodynamic chord length of ≈530 mm (Yokokawa et al., 2008). The dashed white lines identify regions of flow separation in the simulations. Here (a) α = 4°; (b) α = 10°; (c) α = 18°; (d) α = 21°.*

on the inboard separation regardless of whether or not the engine nacelle was installed (Ito et al., 2019). The wind tunnel is extruded ≈5 fuselage lengths in the upstream and downstream directions from the test section to mitigate contamination from the inflow and outflow boundary conditions. At the inflow, a uniform plug flow is prescribed, while at the outflow non-reflecting Navier–Stokes characteristics boundary conditions are applied. The wind tunnel walls are treated as viscous (and wall modelled) because the development of a boundary layer on the tunnel floor is believed to be a key influence on the separation pattern at the wing juncture. The tunnel sidewall is coarsely resolved to capture first-order effects associated with the blockage imposed by the sidewall boundary layer. Still, evaluation of the sidewall boundary layer δ_{99} just ahead of the aircraft installation point reveals agreement to within 10 mm of the experimentally quoted value of 140 mm is achieved with approximately 10 LES grid points spanning its thickness in all three directions (Ito et al., 2019). The tunnel sidewall boundary layer is large compared with the model offset height of 70 mm and for this reason is believed to have a meaningful impact on the quantities of interest. Figure 12 shows a rich turbulent field predicted by these simulations. A wide range of turbulent length scales is readily discernible from small-scale boundary layer turbulence to large-scale phenomena such as a vortex forming over the engine nacelle. A grid refinement study is carried out for all angles of attack considered. The grid sequence is produced by refining the grid isotropically in all three dimensions in the vicinity of the wing and fuselage.

Figure 13 shows the prediction of the lift, drag and pitching moment in comparison with the raw (uncorrected) experimental data. As in the free air case, the medium and fine grids show convergence of the lift coefficient to within the experimental repeatability bounds at maximum lift conditions. Systematic improvement for lift, drag and pitching moment are achieved with additional grid refinement near the point of maximum lift. In the linear part of the lift curve, we observe overprediction of the lift coefficient. We note that the flaps are most highly loaded in a high lift configuration at the lowest angles of attack (Chin, Peters, Spaid, & McGhee, 1993). Examination of the solution on the fine grid at the low angles of attack reveals that the separation bubble observed on the flaps in the experiment and predicted by the coarse and medium grids has been suppressed on the fine grid. Similar separation bubble collapse behaviour with grid refinement using LES has been observed by Iyer and Malik (2021) in a canonical smooth body separation problem. Its cause and remedy are the focus of ongoing efforts. The finest grid now shows convergence of the predicted drag to the experimental values. Lastly, the 'kickback' of the pitching moment near stall is now observed, although with an angle of attack shift of approximately 1°. Figure 14(*d*) shows surface streamlines of skin friction from the LES compared with experimental oil flow visualizations at 21°. Inboard separation is now observed in the simulations consistent with the experimental observations and corroborates the mechanism associated with the kickback in the pitching moment. This inboard separation mechanism was notoriously difficult to capture in prior simulations conducted as part of AIAA HLPW3 (Rumsey et al., 2019). The inclusion of the wind tunnel effects appear to have remedied the highlighted shortcomings of the free air configuration predictions.

Figure 14 shows selected surface flow visualizations compared with the oil flow visualizations from the JAXA experimental campaign. Even the low angles of attack that correspond to the linear regime of the lift curve show flow separation near the trailing edge of the flap. This is reproduced in the simulation results, although the magnitude of the flap separation is slightly reduced (figure 14*a*). The impact of realistic aircraft geometry is visible across the angles of attack; particularly the appearance of wakes on the airfoil main element behind the slat brackets. Surface visualizations from the LES near the maximum lift operating point ($\alpha \approx 18°$) show the separation of the boundary layer near the wing tip consistent with the experimental observations (figure 14*c*). These visualizations provide some qualitative reinforcement that the flow mechanisms associated with the $C_{L,max}$ conditions are appropriately captured by the LES calculations.

4. Remarks on computational cost

The previous sections have established that LES simulations are capable of predicting critical quantities of interest for a full flight vehicle. These calculations can then be utilized as part of the design cycle if

Table 1. *Computational cost summary for a grid sequence using the charLES solver for the JSM configuration at* $\alpha = 18.58°$.

| Cell count (M)	Non-dim. time step	CPU core hours/ flow pass (k)	Wall-clock time to 30 flow passes (using 2000 Ivy bridge Intel CPU cores)	Wall-clock time to 30 flow passes (using 96 NVIDIA V100 GPUs)
32	6.47×10^{-5}	1.4	21 hours	48 minutes
157	3.23×10^{-5}	12.3	7.5 days	7 hours

their computational cost results in turnaround times that are tenable. Table 1 shows the computational cost for the charLES flow solutions on the medium and fine grids (where the calculations are of reasonable accuracy) for the free air configuration, while table 2 shows the computational cost for the Alya solutions. The computational cost for the simulations that include the wind tunnel geometry are similar as the resolution of the nacelle and tunnel marginally increase the grid count and do not impact the time step. Non-dimensional time step refers to the time step normalized by the flow pass time associated with the wing mean aerodynamic chord, while cell counts are listed in millions and CPU core hours are listed in thousands. The charLES grids are composed of isotropic hexagonal close packed cells and are self-similar in that the finer grid is produced by uniform refinement in all three coordinate directions in the vicinity of the solid boundaries, which results in a halving of the time step due to CFL restrictions of the time advancement scheme. The Alya grids have aspect ratios that vary from 1 : 16 (wall normal : streamwise/spanwise direction) on the coarsest grid to 1 : 4 on the finest grid in the boundary layer and are also refined uniformly in all three coordinate directions. The cost estimates are produced for a time horizon of $30c_{mac}/U_\infty$ (convective flow through times over the mean aerodynamic chord) which is deemed sufficient to obtain converged statistics. Time horizons of up to $40c_{mac}/U_\infty$ were used for angles of attack in a post-stall regime due to low frequency oscillations of the aerodynamic loading caused by the massive separation in that flow regime. The wall-clock time is computed based on either the use of 2000 CPU cores (Intel Ivy Bridge generation; NASA's Pleaides cluster) and 96 GPUs (NVIDIA Tesla V100s; Oak Ridge National Laboratory's Summit cluster). This resource allocation is selected as it can be considered reasonably accessible to simulation practitioners, in the experience of the authors. At this level of resource availability, turnaround times of less than a day can be obtained for the medium resolution (facilitated by either CPU or GPU architectures) or on the fine grid (when GPU accelerated). Turnaround times of less than a day were identified by the NASA CFD Vision 2030 report (Slotnick et al., 2014) as a critical threshold that would enable use of LES within the design cycle to complement lower fidelity simulations and experimental campaigns. Reynolds numbers must be scaled up by an order of magnitude to flight test conditions relative to wind tunnel tests such as those used for validation in this study. If we assume a WMLES-type cost inflation of $N \propto Re_c$ as in Choi and Moin (1994), the calculations performed within this study could be weakly scaled (as the grid and computer resources are concomitantly increased) to fit within the CFD Vision 2030 targeted turnaround times on existing GPU supercomputing resources.

Further reduction in the wall-clock time is possible by leveraging additional computational resources. The charLES flow solver can strongly scale (with parallel efficiencies >80 %) down to ≈2000 cells per CPU core or ≈10^6 cells per GPU. For the 157 million cell grid, that would correspond to the use of nearly 80×10^3 cores and 160 GPUs at the scalability limit. That would coincide with a wall-clock time of <4 h.

The tractability of the overall cost envelope can be assessed by comparison of the LES costs with the cost of steady RANS simulations. As a basis for this cost comparison, we choose the OpenFOAM simulations reported from the AIAA HLPW3 (Ashton, Denison, & Zastawny, 2018). The RANS calculations are conducted on a 109M cell grid and warm started (initial guesses of the solution are

Table 2. *Computational cost summary for a grid sequence using the Alya solver for the JSM configuration at α = 18.58°.*

Cell count (M)	Non-dim. time step	CPU core hours/ flow pass (k)	Wall-clock time to 30 flow passes (hr)	Number of CPUs (MareNostrum IV supercomputer)
30	8.2×10^{-4}	0.22	9	720
180	2.5×10^{-4}	4.26	63	2016
2000	1.2×10^{-4}	111.11	35	96 000
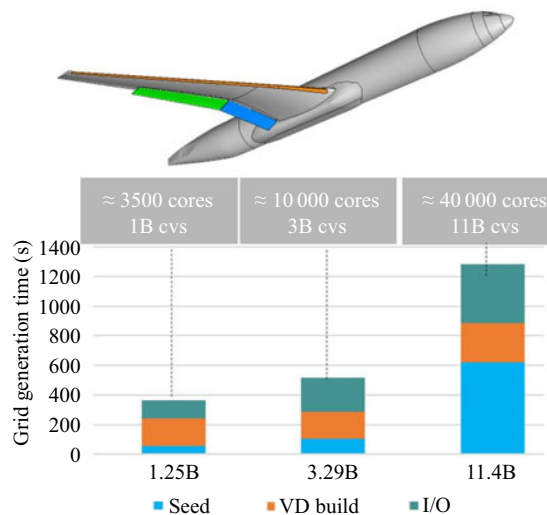
Figure 15. *Grid generation time in seconds for the charLES solver from the second mesh generation workshop for a representative high-lift geometry (the High-Lift Common Research Model). The computer time is broken down between the seeding of the volume domain with cells, building of the Voronoi diagram, and input/output operations (such as writing of the grid to disk). Massive grids ($O(10^{10})$ cells) are routinely built within 30 minutes (Woeber et al., 2019).*

provided from a nearby angle of attack) using a one equation Spalart–Allmaras turbulence model. These simulations take ≈60 × 10³ CPU core hours versus ≈40 × 10³ and ≈360 × 10³ core hours for the medium and fine grid cases, respectively. The medium grid resolution costs are comparable to the RANS solutions and fine grid calculations show an increase of a factor of six.

The overall simulation cost should also consider the time required for the generation of the computational grid. This cost is broken down between the seeding/building of the Voronoi diagram and associated input/output time in figure 15 for very large grids used to highlight the efficiency of the Voronoi grid generation paradigm. The grid generation cost is negligible compared with flow solution time when utilizing the parallel Voronoi diagram mesh generator (for use with the charLES flow solver). The generation of the 157 million cell body-fitted, Voronoi grids is created in ≈2 minutes on 10^3 Ivy Bridge CPU cores. This mesh generator has been effectively utilized to generate grids of $O(10^{10})$ cells on 40 × 10³ cores in less than 30 minutes (Woeber, Masters, & McDaniel, 2019).

Note that this favourable timing pertains only to the actual grid generation. The approximate resolutions as a function of space are to be determined beforehand. This process is guided with little *a priori* information; specifically, it relies primarily on providing a chosen $O(1)$ number of cells per boundary-layer thickness at the trailing edge of the main element, assuming the flow was attached and subjected

to no pressure gradients. This estimate is complemented locally by the experimental data near the leading edge to properly capture the strong acceleration therein (e.g. ensure $\Delta\, \mathrm{d}C_p/\mathrm{d}x$ is sufficiently small for a given numerical scheme). A more rigorous grid estimate from formal boundary-layer thickness estimates can be found in Lozano-Durán et al. (2021).

The JSM case was studied experimentally and numerically at $Re_c = 2 \times 10^6$, which is potentially low enough that the results may still be impacted by low Reynolds number effects compared with flight conditions. The computational cost can be extrapolated based on the scaling arguments advanced by Choi and Moin (2012) where the grid point count is $\propto Re_c$. The overall cost would then scale by $Re_c^{4/3}$, which accounts for the reduction in time step. With this scaling, turnaround times of the order of a few days would be possible on ≈ 100 GPUs for $Re_c \sim 10^7$. Turnaround times of several hours would still be attainable at $Re_c \sim 10^7$ if calculations were conducted at the strong scalability limit of the solver (through the use of additional computational resources). The turnaround times achieved on 96 GPU nodes of Summit shown in table 1 substantiate these estimates. The Reynolds number regime of the current study is in the range of those attained in experimental facilities where considerable data is presently available.

5. Conclusions

The WMLES calculations are conducted of the JSM, which is a representative geometry of a contemporary, commercial high-lift aircraft. These simulations achieve two important milestones. First, the LES calculations demonstrate sufficient accuracy for the prediction of critical quantities of interest of the aerodynamic loading (lift, drag and pitching moment) across the angle of attack range. Of particular note, the maximum lift is computed to nearly within the bounds of experimental repeatability at an angle of attack within a degree of the experimental measurements. Surface pressure and near-wall flow visualizations are compared with experimental measurements, which corroborate that the LES calculations properly replicate appropriate flow structures. These include the prediction of outboard stall near $C_{L,max}$ and the onset of an inboard separation pattern when the wind tunnel effects are properly accounted for. This level of accuracy is on the threshold of meeting industrial requirements. Second, by leveraging modern, massively parallel computer architectures, turnaround times of less than a day are achieved for these calculations with modest resource requirements. This turnaround time is sufficiently rapid to be used as part of the engineering design cycle. It also belies the existing conception that the use of (wall modelled) LES would be prohibitively expensive and intractable. While there remains considerable work to demonstrate the generality of these observations, this simulation campaign represents a significant achievement in the use of high-fidelity simulation approaches for practical aeronautical applications.

Author Contributions. K.A.G. performed charLES JSM calculations. O.L. performed Alya calculations. G.I.P., O.L. and S.T.B collaborated on the initial set of JSM calculations at the 2018 CTR Summer Program. P.M. is the PI for the Boeing and NASA grants and advised on the research.

References

Angelino, M., Fernández-Yáñez, P., Xia, H., & Page, G. (2020). Large-eddy simulation with modeled wall stress for complex aerodynamics and stall prediction. *AIAA Journal*, *59*(4), 1–43.

Ashton, N., Denison, M., & Zastawny, M. (2018). 3rd high-lift workshop summary paper-OpenFOAM, STAR-CCM+ & LAVA simulations on unstructured grids. In *2018 AIAA Aerospace Sciences Meeting, 8–12 January, Kissimmee, Florida* (p. 1253).

Balin, R., & Jansen, K. E. (2018). A comparison of RANS, URANS, and DDES for high lift systems from HiLiftPW-3. In *2018 AIAA Aerospace Sciences Meeting, 8–12 January, Kissimmee, Florida* (p. 1254).

Bodart, J., Larsson, J., & Moin, P. (2013). Large eddy simulation of high-lift devices. In *21st AIAA Computational Fluid Dynamics Conference, 24–27 June, San Diego, California* (p. 2724).

Bose, S. T., & Park, G. I. (2018). Wall-modeled large-eddy simulation for complex turbulent flows. *Annual Review of Fluid Mechanics*, *50*, 535–561.

Bres, G. A., Bose, S. T., Emory, M., Ham, F. E., Schmidt, O. T., Rigas, G., & Colonius, T. (2018). Large-eddy simulations of co-annular turbulent jet using a Voronoi-based mesh generation framework. In *2018 AIAA/CEAS Aeroacoustics Conference, 25–29 June, Atlanta, Georgia* (p. 3302).

Cabot, W., & Moin, P. (2000). Approximate wall boundary conditions in the large-eddy simulation of high Reynolds number flow. *Flow, Turbulence and Combustion*, *63*(1), 269–291.

Cary, A. W., Yousuf, M., Li, P., & Mani, M. (2018). Current practice unstructured grid CFD results for 3rd AIAA High Lift Prediction Workshop. In *2018 AIAA Aerospace Sciences Meeting*, *8–12 January, Kissimmee, Florida* (p. 1037).

Chandrashekar, P. (2013). Kinetic energy preserving and entropy stable finite volume schemes for compressible Euler and Navier–Stokes equations. *Communications in Computational Physics*, *14*(5), 1252–1286.

Charnyi, S., Heister, T., Olshanskii, M. A., & Rebholz, L. G. (2017). On conservation laws of Navier–Stokes Galerkin discretizations. *Journal of Computational Physics*, *337*, 289–308.

Chin, V., Peters, D., Spaid, F., & McGhee, R. (1993). *AIAA Paper 1993-3137*.

Choi, H., & Moin, P. (1994). Effects of the computational time step on numerical solutions of turbulent flow. *Journal of Computational Physics*, *113*(1), 1–4.

Choi, H., & Moin, P. (2012). Grid-point requirements for large eddy simulation: Chapman's estimates revisited. *Physics of Fluids*, *24*(1), 011702.

Clark, A. M., Slotnick, J. P., Taylor, N. J., & Rumsey, C. L. (2020). Requirements and challenges for CFD validation within the high-lift common research model ecosystem. In *AIAA AVIATION 2020 FORUM, 15–19 June, virtual event* (p. 2772).

Coder, J. G., Pulliam, T. H., & Jensen, J. C. (2019). High-lift simulations of the Japan Aerospace Exploration Agency standard model. *Journal of Aircraft*, *56*(5), 1822–1832.

Codina, R. (2001). Pressure stability in fractional step finite element methods for incompressible flows. *Journal of Computational Physics*, *170*, 112–140.

Fu, L., Karp, M., Bose, S. T., Moin, P., & Urzay, J. (2021). Shock-induced heating and transition to turbulence in a hypersonic boundary layer. *Journal of Fluid Mechanics*, *909*, A8.

Germano, M., Piomelli, U., Moin, P., & Cabot, W. H. (1991). A dynamic subgrid-scale eddy viscosity model. *Physics of Fluids A - Fluid Dynamics*, *3*, 1760–1765.

Ghate, A. S., Kenway, G. K., Stich, G.-D., Browne, O. M., Housman, J. A., & Kiris, C. C. (2021). Transonic lift and drag predictions using wall-modelled large eddy simulations. In *AIAA Scitech 2021 Forum, 11–15 & 19–21 January*, *virtual event* (p. 1439).

Goc, K., Bose, S., & Moin, P. (2020). Wall-modeled large eddy simulation of an aircraft in landing configuration. In *AIAA Aviation 2020 Forum, 15–19 June*, *virtual event* (p. 3002).

Gottlieb, S., Shu, C.-W., & Tadmor, E. (2001). Strong stability-preserving high-order time discretization methods. *SIAM Review*, *43*(1), 89–112.

Honein, A. E., & Moin, P. (2004). Higher entropy conservation and numerical stability of compressible turbulence simulations. *Journal of Computational Physics*, *201*(2), 531–545.

Hunt, J. (1988). Eddies, streams and convergence zones in turbulent flows. In *Proceedings of the Summer Program*. Center for Turbulence Research, Stanford University.

Ito, Y., Murayama, M., Yokokawa, Y., Yamamoto, K., Tanaka, K., & Hirai, T. (2019). Wind tunnel interference effects on Japan Aerospace Exploration Agency's standard model. In *AIAA SciTech 2019 Forum, 7–11 January, San Diego, California* (p. 2178).

Iyer, P. S., & Malik, M. R. (2020). Wall-modeled LES of the NASA juncture flow experiment. In *AIAA SciTech 2020 Forum, 6–10 January, Orlando, Florida* (p. 1307).

Iyer, P. S., & Malik, M. R. (2021). Wall-modeled LES of flow over a Gaussian bump. In *AIAA Scitech 2021 Forum, 11–15 & 19–21 January*, *virtual event* (p. 1438).

Jofre, L., Lehmkuhl, O., Ventosa, J., Trias, F., & Oliva, A. (2014). Conservation properties of unstructured finite-volume mesh schemes for the Navier–Stokes equations. *Numerical Heat Transfer, Part B: Fundamentals*, *54*(1), 53–79.

Konig, B., Fares, E., Murayama, M., & Ito, Y. (2018). PowerFLOW simulations for the third AIAA high-lift prediction workshop. In *2018 AIAA Aerospace Sciences Meeting*, *8–12 January, Kissimmee, Florida* (p. 1255).

Larsson, J., Kawai, S., Bodart, J., & Bermejo-Moreno, I. (2016). Large eddy simulation with modeled wall-stress: Recent progress and future directions. *Mechanical Engineering Reviews*, *3*(1), 15–00418.

Lehmkuhl, O., Houzeaux, G., Owen, H., Chrysokentis, G., & Rodriguez, I. (2019). A low-dissipation finite element scheme for scale resolving simulations of turbulent flows. *Journal of Computational Physics*, *390*, 51–65.

Lehmkuhl, O., Park, G., Bose, S., & Moin, P. (2018). Large-eddy simulation of practical aeronautical flows at stall conditions. In *Proceedings of the 2018 Summer Program* (pp. 87–96). Center for Turbulence Research, Stanford University.

Lozano-Duran, A., Bose, S. T., & Moin, P. (2020). Prediction of trailing edge separation on the NASA juncture flow using wall-modeled LES. In *AIAA Scitech 2020 Forum, 6–10 January, Orlando, Florida* (p. 1776).

Lozano-Durán, A., Bose, S. T., & Moin, P. (2021). Performance of wall-modeled les with boundary-layer-conforming grids for external aerodynamics. *AIAA Journal* (in press).

Mittal, R., & Moin, P. (1997). Suitability of upwind-biased finite difference schemes for large-eddy simulation of turbulent flows. *AIAA Journal*, *35*(8), 1415–1417.

Moin, P., Squires, K., Cabot, W., & Lee, S. (1991). A dynamic subgrid-scale model for compressible turbulence and scalar transport. *Physics of Fluids A: Fluid Dynamics*, *3*(11), 2746–2757.

Park, G. I. (2017). Wall-modeled large-eddy simulation of a high Reynolds number separating and reattaching flow. *AIAA Journal*, 3709–3721.

Park, G. I., & Moin, P. (2016). Space-time characteristics of wall-pressure and wall shear-stress fluctuations in wall-modeled large eddy simulation. *Physical Review Fluids*, *1*(2), 024404.

Piomelli, U., & Balaras, E. (2002). Wall-layer models for large-eddy simulations. *Annual Review of Fluid Mechanics*, *34*(1), 349–374.

Rumsey, C. L., Slotnick, J. P., & Sclafani, A. J. (2019). Overview and summary of the 3rd AIAA high lift prediction workshop. *Journal of Aircraft*, *56*(2), 621–644.

Slotnick, J., Khodadoust, A., Alonso, J., Darmofal, D., Gropp, W., Lurie, E., & Mavriplis, D. (2014). *CFD Vision 2030 Study: A Path to Revolutionary Computational Aerosciences* (Tech. Rep. NASA/CR-2014-218178) (pp. 1–51). The Boeing Company.

Spalart, P. R. (1997). Comments on the feasibility of LES for wings, and on a hybrid RANS/LES approach. In *Proceedings of the First AFOSR International Conference on DNS/LES, 4–8 August, Ruston, Louisiana*. Greyden Press.

Tadmor, E. (2003). Entropy stability theory for difference approximations of nonlinear conservation laws and related time-dependent problems. *Acta Numerica*, *12*(1), 451–512.

Trias, F. X., & Lehmkuhl, O. (2011). A self-adaptive strategy for the time integration of Navier–Stokes equations. *Numerical Heat Transfer Part B*, *60*(2), 116–134.

Vazquez, A. M., Houzeaux, G., Koric, S., Artigues, A., Aguado-Sierra, J., Aris, R., . . . Valero, M. (2016). Alya: Multiphysics engineering simulation towards exascale. *Journal of Computational Science*, *14*, 15–27.

Vreman, A. (2004). An eddy-viscosity subgrid-scale model for turbulent shear flow: Algebraic theory and applications. *Physics of Fluids*, *16*, 3670–3681.

Wang, M., & Moin, P. (2002). Dynamic wall modeling for large-eddy simulation of complex turbulent flows. *Physics of Fluids*, *14*, 2043–2051.

Woeber, C., Masters, J. S., & McDaniel, D. R. (2019). Summary of exascale and remeshing efforts for the second geometry and mesh generation workshop. In *AIAA Aviation 2019 Forum, 17–21 June, Dallas, Texas* (p. 3458).

Yang, X. I., & Griffin, K. P. (2021). Grid-point and time-step requirements for direct numerical simulation and large-eddy simulation. *Physics of Fluids*, *33*(1), 015108.

Yokokawa, Y., Murayama, M., Ito, T., & Yamamoto, K. (2006). Experiment and CFD of a high-lift configuration civil transport aircraft model. In *25th AIAA Aerodynamic Measurement Technology and Ground Testing Conference, 5–8 June, San Francisco, California* (p. 3452).

Yokokawa, Y., Murayama, M., Kanazaki, M., Murota, K., Ito, T., & Yamamoto, K. (2008). Investigation and improvement of high-lift aerodynamic performances in lowspeed wind tunnel testing. In *46th AIAA Aerospace Sciences Meeting and Exhibit, 7–10 January, Reno, Nevada* (p. 350).

Yokokawa, Y., Murayama, M., Uchida, H., Tanaka, K., Ito, T., Yamamoto, K., & Yamamoto, K. (2010). Aerodynamic influence of a half-span model installation for high-lift configuration experiment. In *48th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition, 4–7 January, Orlando, Florida* (p. 684).