TOWARDS CERTIFICATION BY ANALYSIS: LARGE-EDDY SIMULATIONS OF COMMERCIAL AIRCRAFT ACROSS THE FLIGHT ENVELOPE

By
Konrad Andrzej Goc and Parviz Moin

Prepared with support from
NASA (Grant Number NNX15AU93A),
Boeing Research & Technology, and
the Franklin P. and Caroline M. Johnson Fellowship

Report No. TF-191

Flow Physics and Computational Engineering Group
Department of Mechanical Engineering
Stanford University
Stanford, CA 94305

March 2023

Towards Certification by Analysis: Large-Eddy Simulations of Commercial Aircraft Across the Flight Envelope

By

Konrad Andrzej Goc and Parviz Moin

Prepared with support from
NASA (Grant Number NNX15AU93A),
Boeing Research & Technology, and
the Franklin P. and Caroline M. Johnson Fellowship

Report No. TF-191

Flow Physics and Computational Engineering Group
Department of Mechanical Engineering
Stanford University
Stanford, CA 94305

March 2023

Abstract

This dissertation presents a summary of numerical simulations performed using several realistic aircraft models, both in landing and cruise configurations and describes the current state and predictive capabilities of high-fidelity Large Eddy Simulation (LES) applied to realistic commercial aircraft configurations. The objective of the selected cases is to demonstrate a level of capability and cost effectiveness of LES for the prediction aeronautical flows of engineering significance that will make LES a useful tool for routine use in the aerospace industry. One of the selected model problems is the Japanese Exploration Agency Standard Model (hereafter JSM). This configuration was selected due to the recent interest garnered by its featuring in the AIAA Third High-Lift Prediction Workshop 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 [18], Lattice-Boltzman VLES [41], or delayed DES [12, 5]), most calculations presented in the workshop were steady RANS simulations deploying variants of the Spalart-Allmaras or SST models for the Reynolds stress closure [63]. A key takeaway from this workshop series has been that the accuracy of steady RANS techniques has plateaued, particularly near stall and in off-design conditions. Large variations from different solvers were also observed, even when the same models/gridding strategies were employed.

The calculations described in this dissertation leverage an LES approach in which prohibitive cost requirements associated with wall-bounded turbulence at realistic Reynolds numbers [81, 15] necessitate the introduction of wall models to ameliorate the stringent grid requirements associated with wall-resolved LES. Equilibrium wall models in which the unsteady, convective, and pressure gradient terms appearing in the turbulent boundary layer equations are assumed to be in balance [38, 11] have been shown to work well in flows of interest [25, 28, 51, 8] and are used in this work. The combination of these physics-based modeling choices along with recent advances in computing hardware (such as GPU-based high-performance computing clusters) and low-dissipation numerical methods for LES [39] have made LES a powerful tool for use in informing design decisions in industry. Validation efforts such as the one presented in this dissertation are a key component towards building confidence in the predictive capability of this emerging computational fluid dynamics (CFD) paradigm.

The studies presented herein which use the JSM configuration in general showed good prediction of the C_L across the lift curve with the coefficient of lift at maximum lift, $C_{L,max}$, being predicted to within 3 lift counts of the experimental value (i.e. within the tolerances required by the aerospace industry of $\Delta C_L \leq 0.03$ at maximum lift[17]). The grid point requirements to achieve this level of accuracy are reduced compared to recent estimates (even for wall modeled LES), with the solutions showing systematic improvement upon grid refinement on grids numbering up to $\approx 150Mcv$. Investigations which included the wind tunnel facility were made in order to address one of the key deficiencies of the free air calculations: the incorrect prediction of the stall mechanism which in the free air cases was missing a large inboard separation. It was found that inclusion of the tunnel facility in the simulations did indeed cause the flow to separate at the leading edge of the wing root and the presence of a "break" in the pitching moment curve associated with this phenomenon was captured, building confidence in the predictive capabilities of LES for high-lift flows, especially in the presence of wind tunnel effects. Another key finding from the JSM calculations was the fact that dynamic subgrid-scale turbulence models such as the DSM model in which the coefficient of the subgrid viscosity is modulated in space and time based on the finest resolved length scales of turbulence on the LES grid [56, 24] were shown to outperform static coefficient models such as the Vreman model [76] in their ability to predict both trailing edge flap separation and juncture flow separation.

In order to establish the robustness of these LES methods with respect to aircraft configuration and Reynolds number, further explorations were carried out in the high-lift flow regime with the High-Lift Common Research Model (hereafter CRM-HL), which has become the new benchmark flow in the community following the Fourth High-Lift Prediction Workshop [64]. The CRM-HL experiment is run at nearly $3\times$ the Reynolds number of JAXA and features a more challenging geometry, with slat suction peaks that reach up to -20 in C_p compared with suctions up to -10 in the case of the JSM. Both of these differences impose stricter grid resolution requirements and pose a greater predictive challenge to LES models (both SGS and wall models). In this effort, we sought to apply lessons related to modeling choices and gridding approach from the JSM case to accurately predict aircraft maximum lift. For the CRM-HL configuration, a more comprehensive grid refinement study was performed compared with the JSM case. In this case, full alpha sweeps were carried out on grids numbering up to 1.5 Bcv. The solutions exhibited a systematic decrease in their sensitivity to grid resolution with each successive refinement until the ≈ 1.5 Bcv and ≈ 400 Mcv were nearly grid-insensitive in their predictions of lift, drag, pitching moment, and surface pressure, giving indication for the first time of what the grid resolution requirements for convergence of engineering quantities of interest for practical aircraft flows are. The ≈ 400 Mcv calculations are carried out on 600 V100 GPU's within ≈ 5 hours of walltime to reach statistical convergence of the forces/moments. For the CRM-HL again wind tunnel effects were explored, this time with a more complete representation of the wind tunnel facility which included not just the test section

as in the case of JSM, but also the inlet, contraction, diffuser sections as well. This necessitated developing a workflow in LES for clearing the long transient timescales associated with the tunnel facility acoustics which was done by grid sequencing solutions from coarser to successively finer grids. The findings from including the calculations which included the tunnel facility were less clear-cut than in the case of the JSM (which is hypothesized to be due to the presence of discrepancies in the tunnel sidewall boundary layer development between the QinetiQ experiment and simulations), but did generally show a trend towards improving the character of the wing root separation pattern at stall. Explorations with numerical boundary conditions to delay nacelle lip separation are made and a novel non-Boussinesq subgrid model is explored to assess its impact on the persistent issue of lift over-prediction by LES at low angles of attack in high-lift flows. The model was shown to largely rectify similar issues in canonical flows [1], but its effect on the prediction of the trailing edge flap flow was to cause over-separation and a corresponding lift under-prediction compared with traditional SGS models such as the Dynamic Smagorinsky.

Finally, an extension to the transonic flow regime was pursued in which LES was applied to the transonic NASA Common Research Model (CRM), which has been the focus of several Drag Prediction Workshops (DPW). Sensitivities to laminar-to-turbulent transition, symmetry plane treatment, and grid topology are established and suggestions for best practices in these simulations are made. Specifically, the inclusion of experimental trip dots is found to be an important component of accurately prediction surface pressures and drag coefficients of transonic low angle of attack flows as the flow in these cases is more streamlined than in their high-lift counterparts and the integration of surface skin friction accounts for nearly $\approx 50\%$ of the drag coefficient compared to only $\approx 5\%$ in high-lift flows. Additionally, an important observation is made regarding symmetry plane boundary conditions in LES. Simpler problems, such as a Blasius boundary layer flow encountering a trip dot on a flat plate are studied in order to draw conclusions pertaining to transition behavior in a controlled environment before deploying the trip dots on the full aircraft configuration. It is also found that the simulations exhibit strong sensitivity to whether they are run with the entire full-span geometry of the aircraft is including or whether a numerical symmetry plane boundary condition is used, with the full span simulations showing much better agreement with the experiment (which uses a full span aircraft). This suggests that in LES, the numerical shortcut of using a symmetry plane on the aircraft center plane may be inappropriate for aircraft simulations, potentially owing to the instantaneous asymmetries that arise when the component of turbulent fluctuations normal to the symmetry plane cause flow to cross this plane. Additionally, anisotropic stranded Voronoi meshes are used and found to outperform their isotropic Voronoi HCP counterparts in this flow as they allow for additional targeted resolution in the wall-normal direction for better grid support of near-wall turbulent structures. Overall, it is found that promoting transition to turbulence via an array of cylindrical trip dots, using the full span of the aircraft, including the experimental sting mounting system in the simulations, and leveraging prismatic boundary layer grids all tend to improve the

quality of the LES solutions. Again, a novel non-Boussinesq subgrid-scale model and sensor-based wall models are applied to this complex external flow over an aircraft. This time, promising results are obtained, particularly on relatively coarse grids numbering about 100 million control volumes (100 Mcv). Finally, the angle of attack of the transonic flows is pushed to high values in order to assess the ability of LES to predict the onset of aerodynamic buffet. This study is performed on a NACA 0012 flow and shows promising early results, though there are discrepancies of $10 - 15\%$ in the precise prediction of the angle of attack of buffet onset and of the frequency of the shock oscillations once buffet is sustained. Still, these efforts offer preliminary indication of the accuracy of LES in this flow regime and propose a tractable workflow for identification of buffet onset using rotating flow solver technology to dynamically modulate the angle of attack of the airfoil section with time. Overall, these efforts serve as important steps towards building confidence in the predictive capabilities of LES applied to complex external aerodynamic flows at the edge of the flight envelope, both in high-lift and transonic regimes, which is a key ingredient in the industry-wide push towards Certification by Analysis.

Acknowledgments

The work carried out in the thesis was made possible by computing resources that were awarded through the both the Argonne and Oak Ridge Leadership Computing Facilities under the INCITE and ALCC Allocation award programs while research funding at Stanford was provided by NASA (Grant Number NNX15AU93A), Boeing Research & Technology, and by The Franklin P. and Caroline M. Johnson Fellowship.

THIS PAGE DELIBERATELY LEFT BLANK

Nomenclature

C_L	=	lift coefficient
$C_{L,max}$	=	maximum lift coefficient
C_D	=	drag coefficient
C_M	=	moment coefficient
C_p	=	pressure coefficient
C_f	=	skin friction coefficient
τ_w	=	wall stress
α	=	angle of attack
ρ	=	density
ν	=	kinematic viscosity
ν_t	=	eddy viscosity
μ	=	dynamics viscosity
T	=	temperature
P	=	pressure
E	=	energy
k	=	thermal conductivity
S_{ij}	=	strain rate tensor
τ_{ij}	=	stress tensor
κ	=	Karman constant
B	=	log law intercept
Q_j	=	heat flux
$(\dot{\cdot})$	=	Reynolds averaging
$(\tilde{\cdot})$	=	Favre averaging
q	=	dynamic pressure
P_t	=	stagnation pressure
u_i	=	velocity components
x_i	=	coordinate directions
y^+	=	viscous length scale

U_∞	=	freestream velocity
u_τ	=	friction velocity
$U_{\|\|}$	=	velocity component parallel to the local surface tangent
U_I	=	inviscid velocity component parallel to the local surface tangent
Re	=	Reynolds number
M	=	Mach number
η	=	semispan fraction
δ	=	boundary layer thickness
δ^*	=	displacement thickness
δ_{99}	=	boundary layer thickness based on 99% of inviscid velocity profile
H	=	shape factor
θ	=	momentum thickness
$h_{0,max}$	=	maximum stagnation enthalpy
Δ	=	grid cell size

Acronyms

LES	=	Large-Eddy Simulation
WMLES	=	Wall-Modeled Large-Eddy Simulation
WRLES	=	Wall-Resolved Large-Eddy Simulation
JSM	=	Japanese Exploration Agency Standard Model
AIAA	=	American Institute of Aeronautics & Astronautics
HLPW	=	High-Lift Prediction Workshop
DPW	=	Drag Prediction Workshop
RANS	=	Reynolds-Averaged Navier Stokes
VLES	=	Very Large-Eddy Simulations
DES	=	Detached Eddy Simulation
SST	=	Shear Stress Transport
CRM	=	Common Research Model
CRM-HL	=	High-Lift Common Research Model
SGS	=	Subgrid-Scale Model
GPU	=	Graphics Processing Unit
CPU	=	Central Processing Unit
Mcv	=	Million Control Volumes
Bcv	=	Billion Control Volumes
CbA	=	Certification by Analysis
ALCC	=	ASCR Leadership Computing Challenge
NTF	=	National Transonic Facility
HCP	=	Hexagonal Close Packed
DNS	=	Direct Numerical Simulation
MD30P30N	=	McDonnell Douglas 30P30N
CFD	=	Computational Fluid Dynamics
NACA	=	National Advisory Committee for Aeronautics
QOI	=	Quantity of Interest
DSM	=	Dynamic Smagorinsky

DTCSM	=	Dynamic Tensorial Coefficient Smagorinsky Model
HIT	=	Homogeneous Isotropic Turbulence
CFL	=	Courant-Friedrichs-Lewy condition
RK3	=	Third Order Runge-Kutta
MAC	=	Mean Aerodynamic Chord
SWM	=	Sensor Wall Model
EQWM	=	Equilibrium Wall Model
NEQWM	=	Nonequilibrium Wall Model
LWT1	=	Low-Speed Wind Tunnel 1
ODE	=	Ordinary Differential Equation

Contents

List of Tables

List of Figures

THIS PAGE DELIBERATELY LEFT BLANK

Chapter 1

Introduction

1.1 Background: High-Lift and Transonic Aircraft Flows

The use of computational fluid dynamics (CFD) for external aerodynamic applications has been a key tool for aircraft design in the modern aerospace industry [53, 67]. In takeoff and landing configurations, predicting the maximum lift an aircraft can produce and the associated onset of boundary layer separation encountered at high angles of attack is critically important [64], while in transonic flow regimes, the onset of shock-induced flow separation and unsteady buffeting motion of the shock has serious implications for design [75, 74, 71]. Figure 1.1 shows a representative "flight envelope" of a commercial aircraft which defines the range of load factors and velocities at which safe operation of the aircraft is possible. The limiting design conditions are usually at the corners of this flight envelope and feature large-scale flow separation at stall or unsteady shock oscillations at buffet. Flow solutions from state-of-the-art solvers are unable to routinely comply with the stringent accuracy and computational efficiency requirements demanded by industry in these regions. Leveraging large eddy simulations (LES) with appropriate wall/subgrid-scale models and low dissipation numerical methods suitable for complex geometries on modern compute architectures offers a tractable path towards meeting these accuracy [17] and affordability requirements [8]. 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; JSM) in less than a day of turnaround time with modest resource requirements [29]. Efforts involving the use of LES in these flow regimes represent an important step towards increased use of computations in the aircraft design cycle and aligns with the broader goal of Certification by Analysis (CbA) in which simulations and experiments will work in unison to mitigate the amount of

resource-consuming flight testing that needs to be conducted in order to certify commercial aircraft for safe operation [53].

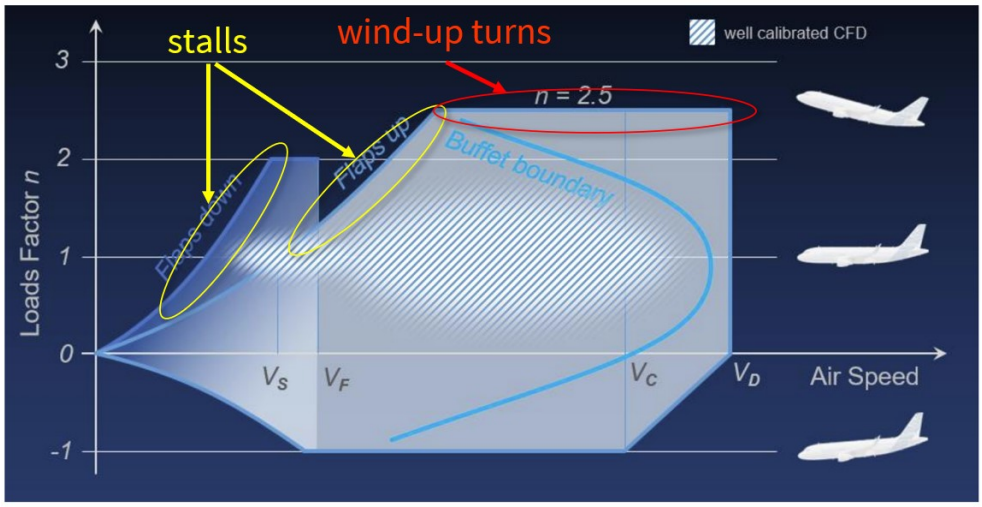

Figure 1.1: Typical velocity/load factor diagram for a commercial aircraft. The area shaded in white refers to the range of safe operating conditions with limiting behavior set by aerodynamic (stalls) and structural (wind-up turns) constraints.

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 [67]: "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/landing and the onset of transonic buffet, 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 fundamentally unsteady simulation techniques such as LES may provide sufficient accuracy. The target date for an LES-based "Technology Demonstration" was estimated to be in the late 2020's. Figure 1.1 shows the results from the Third AIAA High-Lift Prediction Workshop 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 lift, tempering confidence in the predictive capability of steady RANS when flow separation becomes significant and motivating the use of higher fidelity methods for the prediction of these flow conditions.

The accurate simulation of aircraft in high-lift stall conditions is a key pacing item towards CbA

[53] in which high-fidelity simulations would be leveraged in lieu or addition to traditional wind tunnel and flight testing as part of the aircraft certification process. To meet these objectives, the aerospace industry imposes strict accuracy requirements [17] on the predictions of computational methods applied to realistic external aerodynamic flows, including the prediction of maximum aircraft lift coefficient to within 0.03 and the angle of attack at stall to within 1 degree. In aircraft flows, key engineering quantities of interest (QOI's) include the lift/drag forces, pitching moment, pressure distributions, as well as flow separation patterns which are used to identify the stall characteristics of aircraft. Recent efforts at prediction of aircraft high-lift have underscored the deficiency of steady Reynolds Averaged Navier Stokes (RANS) approaches for accurate prediction of QOI's in this flow regime [63], suggesting the need for higher fidelity, eddy-resolving tools such as Large-Eddy Simulation (LES) [29, 40].

Computational cost is the principal uncertainty of when higher fidelity simulation techniques can be adopted. There is little doubt that a first-principles, direct numerical simulation would provide accurate predictions, but the resolution requirements associated with the near-wall, viscous length scales at high Reynolds numbers would be prohibitive. Wall modeled approaches for LES in which the ensemble effect of the unresolved near-wall structures are modeled while eddies that scale with the boundary layer thickness are directly resolved [58, 44, 8] can ameliorate these resolution requirements. And, while LES wall models have been constructed from equilibrium boundary layer arguments, there is theoretical justification that these approximations should be valid if sufficient resolution of the boundary layer thickness, δ, is provided (specifically, the displacement and momentum thicknesses) [8]. 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 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).

1.2 Accomplishments

The following list summarizes the key contributions of this work.

- Establishing robustness with respect to Reynolds number and aircraft configuration of LES methods by systematically assessing the impact of grid resolution, subgrid and wall modeling approach, and geometric modeling in high-lift flows.
- Identification of best practices for accurate and consistent prediction of aircraft maximum lift coefficient and angle of attack at stall to within the strict tolerances required by the aerospace industry.
- Establishing the superiority of dynamic subgrid models relative to static ones in the prediction

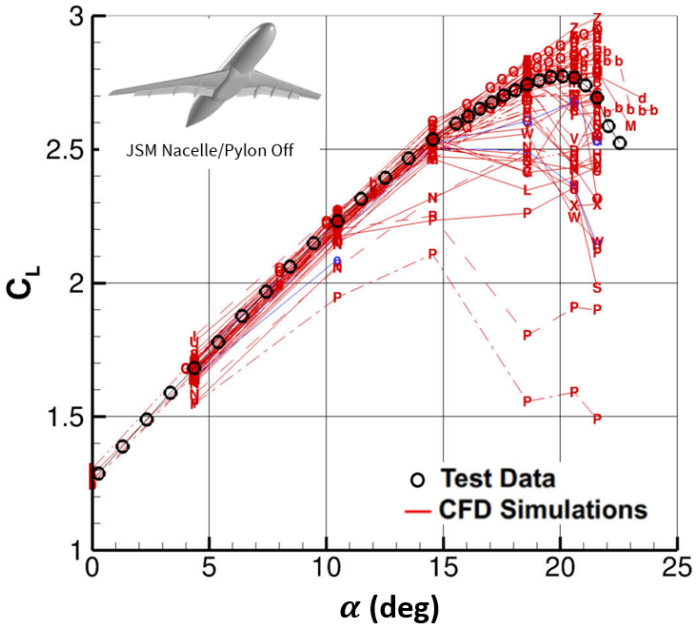

Figure 1.1: Collection of the results of 30+ participants' data submissions to the Third AIAA High-Lift Prediction Workshop. A variety of solvers and gridding strategies are represented, however the majority of the participants used steady RANS tools to generate their submissions. As a whole, the predictive 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 [63].

of juncture and smooth body separation flow regimes in practical aircraft flows.

- Assessment of wind tunnel facility interference effects in validation efforts of half-span sidewall-mounted wind tunnel experiments of aircraft.
- Identification of best practices related to grid topology and wall modeling approach required for accurate prediction of shock location in transonic cruise-type aircraft configurations.
- Achieved the prediction of initial buffet onset in transonic flows, a critical design condition at the corner of the flight envelope.

Chapter 2

Governing equations, physical models, & 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:

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

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

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

where ρ, P, T, and u_i refer to the fluid density, pressure, temperature, and velocity vector respectively. The total energy is defined as $E = \overline{\rho}\tilde{e} + \overline{\rho}\tilde{u}_k\tilde{u}_k/2$, $\tilde{\sigma}_{ij} = \mu(\tilde{T})(\tilde{S}_{ij} - 1/3\tilde{S}_{kk}\delta_{ij})$ 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 $\overline{(\cdot)}$ 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:

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

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

In the present calculations, the static coefficient Vreman eddy viscosity model [76] is used for the JSM free air cases, while the dynamic Smagorinsky model [24, 56] is employed for the wind tunnel and cases with nacelle. The JSM experiment is conducted with a freestream Mach number of 0.172 and a $Re_c = 1.96 \times 10^6$ (based on the mean aerodynamic chord and the freestream 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. The compressible Navier-Stokes equations are temporally integrated due to the presence of this strong flow acceleration.

2.2 Wall modeling

Viscous length scales near a no-slip wall introduce resolution requirements that scale with $l_\nu = \frac{\nu}{u_\tau} = \frac{\nu\sqrt{\rho}}{\sqrt{\tau_w}}$, where $u_\tau = \frac{\sqrt{\tau_w}}{\sqrt{\rho}}$, 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 ($Re_\tau \approx 5,000$ 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 to 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 (although these near wall streaks do exhibit some level of anisotropy [37]. Several estimates of the required number of grid points have been made for realistic aircraft based on experimental correlations for the skin friction [16] or RANS solutions [68] suggesting that direct resolution would be untenable using available supercomputers until 2050-2070. 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 modeled, referred to as a wall modeled LES. In this limit, the grid point count for the turbulent flow regime would scale linearly with respect to the chord-based Reynolds number, $N_{grid} \propto Re_c$, which would potentially be feasible using present computing capacity.

The most common approach for the near-wall modeling is to assume that the thin boundary layer equations are valid. Assuming a local balance between the pressure gradient/streamwise convection and that the spatio-temporal resolution of the LES is large compared to 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:

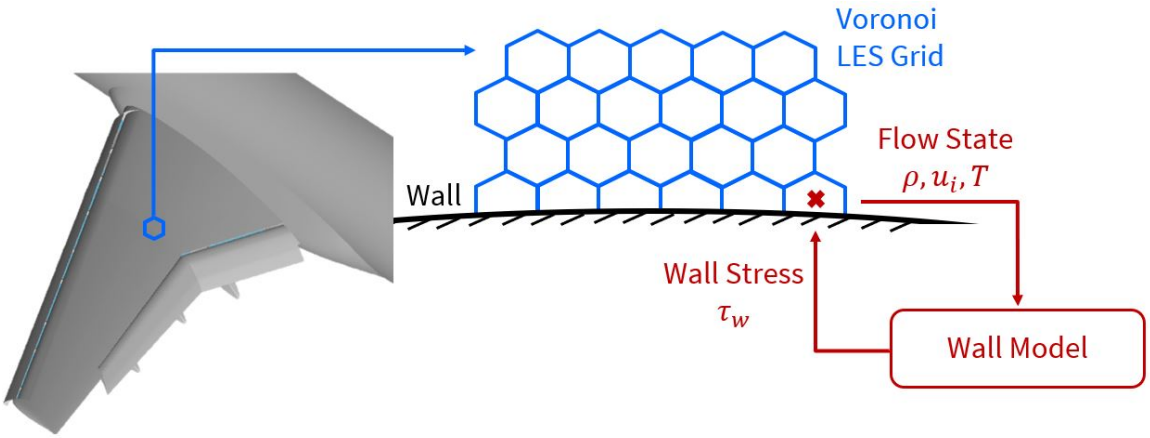

Figure 2.1: Schematic showing the the procedure for wall flux modeling in LES. The LES solution (blue) supplies the top boundary condition to an auxiliary wall model (red), whose role is to deliver wall fluxes to the LES in the form of a Neumann boundary condition coming from a solution to the wall model equations.

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

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), Eqn. 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 2.1. For additional details regarding the construction and implementation of wall models for LES, the reader is referred to recent reviews [58, 44, 8]. Details of the particular wall modeling methodologies employed in the present study can be found in [45].

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 based estimates were designed based around reasoning 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), or finite span phenomena (e.g., tip vortices). Second, many assumptions in the turbulent 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 modeled LES calculations that suggest that the equilibrium approximations are sufficient for flows over airfoils/aircraft models [77, 7, 57]. 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 [8] 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 Subgrid-scale modeling

Calculations in this work use a variety of SGS closures, including eddy viscosity-based models that invoke the Boussinesq assumption involving the alignment of the SGS stress tensor with the filtered strain (i.e. the static coefficient Vreman and Dynamic Smagorinsky, DSM, models), and a non-Boussinesq model referred to as the Dynamic Tensorial Coefficient Smagorinsky Model (DTCSM). The dynamic models account for the deviatoric part of the SGS stress tensor (Eq. 2.5) using either a single or tensorial coefficient eddy-viscosity model with the coefficient computed using a variation of the procedure of [24] and [49] based on local time-averaging (with time scale set using an inverse of the local strain rate) to regularize the least-squares problem. The isotropic component of τ_{ij}^{sgs} is absorbed into a modified pressure.

The single coefficient DSM model is defined as

$$\tau_{ij}^{sgs} - \frac{\tau_{kk}^{sgs}}{3}\delta_{ij} = \tau_{ij}^{DSM} = -C\Delta^2\bar{\rho}\tilde{S}_{ij}|\tilde{S}| \tag{2.7}$$

where Δ is a measure of the grid cell size, a C is a constant which is unknown *a priori*. The combined coefficient $C\Delta^2$ is solved for using the locally/temporally-averaged least-squares Germano identity. The tensorial coefficient eddy-viscosity model is the DTCSM of [2], where the Boussinesq approximation is relaxed and is defined as

$$\tau_{ij}^{sgs} - \frac{\tau_{kk}^{sgs}}{3}\delta_{ij} = \tau_{ij}^{DTCSM} = -(C_{ik}\tilde{S}_{kj} + C_{jk}\tilde{S}_{ki})\Delta^2\bar{\rho}|\tilde{S}| \tag{2.8}$$

with the tensorial coefficients defined as follows

$$C_{11} = C_{22} = C_{33}; C_{ij} = -C_{ji} (i \neq j) \tag{2.9}$$

The reasoning behind these modeling choices can be found in [2]. The remaining four coefficients in the DTCSM model are also solved for using the locally temporally averaged least-squares Germano identity.

An alternate SGS modeling approach involves the use of a static model coefficient calibrated based on a reference flow that is not adjusted in space and time based on the resolved flowfield. An example of such a model is the eddy viscosity-based static Vreman model [76] which is also used selectively in this thesis. The Vreman model is defined as

$$\tau_{ij}^{sgs} - \frac{\tau_{kk}^{sgs}}{3}\delta_{ij} = \tau_{ij}^{Vreman} = -2\nu_t S_{ij} \tag{2.10}$$

where $S_{ij} = \frac{1}{2}\partial_i \tilde{u}_j + \frac{1}{2}\partial_j \tilde{u}_i$ and the eddy viscosity, ν_t is defined as

$$\nu_t = c\sqrt{\frac{B_\beta}{\alpha_{ij}\alpha_{ij}}} \tag{2.11}$$

with

$$\alpha_{ij} = \partial_i \tilde{u}_j = \frac{\partial \tilde{u}_j}{\partial x_i} \tag{2.12}$$

and

$$\beta_{ij} = \Delta_m^2 \alpha_{mi} \alpha_{mj} \tag{2.13}$$

$$B_\beta = \beta_{11}\beta_{22} - \beta_{12}^2 + \beta_{11}\beta_{33} - \beta_{13}^2 + \beta_{22}\beta_{33} - \beta_{23}^2 \tag{2.14}$$

The key point here is that with the static Vreman model, the model coefficient, $c = 2C_s^2$, is related to the Smagorinsky constant (C_s), which is known based on a calibration to the decay rate of homogeneous isotropic turbulence (HIT). This model is therefore easy to apply in practical LES calculations as it only requires the local filter width and first-order derivatives of the velocity field. Its disadvantage is that the model coefficient is invariant in space and time which may not be robust in more complex flow scenarios.

2.4 Flow solver & numerical methods

The flow solver has been utilized for the investigations presented herein: charLES. charLES is a massively parallel, finite volume, compressible flow solver. It utilizes a low dissipation spatial discretion based on principles of discrete entropy preservation [69, 33, 13] where the fluxes are

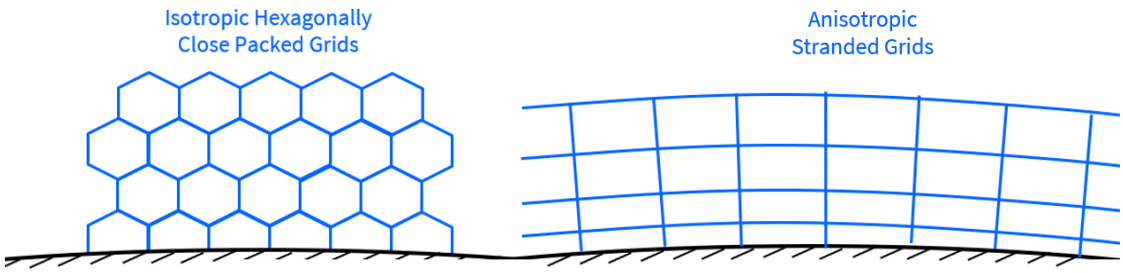

Figure 2.2: Schematic showing the different grid topologies explored in this thesis. On the left, a sketch of the baseline meshing approach in charLES, isotropic HCP is shown, while on the right, a sketch of a prismatic ("stranded") grid of moderate anisotropy is shown.

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 from 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 [30], and the spatial discretization is formally second order accurate. Additional details of the numerical discretization and grid generation can be found in [23], [50], and [9].

One distinguishing feature of the numerical approach in charLES is an emphasis on low-dissipation schemes. Numerical dissipation is known to be extremely detrimental to the simulation of turbulent flows [55], 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 limiting 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.

2.5 Spatial discretization

This section describes the unstructured grids used in charLES to carry out aircraft simulations. The grids are body-fitted and are computed using a Voronoi tessellation of the volume [20]. The generation of the Voronoi diagram in the volume around the aircraft typically uses a staggered point seeding, which results in a Hexagonally Close-Packed (HCP) grid topology. In this approach,

the cells are locally isotropic, and refinement windows are set based on the distance to the nearest boundary, with each window representing a 2× uniform refinement/coarsening relative to the previous resolution window. This meshing approach is advantageous for LES because it is conducive to low numerical dissipation calculations as cell faces always perpendicularly bisect the line connecting adjacent control volumes, resulting in low dissipation calculation of the spatial fluxes. Figure 2.2 shows a sketch of the HCP meshes used predominantly in this work. Also explored were prismatic meshes, hereafter stranded meshes, of moderate anisotropy (up to 4:1). These meshes also represent a valid Voronoi partitioning of the volume, but have the ability to be refined non-uniformly in the wall-normal direction to achieve additional refinement at moderate cost (since the stream/span-wise directions can remain unchanged relative to their coarse grid counterpart). Because realistic turbulent structures have moderate anisotropy ratios, this means of mesh refinement can be used only sparingly.

Chapter 3

Free air large-eddy simulations of high-lift aircraft flows

3.1 Background

In this chapter, a systematic study of the predictive capability of LES across a range of angles of attack (including maximum lift and post-stall regimes) is carried out. The robustness of the predictions to grid resolution, aircraft configuration (JAXA Standard Model and CRM-HL), and Reynolds number are assessed. Integrated engineering quantities of interest, such as lift, drag, and pitching moment are compared against available experimental data, while sectional pressure measurements along the span of the wing are used to corroborate the accuracy of the integrated quantities. In general, 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 3 lift counts of the experimental value. The grid point requirements to achieve this level of accuracy are reduced compared to recent estimates (even for wall modeled LES), with the solutions showing systematic improvement upon grid refinement. For the CRM-HL cases, the solutions show decreasing sensitivity to the grid with each successive refinement level and systematically approach the experimental lift coefficient data as the grid is refined, with the 1.5 billion CV case showing excellent agreement with the corrected experimental data. Turnaround times on the order of a day are made possible even at this grid resolution level in part by algorithmic advances made to leverage graphical processing units (GPU's). The results presented herein suggest that this combined approach (meshing, numerical algorithms, modeling, efficient computer implementation) is on the threshold of readiness for industrial use in aeronautical design. A key development in these simulations is establishing robustness of the methods to Reynolds number and geometry relative to previous studies which focused exclusively on the JAXA Standard Model. Also described in this chapter are subgrid-scale (SGS) modeling sensitivities on two grids for the JSM case, one numbering 10 million control volumes (Mcv) and

another numbering 42 Mcv. Two eddy-viscosity-based SGS modeling approaches are considered, the static coefficient Vreman model [76] and the dynamic coefficient Smagorinsky model [24] in this context, while in the CRM-HL flow, a novel non-Boussinesq model is employed at low angles of attack to assess its ability to predict smooth body separation over the flap [2]. The details of these modeling approaches were described in Section 2.3. All cases in this chapter were run in a free air setting with the appropriateness of this assumption being revisited in the next chapter.

3.2 Results

3.2.1 JAXA Standard Model

In this section we detail the setup and results for the simulations of the JAXA Standard Model, which was the focus of the Third AIAA High-Lift Prediction Workshop (HLPW3). Computational results are compared with the experimental data originally collected by [82] and later expanded upon by [83, 84]. In this chapter, the aircraft is considered only in free air setting while in the subsequent chapter a domain that represents the geometry of the JAXA experimental campaign (i.e. includes wind tunnel walls and sidewall offset) is simulated. The free air simulation results for the integrated lift, drag, and moments are compared against 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 against experimental oil film visualizations and surface pressure measurements, but no wind tunnel corrections are available for the surface data.

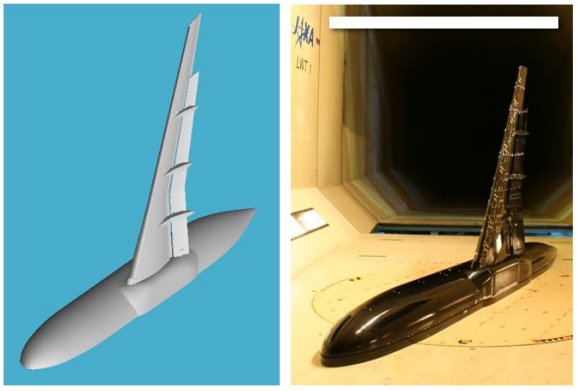

(a) (b)

Figure 3.1: The computational representation (a) of the JAXA Standard Model and (b) the experimental wind tunnel test apparatus, with white scale bar in the upper right corner of the experimental image corresponding to 1 wing semi-span length of 2.3 m.

The computational geometry and domain is shown alongside the experimental model in Figure 3.1 with a slice depicting the mesh utilized from the charLES flow solver in Figure 3.2. The grid shown

corresponds to a total grid count of 32 million control volumes, which was sized to fit approximately 10 points across the trailing edge boundary layer thickness 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 majority of grid points for the baseline grid resolution 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). Sensitivity of the results to grid refinement will also be explored.

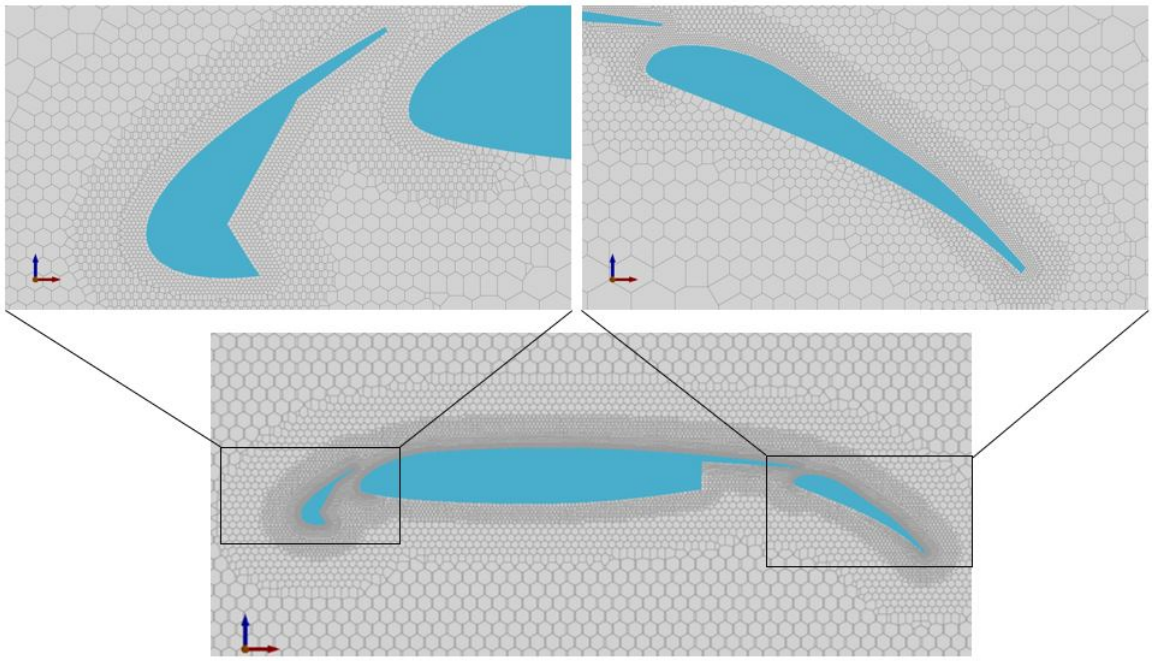

Figure 3.2: Spanwise slices of the baseline 32 Mcv grid used for JSM free air calculations including zoomed views of the slat/main element leading edge and the main element trailing edge/flap. This Voronoi grid is based on a hexagonally close packed point arrangement.

Baseline Calculations

The integrated forces are shown versus the angle of attack for the free air configuration using the baseline 32 Mcv resolution in Figure 3.3. Error bars in the experiment denote the quoted experimental repeatability at select angles of attack. This first-pass simulation represents an extension of the approach of [45] across the angle of attack curve and represents a reasonable prediction of the integrated forces. Notable deficiencies of these simulations include: the lift over-prediction at low angles of attack ($\Delta C_L \approx 0.05 - 0.08$, Figure 3.3a, the over-prediction of drag coefficient (Figure 3.3b, and the lack of a distinction pitching moment break post-stall Figure 3.3c. It is encouraging that the lift predictions from the charLES solutions at this resolution predict both the maximum

lift and the onset of the post-stall regime appropriately.

The drag polar is shown in Figure 3.3b. 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 to corrected wind tunnel data [63]. A reason for this was not cited by AIAA, 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 against uncorrected wind tunnel data. These will be discussed in the subsequent chapter.

Figure 3.3c shows the pitching moment coefficient for the JSM configuration, whose reference center is at $x \approx 2.38$ m, $y = 0$, $z = 0$, or at about mid-chord of the inboard wing. The prediction of the pitching moment is of equal engineering 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 LES 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. [36] suggests 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 Section 4.3.2 where the tunnel geometry is included in the calculations.

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 considered in section 3.2.1 suggested that the $C_{L,max}$ and post-stall regimes could be approximately captured with relatively coarse resolutions. To assess the robustness of these predictions, a grid refinement study is conducted at $\alpha = 18.58°$ in which both a coarser 9 Mcv and a finer 157 Mcv mesh is considered.

Sectional pressure predictions at midspan and near the wingtip from a finer resolution compare favorably with the experimental measurements at $\alpha = 18.58°$. This suggests that the accuracy of the lift predictions at stall are not due to fortuitous error cancellation along the wing. Sectional pressures from a series of grids at $\alpha = 18.58°$ at various spanwise locations are shown in Figure 3.4. 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

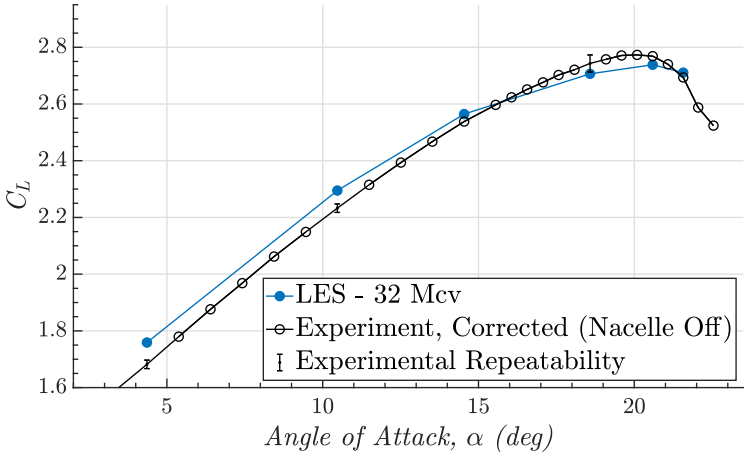

(a) C_L vs. α

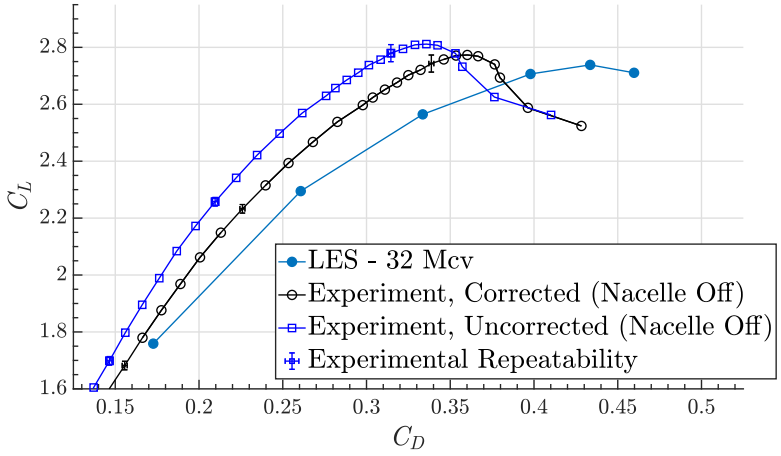

(b) Drag Polar

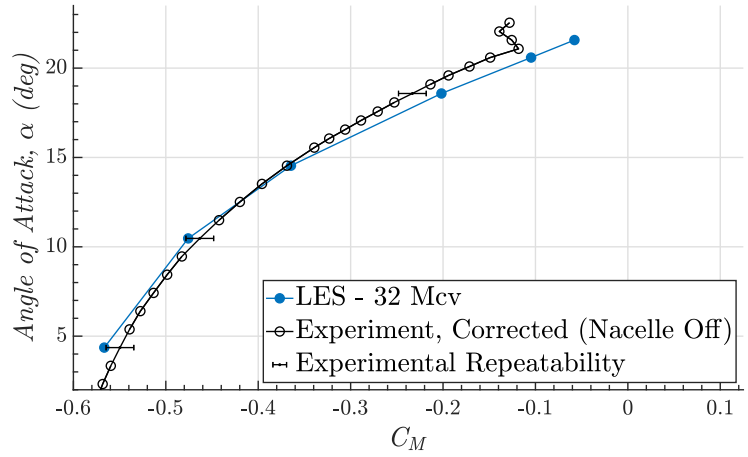

(c) α vs. C_M

Figure 3.3: (a) Lift curve, (b) drag polar , and (c) pitching moment coefficient from baseline resolution simulations compared against experimental measurements from JAXA.

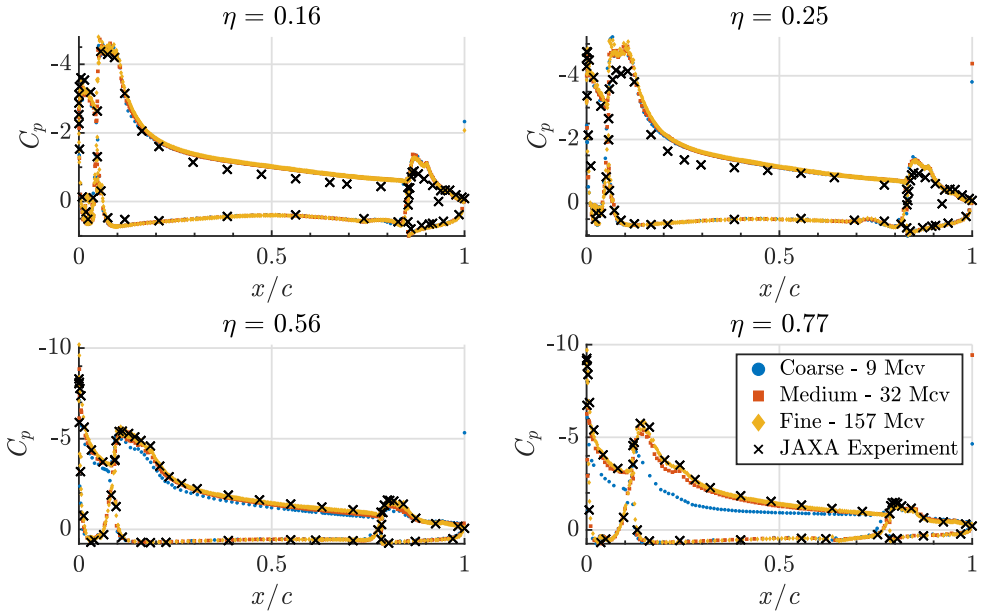

Figure 3.4: Sectional pressure measurements from a grid refinement sequence using the charLES solver at $\alpha = 18.58°$.

and main elements are notably reduced compared to the experimental measurements on this coarse 9 Mcv mesh. 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 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 favorably with the experimental measurements. Figures 3.5 and 3.6 show the lift and drag coefficients corresponding to this grid refinement sequence near stall. It is clear that the lift is approaching the experimental observations from below, consistent with the excess separation seen in the coarse grid pressure data, and the the solution is changing less with each refinement. A more complete picture of the sensitivities associated with grid resolution will be drawn for the CRM-HL high-lift configuration in which a full alpha sweep of the coarse and fine grids will be completed. The question of grid-convergence will also be explored by means of a grid numbering more than 1 Bcv.

Although this limited grid refinement study built confidence that the prediction of lift near stall is for the right reasons, further work is needed to address the key deficiencies of the baseline free air simulations, which include the lift over-prediction at low angles of attack (addressed via SGS modeling), the drag over-prediction (likely a deficiency of experimental correction techniques rectified by running simulations that include the wind tunnel and compare against uncorrected experimental

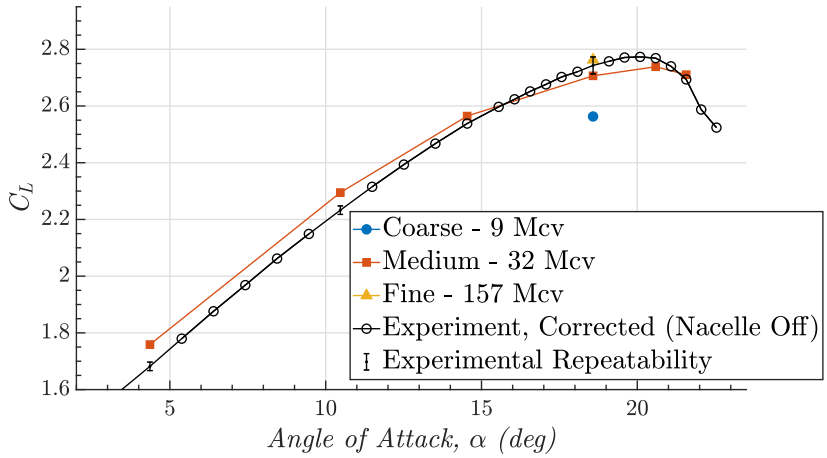

Figure 3.5: The lift curve predicted by the LES simulations in free air. A grid refinement sequence consisting of a coarse, medium, and fine mesh is completed at the stall condition near $C_{L,max}$ ($\alpha = 18.58°$).

data), and the lack of pitching moment shift (also thought to be tied to wind tunnel installation effects including tunnel blockage the the build-up of a sidewall boundary layer).

Remarks on computational cost

We take a brief aside in this section to discuss the computational cost of the free air JSM calculations in this section. The previous section 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 their computational cost results in turnaround times that are tenable. Table 3.1 shows the computational cost for the charLES flow solutions on the medium and fine grids (where the calculations are of sufficient accuracy near stall) for the the free air configuration. 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. The grids are self-similar in that the finer grid is produced by isotropic refinement of the coarser mesh in all directions in the vicinity of the solid boundaries, which results in a halving of the time step due to CFL restrictions of the explicit RK3 time advancement scheme. The cost estimates are produced for a time horizon of $30c_{mac}/U_\infty$ (convective flow through times over the mean aerodynamic chord) which was deemed sufficient to obtain converged statistics. Time horizons of up to $40c_{mac}/U_\infty$ were used for averaging solutions at angles of attack in the post-stall regime due to the presence of low frequency oscillations in the aerodynamic loading caused by 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) or 96 GPUs (NVIDIA Tesla V100s;

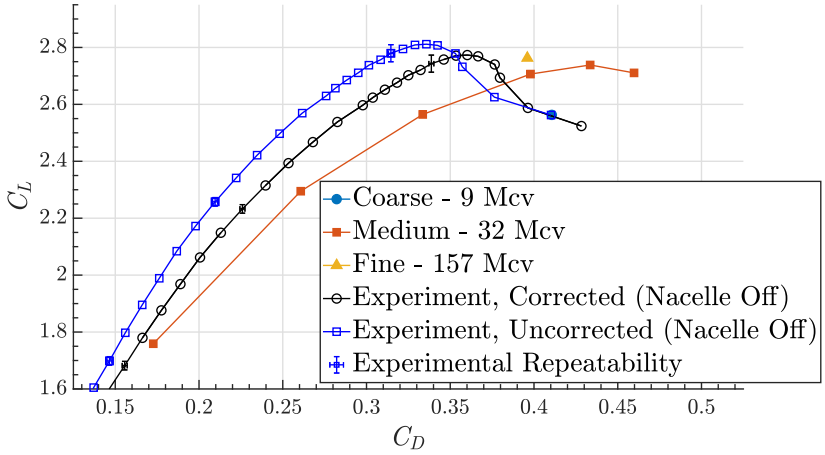

Figure 3.6: The drag polar predicted by the LES simulations. A grid refinement sequence consisting of a coarse, medium, and fine mesh is completed near $C_{L,max}$ ($\alpha = 18.58°$).

Oak Ridge National Laboratory's Summit cluster). This resource allocation is selected as it can be considered reasonably accessible to simulation practitioners. At this level of resource availability, turnaround times of less than an hour can be obtained for the medium resolution (facilitated by either CPU or GPU architectures) or less than half of a day on the fine grid (when GPU accelerated). Turnaround times of less than a day were identified by the NASA CFD Vision 2030 report [67] as a critical threshold that would enable use of LES within the design cycle to complement lower fidelity techniques and experimental campaigns.

Table 3.1: Computational cost summary for a grid sequence using the charLES solver for the JAXA Standard Model configuration at $\alpha = 18.58°$.

Cell Count (M)	**Time Step (ns)**	**CPU core-hours/ flow pass (k)**	**Wall Clock Time to 30 Flow Passes (Using 2,000 Ivy Bridge Intel CPU cores)**	**Wall Clock Time to 30 Flow Passes (Using 96 NVIDIA V100 GPU's)**
32	580	1.4	21 hours	48 min
157	290	12.3	7.5 days	7 hours

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 $\approx 10^6$ 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 hours.

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 [3]. The RANS calculations are conducted on a 109M cell grid and warm started (initial guesses of the solution are provided from a nearby angle of attack) using a one equation Spalart-Allmaras turbulence model. These simulations take $\approx 60 \times 10^3$ CPU core hours versus $\approx 40 \times 10^3$ and $\approx 360 \times 10^3$ core hours for the medium and fine grid cases, respectively. The medium grid LES costs are comparable to the RANS solutions and fine grid calculations show an increase of a factor of 6. As RANS simulations are routinely conducted, this cost differential is sufficiently small that the use of targeted LES calculations as a companion to RANS simulations would be presently feasible in industry. The overall simulation cost should also consider the time required for the generation of the computational grid. This cost is negligible compared to 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 $\mathcal{O}(10^{10})$ cells on 40×10^3 cores in less than 30 minutes [80].

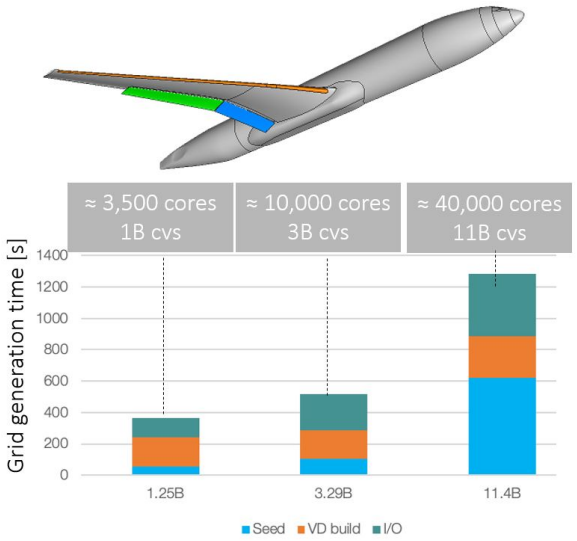

Figure 3.7: Grid generation time in seconds for the charLES solver from the 2^{nd} mesh generation workshop for a representative high-lift geometry (the High-Lift Common Research Model). The compute time is broken down between the seeding of the volume domain with cells, building of the Voronoi diagram, and input/output operations (such as writing of the grid to disk). Massive grids ($\mathcal{O}(10^{10})$ cells) are routinely built within 30 minutes [80].

The JAXA Standard Model 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 to flight conditions. The computational cost can be extrapolated based on the scaling arguments advanced by [16] 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 on the order of a few days would be possible on ≈ 100 GPU's for $Re_c \sim 10^7$. Turnaround times

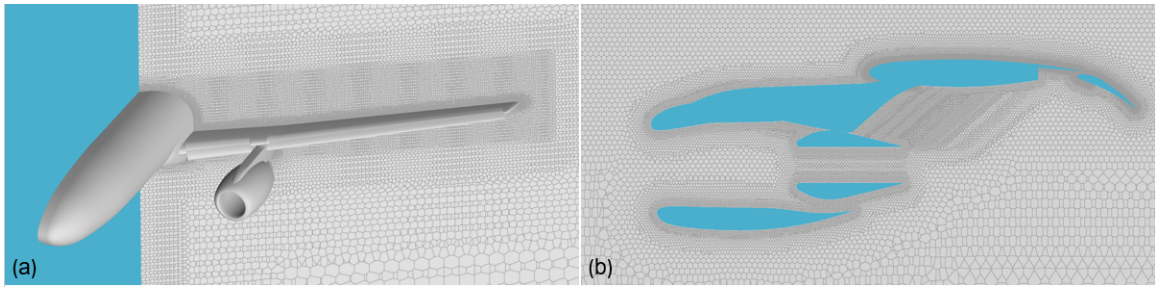

Figure 3.8: Slices of the hexagonal close-packed grid from the 42 Mcv grid along the span of the wing (a) and through the nacelle/pylon region (b). A refinement box around the wing adds grid support for off-body flow features such as viscous wakes and shear layers.

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). This Reynolds number regime is in the range of those attained in experimental facilities where considerable data is presently available. Reynolds number scaling will be further evaluated *aposteriori* when the CRM-HL calculations are discussed in Section 3.2.2.

Subgrid-scale modeling sensitivities: dynamic Smagorinsky vs. static Vreman

Aimed at addressing several of the issues found in the baseline static Vreman simulations performed in free air discussed above, a sensitivity assessment to SGS modeling approach is made. Slices of the two grids used for this study are shown in Figures 3.8 and 3.9. The 42 Mcv grid is obtained from the 10 Mcv grid by adding a layer of twofold refined isotropic HCP cells adjacent to the solid boundary of the domain, accounting for the additional 32 Mcv in the finer grid. An *a posteriori* estimate of the number of points per trailing-edge boundary-layer thickness on the inboard part of the wing revealed $\delta/\Delta \approx 10$ on the 10 Mcv grid and $\delta/\Delta \approx 20$ on the 42 Mcv grid, with δ the local boundary-layer thickness and Δ the local cell length scale. In viscous units, the first cell achieved a maximum $y^+ \approx 320$ on the 10 Mcv grid and a maximum $y^+ \approx 160$ on the 42 Mcv grid at the wing leading edge, which is within the suitable range for wall-modeled LES.

The discussion of the results begins with the 42 Mcv grid, which was built so as to fit $\delta_{chord}/\Delta \approx 10$, δ_{chord} being the boundary-layer thickness based on the mean aerodynamic chord, estimated using a flat-plate turbulent correlation for the boundary-layer thickness. This estimate proved to be conservative, as mentioned before, as the output of the calculation revealed $\delta_{chord}/\Delta \approx 20$ at the trailing edge on the inboard part of the wing. Exploration of SGS modeling sensitivities was motivated by the inadequate (>5 lift counts of error) prediction of lift in both the linear and near-stall regime by the Vreman SGS model on the 42 Mcv grid. The dynamic Smagorinsky model (DSM) [24] with the modification by [49] was chosen as a viable alternative model because its coefficient is set using a dynamic procedure that leverages information in the resolved scales to adjust the model

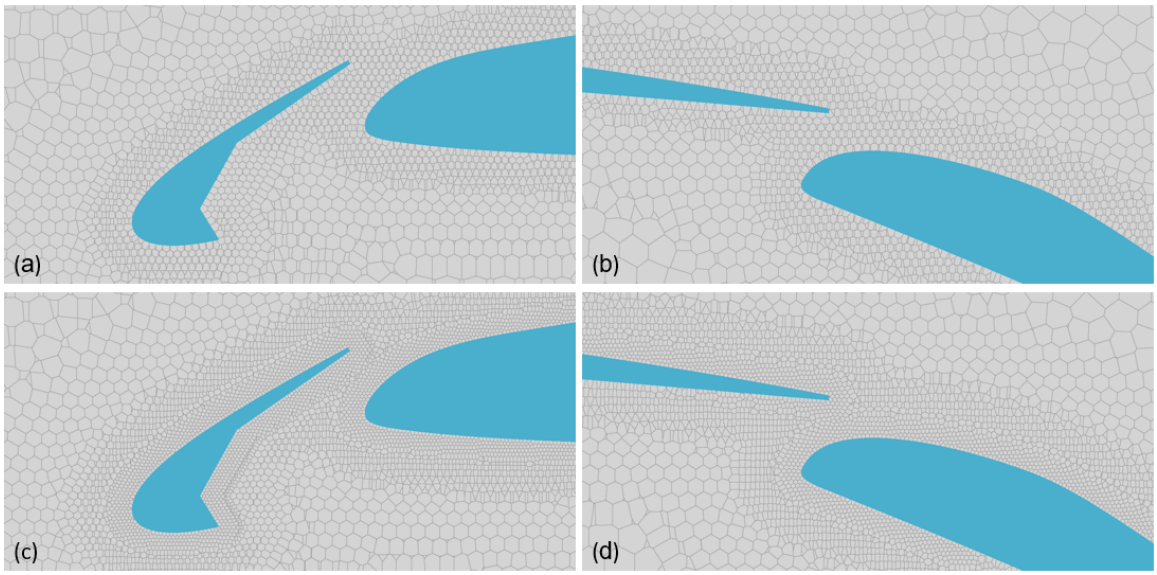

Figure 3.9: Slices of the hexagonal close-packed grid near the slat (a,c) and flap (b,d) for the 10 Mcv (a,b) and 42 Mcv (c,d) grids. The grids differ in that the finer grid contains a layer of refinement in the wall-adjacent viscous region.

coefficient in time and space. In contrast, the static Vreman model uses a coefficient obtained via calibration to homogeneous isotropic turbulence [76] as discussed in Section 2.3. We hypothesized that the dynamic procedure would result in greater predictive capability of the SGS model because the coefficient is free to vary in a way that better accounts for the complex phenomena experienced in the flow over an aircraft. Indeed, in both the linear region of the lift curve shown in Figure 3.10 and in the near-stall region, DSM systematically outperforms Vreman on the 42 Mcv grid. The same improvement holds when considering drag coefficient (Figure 3.11), while improvement in pitching moment (Figure 3.12) is not readily discernible.

Further efforts were made to evaluate the specific regions on the aircraft in which differences between the predictions of the models exist. This may highlight particular aspects of the flow that for which the DSM model outperforms the static Vreman model. To that end, Figure 3.13 is used, which reports the total lift force, as well as the contribution of various components of the aircraft to the total. As expected, the main element of the wing contributes the most significantly to the total lift coefficient, C_L. DSM outperformed static Vreman in both the linear range of the lift curve, lessening the lift over-prediction and in the region near stall, ameliorating the lift under-prediction by Vreman. Figure 3.13 shows that this is because of an improved prediction of the flap separation at low angle of attack and because of an improved characterization of the main element/fuselage flow near stall.

We find that in the linear regime, the lift over-prediction achieved with the static coefficient

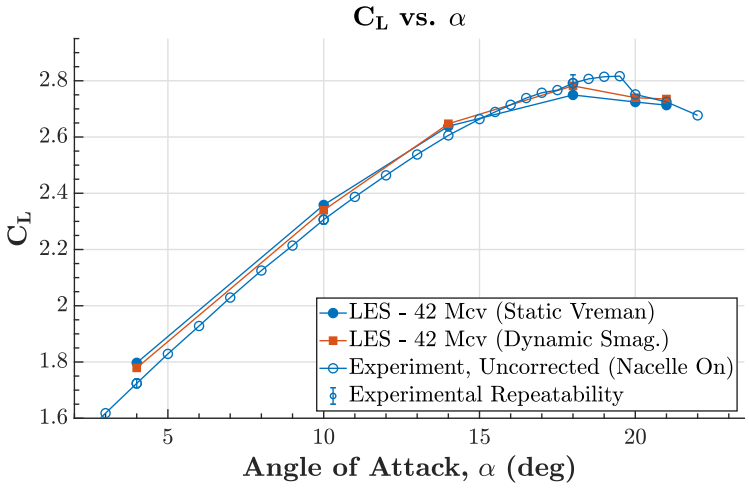

Figure 3.10: Lift coefficient (C_L) plotted against aircraft angle of attack (α) in degrees for the JAXA standard model as compared with the predictions of LES using two SGS modeling approaches: static Vreman and dynamic Smagorinsky. The prediction of lift at stall is improved from an error of 0.08 in C_L with static Vreman to 0.01 with DSM (18°). At low alpha, the error is improved from 0.08 in C_L with static Vreman to 0.04 with DSM (4°). Uncorrected experimental data is plotted for comparison against LES results that include the wind tunnel test section and peniche geometry.

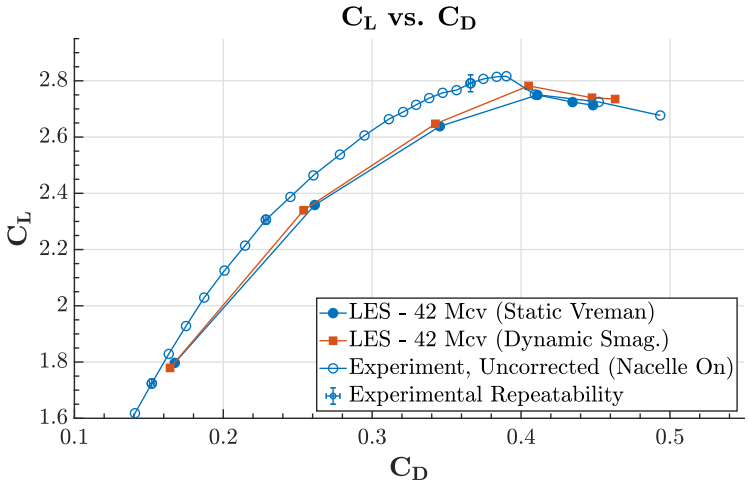

Figure 3.11: Drag polar ($C_L vs. C_D$) of the JAXA standard model (nacelle/pylon on configuration) as compared with the predictions of charLES using two SGS modeling approaches: static Vreman and dynamic Smagorinsky.

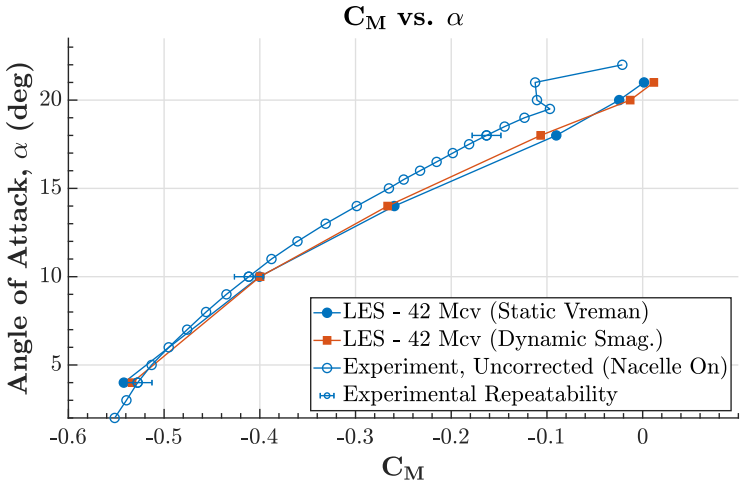

Figure 3.12: Aircraft angle of attack (α) in degrees plotted against aircraft moment coefficient (C_M) for the JAXA standard model (nacelle/pylon on configuration) as compared with the predictions of charLES using two SGS modeling approaches: static Vreman and dynamic Smagorinsky.

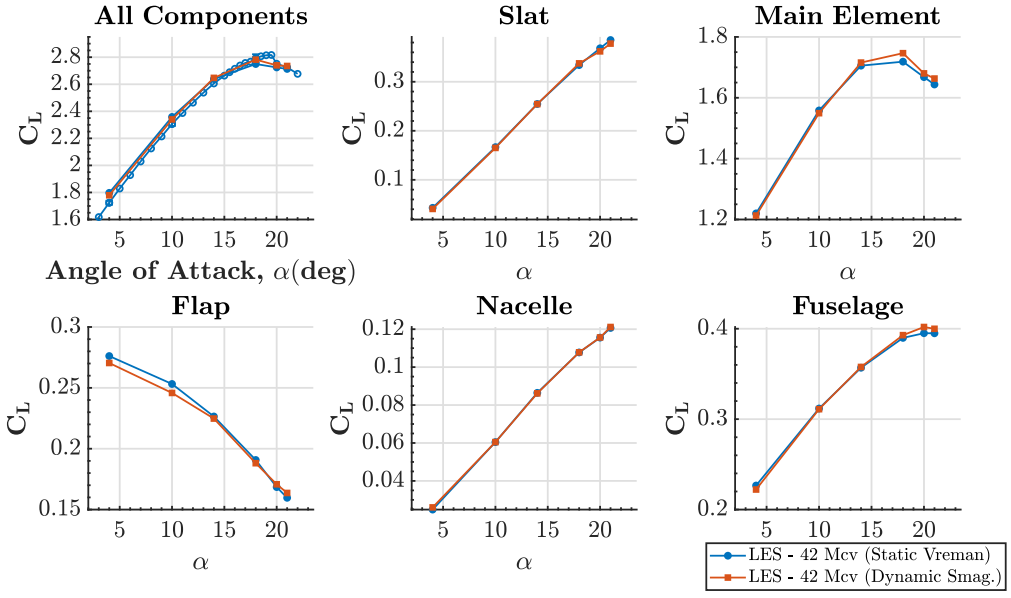

Figure 3.13: Breakdown of the total C_L into contributions from various components of the aircraft, including sub-components of the wing (slat, main element, flap), the nacelle, and the fuselage as predicted by LES using two SGS modeling approaches: static Vreman and dynamic Smagorinsky across all simulated angles of attack, α (degrees).

Vreman model is ameliorated by DSM due to the prediction of less lift on the trailing-edge flaps, which, somewhat counter-intuitively, are most highly loaded when the aircraft is at a low angle of attack, and therefore most likely to separate [14]. We investigate further by considering the sectional pressure coefficients at six spanwise stations along the wing shown in Figure 3.14 and find in general very good agreement between the two modeling approaches and the experimental measurements, except on the flap at the 16% & 25% semispan locations, where the flap suction is over-predicted by both models. Clearly, though DSM is lessening the over-prediction of trailing edge suction on the flap, both modeling approaches continue to predict a delayed separation on the trailing edge flap at low angles of attack. The average velocity contours shown in Figure 3.15, can be used to conclude that indeed DSM predicts lower lift than the Vreman model at $\alpha = 4^\circ$ due to a prediction of increased flap separation, indicated in the figure by the darker blue regions aft of the flap which reflect the region of slow-moving fluid that has separated. The merging of the shear layer from the main element of the wing with the flap wake is also visible in Figure 3.15. This combination of evidence points to the fact that DSM outperforms static Vreman in the linear region of the lift curve due to a prediction of increased flap separation, which is consistent with the trend observed in the experimental pressure measurements. In an absolute sense, however, this correction while in the right direction, is of insufficient magnitude to fully reconcile the prediction of the LES simulations with experimental observations.

Additionally, it is interesting to identify the reason behind the superior performance of DSM relative to Vreman in the near-stall region of the lift curve, in particular the $\alpha = 18^\circ$ condition. In this case, the static Vreman model was deficient in lift relative to the experiment, as shown in Figure 3.10. Figure 3.13 shows that the main element is primarily responsible for the augmented lift that the DSM model predicts. Sectional pressure measurements in Figure 3.16 do not reveal glaring differences between the predictions of the two models, which suggests that the difference in lift is accounted for by a region which is not sampled by a pressure belt, such as the wing/body juncture. Since the wing area is quite large near the wing root, slight differences may integrate out into meaningful differences in C_L. Indeed, when we consider a qualitative comparison between an experimental oil-flow image and average skin friction contours in the wing/body juncture region, we find differences between the two models (Figure 3.17). In particular, the DSM model predicts a larger region of flow turning at the juncture, which extends almost to the first inboard pressure belt (dashed line), in contrast to static Vreman, in which the region of flow turning is very narrow. The prediction of DSM is closer to that observed in the experiment, as the region of flow turning visible in the oil flow also extends nearly to the first pressure belt.

An additional study was performed on a coarser grid because the influence of SGS model was thought to increase with coarsened grid resolution. The grid was coarsened by removing the finest layer of grid refinement at the wall in the 42 Mcv grid, resulting in a grid of 10 Mcv. Surprisingly, the solution on the coarser grid showed less sensitivity to the choice of SGS model than on the finer grid, potentially because at these resolution levels, the grid is so coarse that the input to the SGS

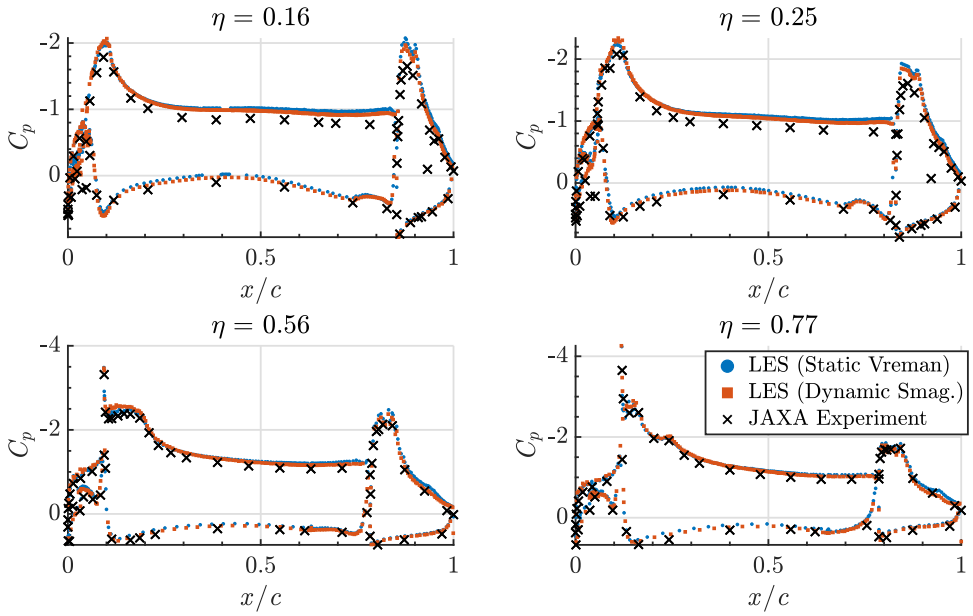

Figure 3.14: Comparison of sectional pressure coefficient, C_p, at four spanwise pressure belts (16% semispan, 25% semispan, 50% semispan, and 77% semispan) at $\alpha = 4°$ as predicted by LES using two SGS modeling approaches: static Vreman and dynamic Smagorinsky on the same grid. Data are plotted against the local chordwise coordinate normalized by the local chord, x/c. Differences are subtle, with dynamic Smagorinsky predicting slightly less suction on the main element of the wing and the flap at the inboard stations (16% semispan and 25% semispan).

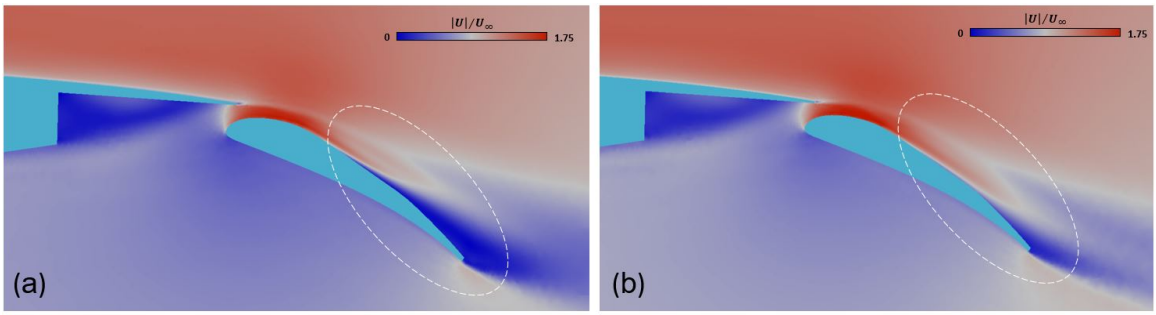

Figure 3.15: Velocity magnitude slice through the flap near the Yehudi break showing increased flap separation predicted by (a) dyanmic Smagorinsky relative to (b) static Vreman.

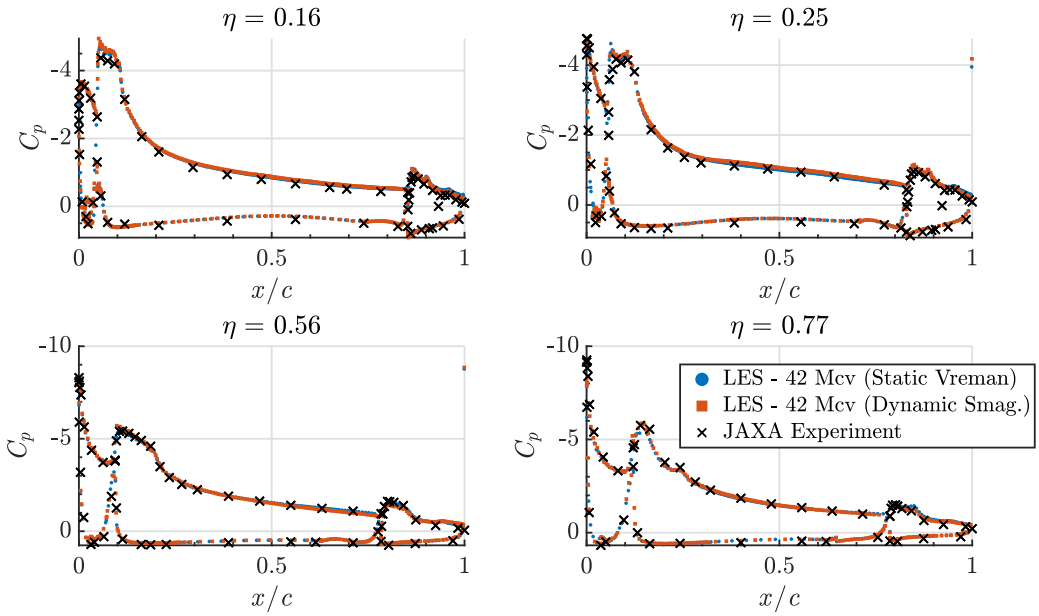

Figure 3.16: Comparison of sectional pressure coefficient, C_p at four spanwise pressure belts (16% semispan, 25% semispan, 50% semispan, and 77% semispan) at $\alpha = 18°$ as predicted by LES using two SGS modeling approaches: static Vreman and dynamic Smagorinsky on the same grid. Data are plotted against the local chordwise coordinate normalized by the local chord, x/c. Differences are not readily discernible, suggesting the improved performance of dynamic Smagorinsky near stall is accounted for by regions that are under-sampled by the experimental pressure belts.

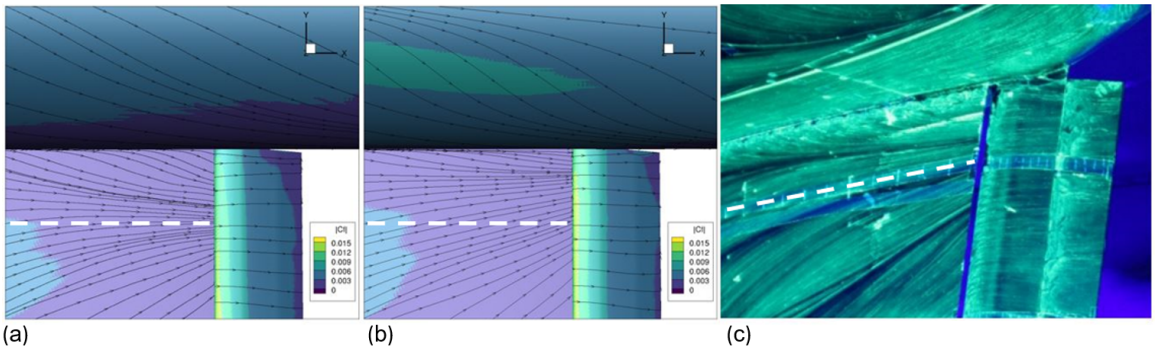

Figure 3.17: Average skin friction contours (a-b) from LES predicted using two different SGS models and an experimental oil flow visualization (c) at a near-stall angle of attack ($\alpha = 18°$). The dashed white line indicates the location of the first pressure belt. The (a) Dynamic Smagorinsky model better reproduces the character of the flow at the wing/body juncture observed in the experiment than does (b) static Vreman.

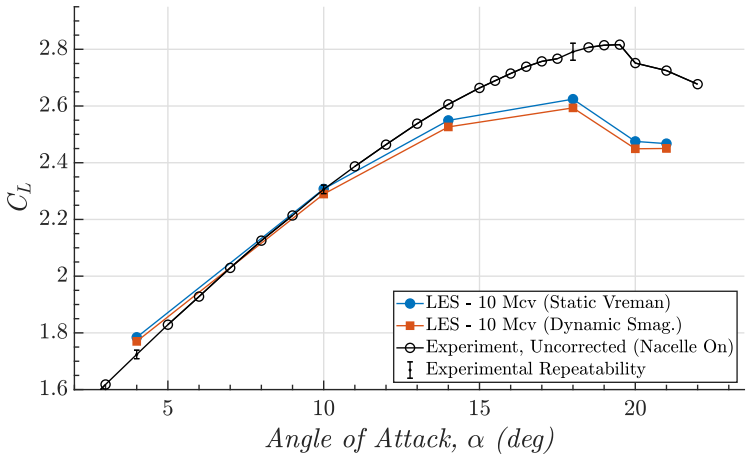

Figure 3.18: Lift coefficient (C_L) plotted against aircraft angle of attack, α. in degrees for the JAXA standard model (nacelle/pylon on configuration) as compared with the predictions of LES using two SGS modeling approaches: static Vreman and dynamic Smagorinsky on a coarse 10 Mcv grid.

model from the resolved LES field is too crude to expose model performance in a meaningful way. Figure 3.18 shows that the lift curves predicted by both modeling approaches are deficient in lift and nearly identical in character across the lift curve. Sectional pressure coefficients in Figure 3.19 reveal that the grid refinement primarily leads to an improved prediction of lift because of an improved characterization of the flow on the wingtip.

A high-lift aircraft configuration at two resolution levels was simulated to assess the sensitivity of the prediction of integrated forces and moments to the choice of SGS model. The two grids numbered 10 Mcv and 42 Mcv and contained approximately $\delta_{chord}/\Delta = 10$ and $\delta_{chord}/\Delta = 20$, meaning that at the trailing edge of the main element, between 10 and 20 points were contained per boundary-layer thickness in all directions (the grids are isotropic). The simulations on the 42 Mcv grid highlighted the superior performance of DSM, whose model coefficient is dynamically set and varies in time and space, relative to the static coefficient Vreman model, whose model coefficient is set based on a calibration to HIT, as described in Section 2.3. The superiority was not limited to a one particular flow phenomenon as DSM outperformed static Vreman in the characterization of both the trailing edge flap separation at low α and the wing/body juncture flow near $C_{L,max}$, suggesting it has a broader predictive scope than static Vreman. Observations of the flap and juncture region flow obtained with the DSM model were in better agreement with experimental observations. The improved predictions of DSM came at a cost increase of ≈15% on a core-hour per flow-pass basis relative to static Vreman, owing to the additional computational expense of the dynamic procedure. The belief of the author is that this improvement in accuracy justifies this cost.

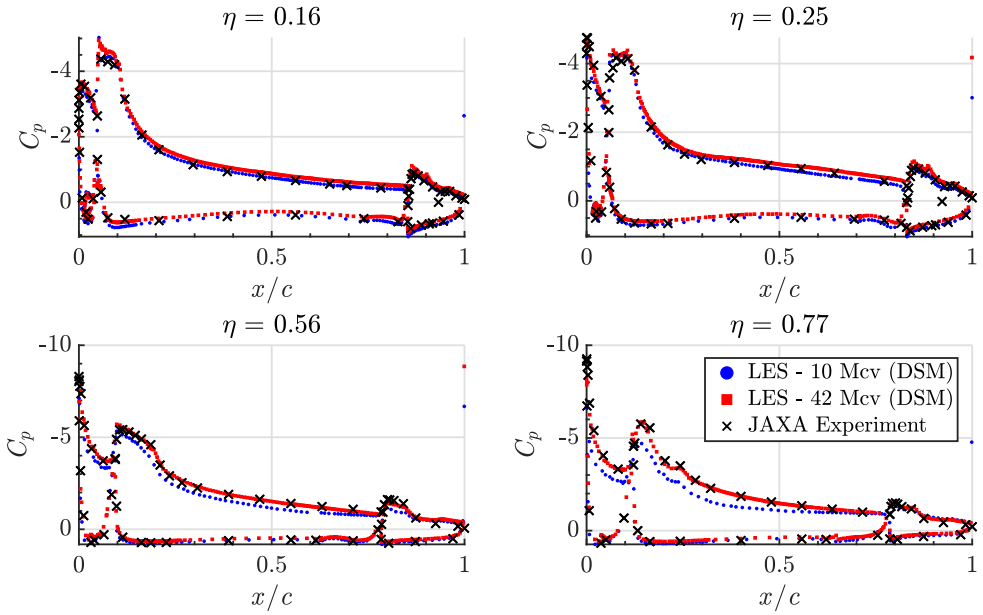

Figure 3.19: Comparison of sectional pressure coefficient, C_p, at four spanwise pressure belts (16% semispan, 25% semispan, 50% semispan, and 77% semispan) at $\alpha = 18°$ as predicted by charLES using the dynamic Smagorinsky SGS model on a coarse and a refined grid. Data are plotted against the local chordwise coordinate normalized by the local chord, x/c. The refined grid improves the predictions, particularly near the wingtip.

3.2.2 High-Lift Common Research Model

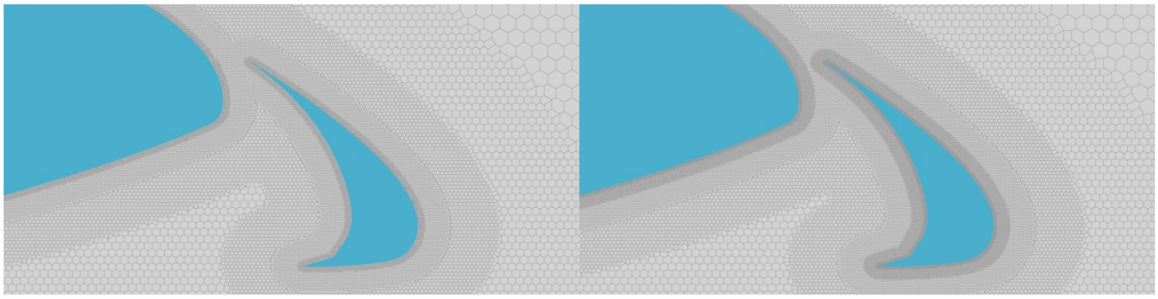

Figure 3.20: Slices of the hexagonal close-packed grid from the Xfine grid through the leading edge slat (a) and from the XXfine grid at the same location (b).

In this section, free air calculations of the NASA CRM-HL aircraft configuration are performed. The CRM-HL which was the focus of the Fourth AIAA High-Lift Prediction Workshop and is set to be the benchmark flow in the high-lift aerodynamics community for the foreseeable future. The simulation geometry is designed to closely match the test article and includes the bracketry associated with deployed high-lift devices (flaps/slats) as well as a flow-through nacelle mounted on the underside of the wing about a third of way along the wing semispan. The reference Reynolds number is 5.49 million, based on the mean aerodynamic chord (MAC), and the freestream Mach number is 0.20, both chosen to match the experimental conditions of [22]. The design choices made in creating the CRM-HL are described in detail by [42], but in general are meant to be representative of a modern commercial airliner such as the Boeing 777 in landing configuration.

The treatment of boundary conditions is as follows: a uniform plug flow is prescribed at the inlet which matches the test conditions of [22]. A viscous wall-modeled algebraic equilibrium boundary condition is applied on the aircraft surface. At the outflow, we have prescribed a characteristic non-reflecting characteristic boundary condition, in which the freestream outlet pressure is specified [59]. The symmetry plane is treated as a slip wall.

Slices of the grid are shown in Figure 3.20. The grids are generated by computing a Voronoi diagram of the water-tight volume around the aircraft [20], using a staggered point seeding, which results in the HCP topology shown in the figure. The cells are locally isotropic, and refinement windows are set according to the distance to the nearest boundary. This topology is conducive to low numerical dissipation calculations as cell faces always perpendicularly bisect the line connecting adjacent control volume centroids, which has favorable implications for the calculation of spatial fluxes.

The 1.5 Bcv XXfine grid in Figure 3.20 is obtained from the 384 Mcv Xfine grid by inserting a 5-cell-thick layer of twofold-refined isotropic HCP cells adjacent to the solid boundary of the domain,

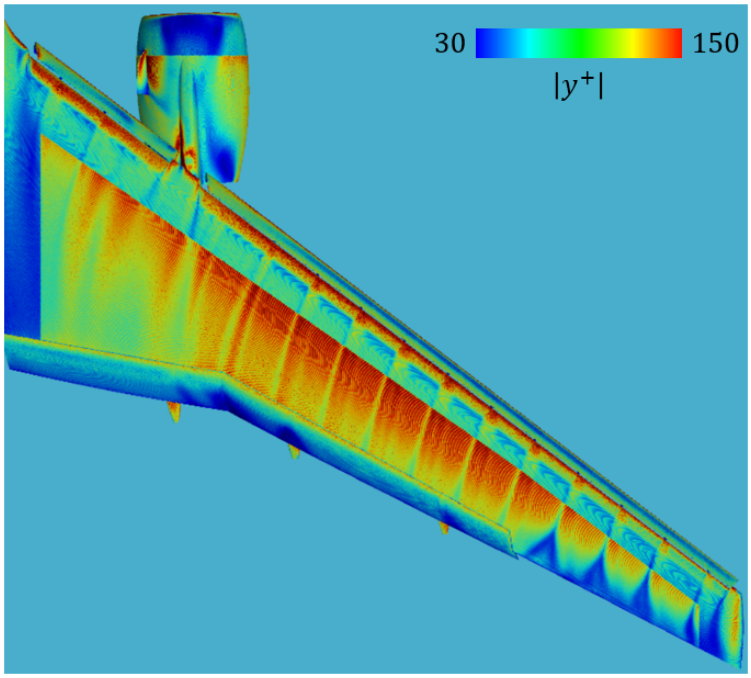

Figure 3.21: Contour plot of y^+ on the Xfine grid for the free air charLES simulations at an angle of attack of 19.57°. The contour levels vary from 30 to 150 in viscous units. Abrupt changes in the value of y^+ reflect changes to the local grid spacing associated with targeted refinement regions at the leading edges of the wing/nacelle and at the wing root/tip.

accounting for the additional ≈ 1.1 Bcv. All grids (Coarse-XXfine) were refined in this manner and the refinement sequence can be considered a grid family. The finest layer of grid adjacent to the body in the XXfine mesh is such that there are 4096 points per MAC at that resolution level. Figure 3.21 shows the achieved resolution of the Xfine grid in viscous units. This figure shows the y^+ associated with the first cell centroid, which coincides with the wall-model/LES exchange location. In viscous units, the first cell achieved a maximum $y^+ \approx 150$ on the Xfine grid and a minimum $y^+ \approx 30$ on the same grid. Clearly the y^+ varies significantly over the wing as the flow encounters favorable/adverse pressure gradients, flow separation, etc. In general, we conclude that the Xfine and XXfine grids have resolutions that are adequate for WMLES, as the first cell lies approximately within the logarithmic region of the boundary layer (the XXfine mesh was refined uniformly by $2\times$ in all directions relative to the Xfine mesh and so the y^+ values are approximately half of those achieved on the Xfine mesh).

Figure 3.22 shows the turbulent flowfield computed on the Fine mesh at 17.05°, just ahead of stall, through the lens of the Q-criterion [34], shaded according to the normalized velocity. The flowfield is rich, with multiscale flow phenomena readily observable, from the strong chine vortex and wing juncture separations to the fine-scale boundary layer turbulence over the wing.

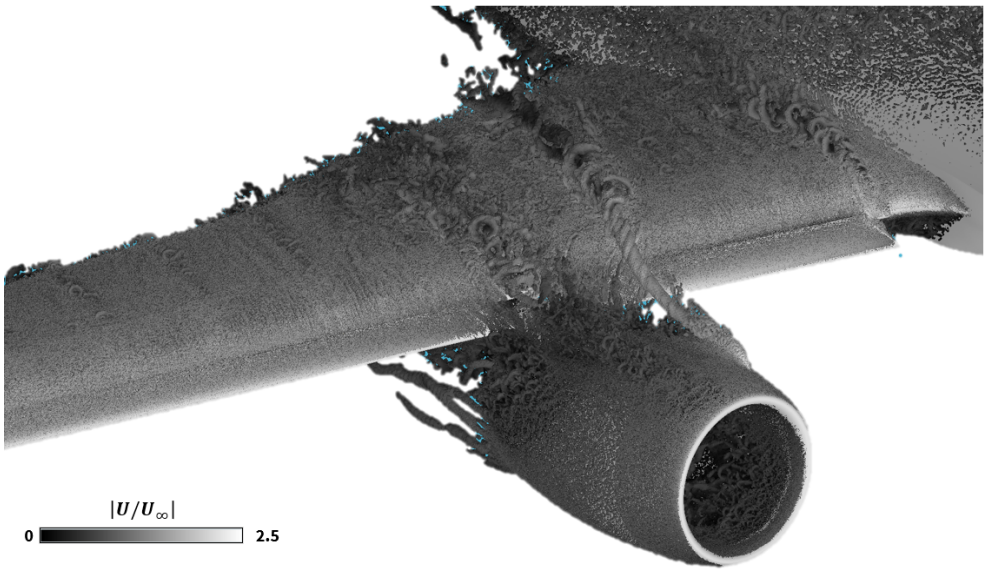

Figure 3.22: Visualizations of iso-surfaces of the Q-criterion of the flow around the CRM-HL at an angle of attack of 17.05°.

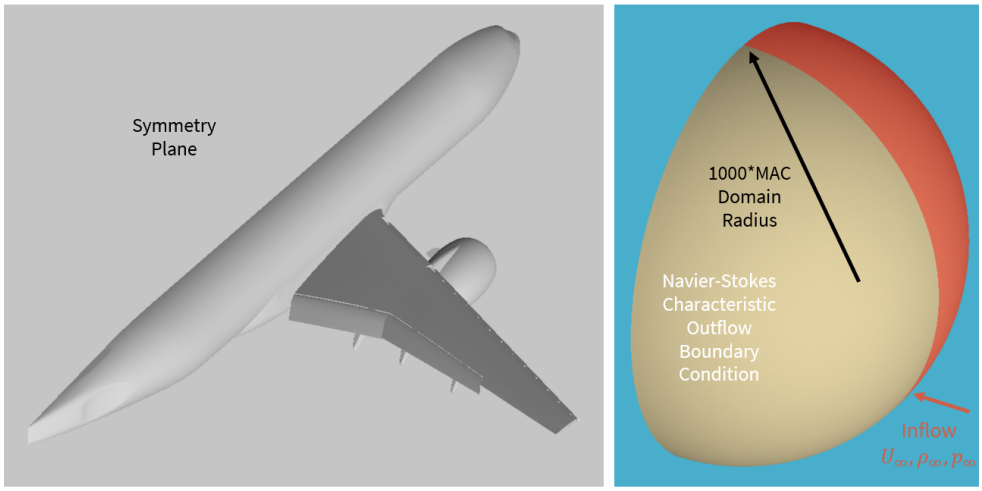

Figure 3.23: Schematic showing the computational geometry for the CRM-HL free air configuration. Also noted are inflow/outflow/symmetry boundary conditions and the radial extent of the hemispherical domain.

It is the long-standing practice of the computational aerodynamics community to compare corrected experimental data to half-span simulations with a large computational domain and a symmetry plane along the centerline [65, 62, 63, 67]. We adopt this approach for this study, though it will be revisited in the transonic CRM calculations. Figure 3.23 shows the domain that was considered

for this study. A large hemispherical enclosure of radius equal to 1000 MAC lengths is chosen in which a uniform freestream inlet is prescribed over the forward half of the domain, while a characteristic non-reflecting boundary condition is applied over the aft half. The center plane is modeled as a slip wall.

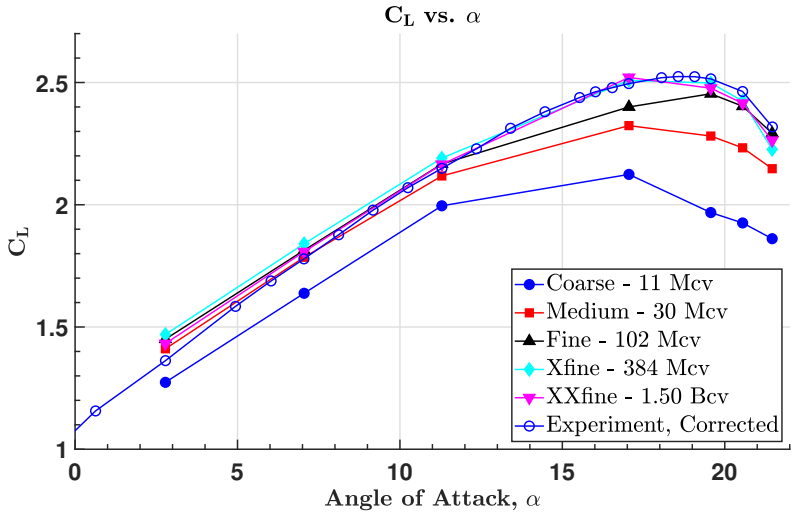

Figure 3.24: Corrected lift coefficient (C_L) plotted against aircraft angle of attack (α) in degrees for the CRM-HL free air configuration as compared with the predictions of charLES using five different grid resolution levels. Corrected experimental data [22] are plotted for comparison because the charLES results exclude the wind tunnel geometry.

Quantitative results are presented in Figures 3.24, 3.25, and 3.26. These include the average lift, drag, and pitching moment. The coarse and medium meshes are shown only on the plot of lift and omitted from there on out to simplify the plots. The average skin friction surface streamlines are compared with experimental oil flow visualizations in Figure 3.27. The raw lift force histories are shown in Figure 3.28. The calculations in free air were impulsively started, and averaging of the statistics was started after 20 convective flow passes based on the MAC had elapsed (except for the 21.46° case, where the startup transient lasted 30 flow passes). The time horizon of the simulation depended on the angle of attack, with the high angles generally being run for longer because increased unsteadiness associated with flow separation appeared at the higher angles. The calculations were considered converged when the running average of the lift coefficient changed by less than 0.01 over the course of the last 10 flow passes. This convergence metric identifies when the primary quantity of interest (lift) was changing by significantly less than the accuracy of significance (reported as 0.03 by [17]) over the course of a sufficiently long temporal integration window (10 flow passes). For the 21.46° case, an abrupt loss of lift was observed and was associated with the appearance of a large inboard stall. This is an unsteady phenomenon with long time scales, and the

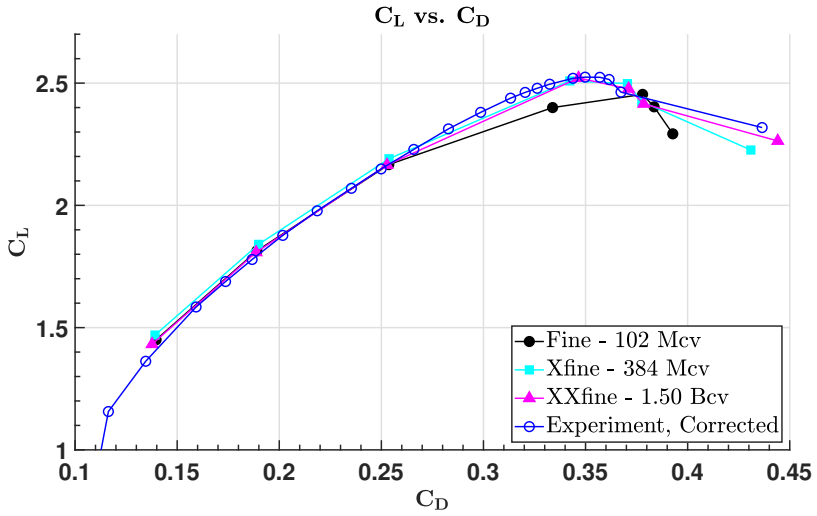

Figure 3.25: Corrected drag polar (C_L versus C_D) of the CRM-HL free air configuration as compared against the predictions of charLES using the finest three grid resolution levels considered.

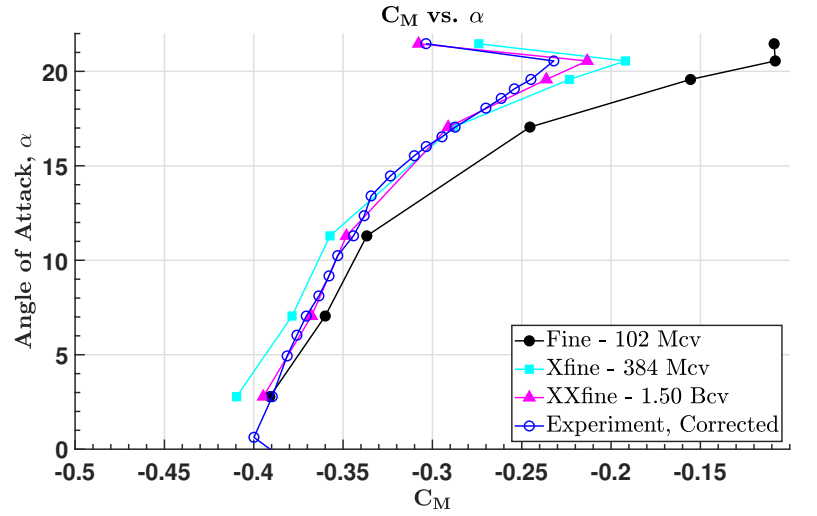

Figure 3.26: Aircraft angle of attack (α) in degrees plotted against corrected aircraft pitching moment coefficient (C_M) for the CRM-HL free air configuration as compared with the predictions of charLES using the finest three grid resolution levels considered.

quantities shown in this chapter for the 21.46° case are averaged over the entire simulation window shown in Figure 3.28, excluding a startup transient of 30 MAC-based flow passes.

The agreement of the LES solution with corrected experimental data for the free air configuration is excellent on the finest grid considered. Quantities such as the lift coefficient (Figure 3.24), the

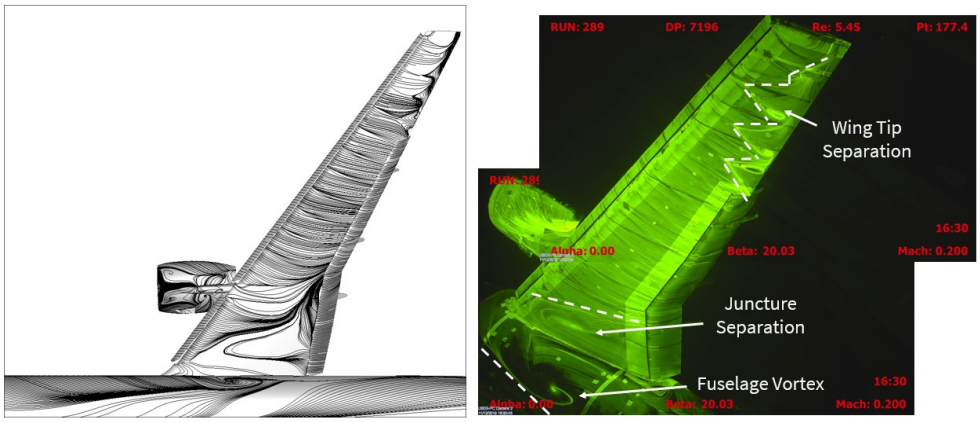

Figure 3.27: (a) Average wall shear stress streamlines on the suction side of the CRM-HL at an angle of attack after $C_{L,max}$, 21.46° from the free air charLES simulations compared against (b) oil flow visualizations obtained from the QinetiQ wind tunnel experiments.

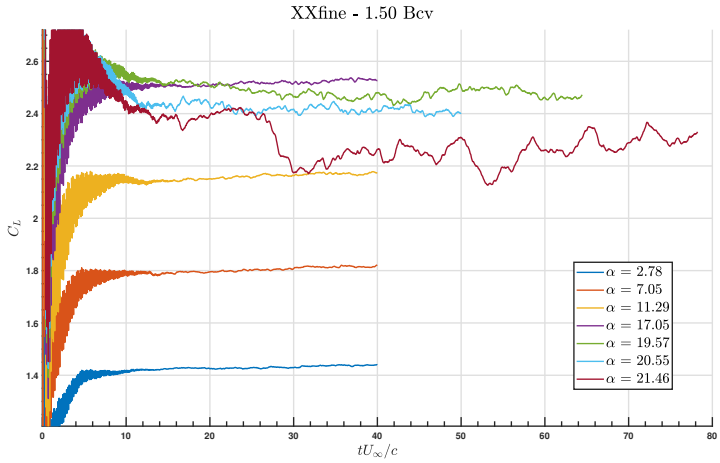

Figure 3.28: Lift force convergence time histories for the CRM-HL free air calculations plotted against the non-dimensional flow pass time. A startup transient of at least 20 flow passes is visible in these impulsively-started simulations and is excluded from the average statistics.

drag polar (Figure 3.25), and the pitching moment (Figure 3.26) all demonstrate excellent agreement with corrected experimental measurements. The accuracy of the integrated lift measurement is corroborated at the highest angle of attack in Figures 3.29 and 3.30, where the sectional pressure measurements on the Xfine and XXfine grids also reveal excellent agreement with experiment, showing that the prediction of integrated quantities is not subject to fortuitous error cancellation along the span. A promising result is the systematic improvement observed in the solution with grid

refinement, with the exception of some non-monotonic convergence observed in the linear regime of the pitching moment curve (the two lowest angles of attack in Figure (3.26) on the Xfine grid in which the moment becomes more nose down relative to the Fine and XXfine grids. This drop in pitching moment was associated with a non-physical collapse of the flap separation pattern and although the pitching moment trended towards the test data on the XXfine grid, promising novel subgrid and wall modeling approaches for LES may be the key to rectifying this undesirable grid convergence behavior in the future [1].

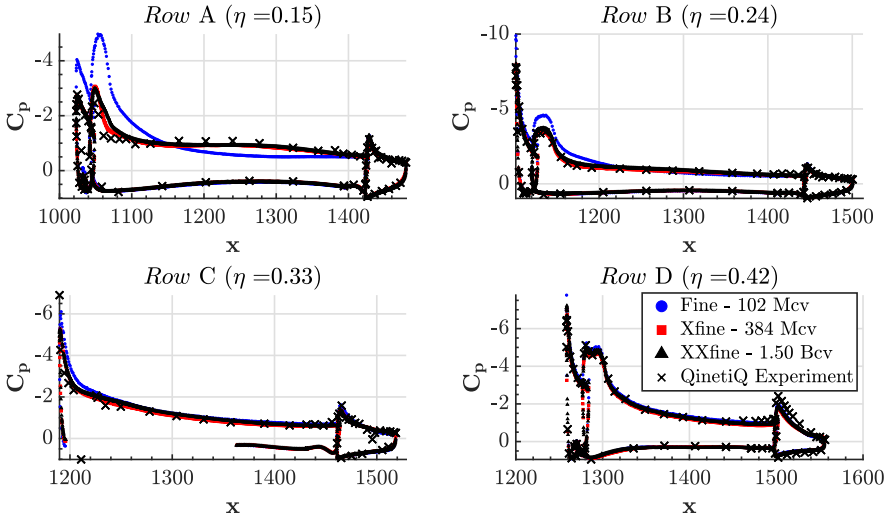

Figure 3.29: Pressure coefficient (C_p) plotted against chordwise location (x) for the CRM-HL configuration, showing a comparison of the predictions of charLES using three grid resolution levels, compared with experimental measurements for the four inboard pressure belts at a corrected angle of attack of 21.46°.

We emphasize that the augmentation in lift (Figure 3.24) and drop in pitching moment (Figure 3.26) observed on the Xfine grid have to do with the non-physical collapse of the flap separation bubble predicted by WMLES at these flow conditions. Because the flap lies aft of the point about which the pitching moment is computed and because the sign convention is that negative pitching moment is in the nose-down direction, the spuriously attached flap flow contributes increased nose-down pitching moment relative to the stalled flap case. Experimental oil flow observations confirm that at low angle of attack, this flap flow should indeed be separated [22]. This phenomenon has now been observed across various high lift configurations, including the JAXA Standard Model, and is discussed in greater detail by [29].

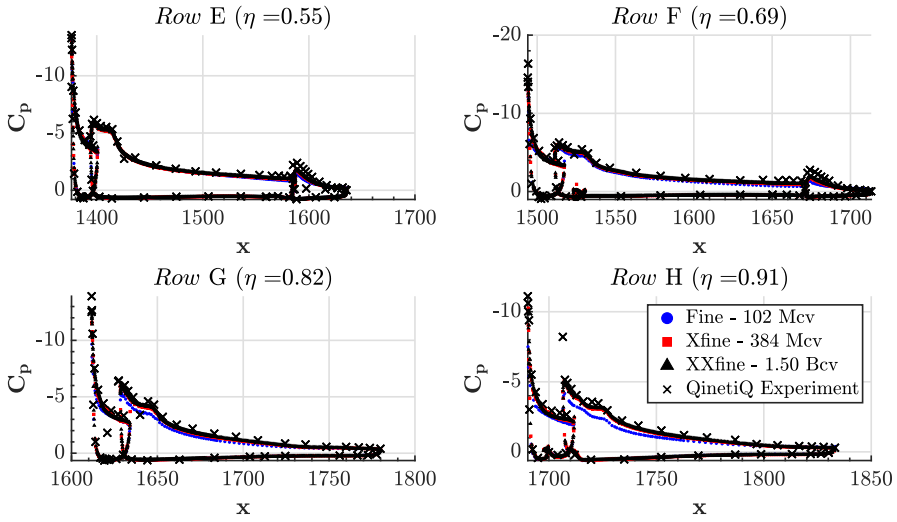

Figure 3.30: Pressure coefficient (C_p) plotted against chordwise location (x) for the CRM-HL configuration, showing a comparison of the predictions of charLES using three grid resolution levels, compared with experimental measurements for the four outboard pressure belts at a corrected angle of attack of 21.46°.

Application of a Non-Boussinesq Subgrid-Scale Model

This subsection describes the application of the novel non-Boussinesq Dynamic Tensor Coefficient Smagorinsky model (DTCSM) along with a suitable sensor-based wall model (SWM) [2, 1] to the CRM-HL, specifically aimed at addressing the lift over-prediction at low angles of attack. This modeling approach has been shown to rectify non-monotonic grid convergence of the separation bubble over a Guassian bump [79], a phenomenon closely tied to that of the separation collapse over the trailing edge flap observed in high-lift flows. The numerical experiment carried out in this section was performed on the Xfine mesh numbering 384 Mcv, because the collapse of the flap separation bubble is most pronounced at this grid resolution (see Figure 3.24 where the cyan Xfine curve lies clearly above the rest of the meshes in the low angle of attack regime of the lift curve). Figure 3.31 shows that the effect of the DTCSM + SWM is to decrease the lift, which is a trend in the right direction, but that the effect is too drastic. Further diagnostics were evaluated to identify the culprit of this aggressive drop in lift, including the sectional pressure coefficients (Figure 3.32) and the surface streamline patterns (Figure 3.33). It is clear from both of these metrics that the DTCSM + SWM predicts too much flap separation, in contrast to the DSM + EQWM approach which predicted too little. Figure 3.32 shows that the strength of the trailing edge flap suction peak is too weak relative to the experimental observations while Figure 3.33 shows wider extents of area on the flap where there are no attached streamlines from the DTCSM + SWM simulation compared

to the DSM + EQWM case. Since the streamline traces are initialized at the leading edges of the flaps, regions in which there is no attached surface streamline can be deemed separated. Overall, we conclude from this experiment that the full benefit achieved on the canonical Gaussian speed bump case using the DTCSM + SWM approach was not realized on the full aircraft configuration, though in general the predicted trend towards more flap separation was correct, the size of the separation predicted was too large and the issue of trailing edge flap separation bubble non-monotonic grid convergence persists, even when leveraging recent advances in physics-based subgrid-scale modeling.

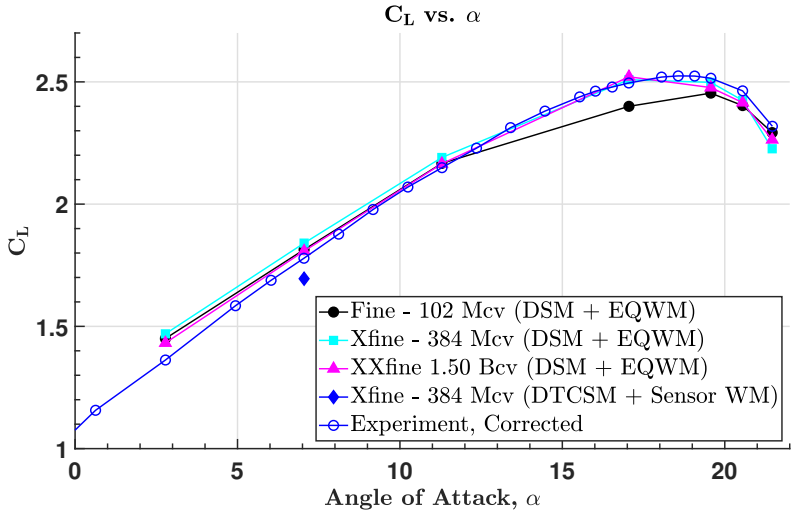

Figure 3.31: Corrected lift coefficient (C_L) plotted against aircraft angle of attack (α) in degrees for the CRM-HL free air configuration as compared with the predictions of charLES using the finest three grid resolutions considered with the standard DSM + EQWM modeling approach along with the novel DTCSM + Sensor Wall Model approach.

Note on computational cost

The CRM-HL calculations described were run primarily on the Oak Ridge Summit cluster. The GPU-accelerated charLES solver developed at Cascade Technologies was leveraged to enable turnaround times and grid resolution levels not attainable on typical CPU-based computing platforms. Table 3.2 shows the cost of each case of the free air runs. Each run of the Fine-XXfine meshes was conducted on 100 NVIDIA V100 GPU nodes, each of which contains 6 GPU cards. At the Xfine resolution, whose quantitive results agreed quite well with the 1.5 Bcv mesh, statistically converged simulations were achievable within a few hours on 100 GPU nodes of the Summit cluster, which accounts for about 2% of the machine. Further speedup in the Xfine-XXfine simulations can be achieved by increasing the processor count beyond 600 GPU's (since the solver continues to scale well at least to a loading of 1 Mcv/GPU, the 1.5 Bcv mesh solution in particular could be sped up significantly);

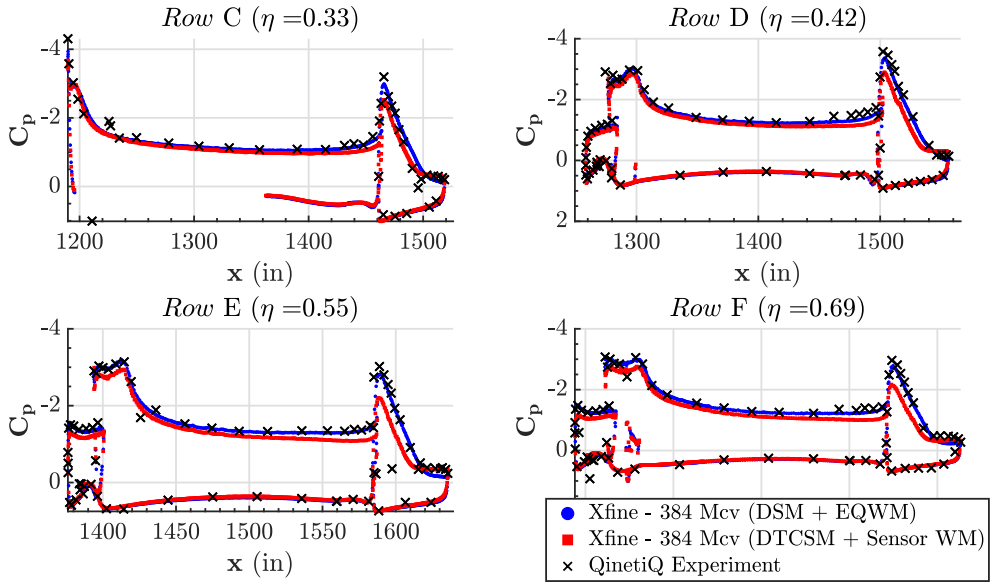

Figure 3.32: Pressure coefficient (C_p) plotted against chordwise location (x) for the CRM-HL configuration, showing a comparison of the predictions of charLES with the standard DSM + EQWM modeling approach along with the novel DTCSM + Sensor Wall Model approach, compared with experimental measurements for four mid-span pressure belts at a corrected angle of attack of 7.05°.

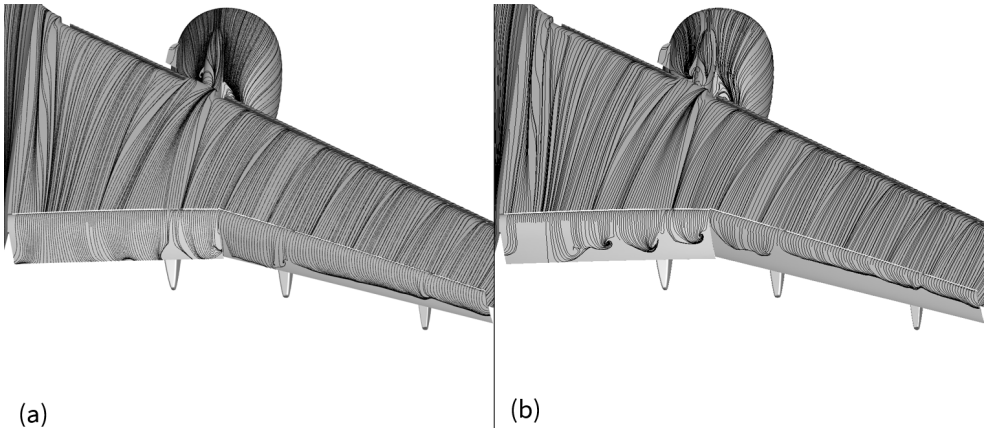

Figure 3.33: Average wall shear stress streamlines on the suction side of the CRM-HL at an angle of attack of 7.05° from Xfine mesh simulations using (a) traditional DSM + EQWM turbulence modeling approach and (b) using a non-Boussinesq dynamic tensorial coefficient SGS model with a sensor-based wall model.

however, practical queuing considerations on the Oak Ridge Summit cluster led to this choice of maximum processor count. Time to solution is defined as 30 convective flow passes. This is chosen as the minimum viable solution time in which the first 20 flow passes are cleared due to startup transient effects and statistics are collected over the final 10 flow passes of the simulation. These results demonstrate that affordable and accurate high-lift aircraft simulations are achievable with GPU-acceleration on modern computing systems.

Grid	**Cell Count (M)**	**Wall Time to Solution**	**Number GPUs**
Coarse	11	31 min	120
Medium	30	43 min	300
Fine	102	1 hr 10 min	600
Xfine	384	5 hr 13 min	600
XXfine	1500	39 hr 2 min	600

Table 3.2: Computational cost summary for a grid sequence using the charLES solver for the CRM-HL configuration at angle of attack of 7.05°.

3.3 Summary

In this chapter, wall modeled LES calculations were conducted of both the JAXA Standard Model and the CRM-HL, which are representative geometries of commercial high-lift aircraft. These simulations achieved two important objectives. First, the LES calculations demonstrated 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 within the bounds of experimental repeatability at an angle of attack within a degree of the experimental measurements for the JSM case. Surface pressure and near wall flow visualizations were compared against experimental measurements, which corroborated that the LES calculations properly replicate appropriate flow structures for both cases. These include the prediction of separations at the outboard slat brackets near $C_{L,max}$ and the onset of an inboard separation pattern at high angle of attack for the CRM-HL case. 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 modeled) LES would be prohibitively expensive and intractable.

Furthermore, the solutions for both configurations generally exhibited improving agreement with grid refinement. A grid convergence study in the CRM-HL case was considered at 7 angles of attack

and at 5 grid resolution levels each, with the finest mesh considered being nearly collapsed with the test data. For the CRM-HL near stall, the solutions are in excellent agreement with experiment as measured by force/moment, sectional pressure, and surface streamlines. The presence of a large inboard separation pattern was observed in the simulations, suggesting that some key aspects of the separation mechanism are able to be predicted even in the absence of wind tunnel installation effects, though this conclusion will be expanded upon in the subsequent chapter. The most significant grid sensitivities between the grids of increasing refinement level were concentrated both towards the leading edges of the slat and of the main element towards the wingtip and at the inboard part of the wing where there is a large separation. This suggests that grid sensitivities in the present simulations are tied to both inviscid effects associated with fine-scale features in the geometry at the wingtip and with resolution of the flow physics of the juncture flow. Lift over-prediction at the low angles of attack was observed in both configurations and was attributed to a flap separation prediction that is too small relative to the experimental observations. A novel non-Boussinesq SGS modeling approach which has showed significant promise in canonical flow environments was explored in an attempt to rectify this issue, however the effect of this approach was to instead predict over-separation of the flap flow. Overall, this simulation campaign represents a significant achievement in the use of high fidelity simulation approaches for practical aeronautical applications.

Chapter 4

Wind tunnel effects in large-eddy simulations of high-lift flows

4.1 Background

This chapter describes the study of wind tunnel installation effects in two high lift aircraft in landing configuration (the JAXA Standard Model and the High-Lift Common Research Model, as before), a complex external aerodynamic flow configuration with deployed slats, flaps, a flow-through nacelle, and the associated brackets/fairings on the high-lift devices. the bulk Mach number for both cases is low (0.2) and the MAC-based Reynolds number is typical of a wind tunnel experiment (1.93×10^6 - 5.49×10^6). The calculations of the CRM-HL case are carried out in a setting and which includes the QinetiQ wind tunnel facility at seven angles of attack at five grid resolution levels, ranging from $\approx 10 - 1500$ million control volumes (Mcv), while the JAXA calculations include the JAXA LWT1 wind tunnel facility at six angles of attack at three grid resolution levels, ranging from $10 - 185$ Mcv. In both cases the grid densities are designed to be identical to their corresponding Coarse/Medium/Fine free air counterparts discussed in the previous chapter. The grid counts differ because of the additional resolution required to captured the tunnel sidewall boundary layers in the in-tunnel cases. A key flow feature observed in the experiments near stall was the appearance of a large inboard separation. The LES simulations performed in the wind tunnel predict a stall mechanism featuring a large inboard juncture stall and an associated nose-down break in the pitching moment curve for both configurations. In the case of the JSM, this separation was absent from free air calculations of the same configuration. The accuracy of the simulations is once again assessed via comparisons of integrated forces/moments, surface pressures, and surface skin friction visualizations. Simulations that include the wind tunnel geometry facilitate more direct comparisons against the raw experimental measurements because experimental corrections meant to account for facility-specific effects to the wind tunnel force/moment data need not be taken into account and

therefore uncorrected data are used as a reference in this chapter. The presence of the wind tunnel installation does not change the conclusion drawn in the previous chapter that GPU-acceleration of the charLES solver results in tractable turnaround times that make LES a useful tool in the industry design cycle. We once again find that high-lift aircraft flows with solutions that are nearly converged in lift/drag/moment are attainable within several hours of walltime on 600 GPU cores, the only difference relative to the free air simulations being the need to perform grid sequencing to clear the long startup transient associated with the trapping of acoustic waves in the confined wind tunnel structure. This transient is efficiently cleared on coarse meshes and therefore does not contribute meaningfully to the total cost of the fine mesh simulations. This chapter benefits from the learnings of earlier research at the Stanford Center for Turbulence Research [45, 26, 27, 28, 29], where wall-modeled large-eddy simulation (WMLES) was used to simulate a realistic aircraft in landing configuration across the lift curve.

4.2 Wind tunnel calibration procedure

4.2.1 Achievement of test section target Mach number

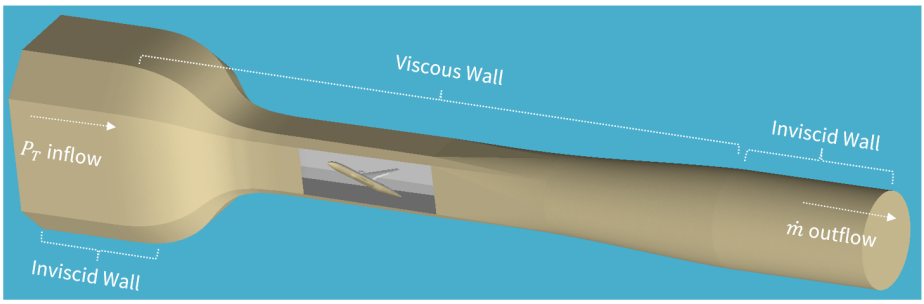

Figure 4.1: Schematic showing the computational geometry for the CRM-HL in wind tunnel configuration. Also noted is the treatment of the viscous/inviscid tunnel sidewalls and the inflow/outflow boundary conditions.

In this section, we discuss efforts at the simulation of the CRM-HL configuration in a setting in which the wind tunnel facility is reproduced in the computational model. The geometry of the QinetiQ 5m experimental facility in which the half-span model was tested was provided by the Fourth AIAA High-Lift Prediction Workshop committee. The variable cross-sectional area facility, including the inlet, contraction, test section, and diffuser are included. Figure 4.1 shows the geometry with the airplane mounted in the test section along with the boundary conditions prescribed. The inlet and outlet are artificially extended in the upstream and downstream directions to mitigate numerical contamination from the boundary conditions. These extension regions are treated inviscidly, as shown in Figure 4.1. Elsewhere, the tunnel sidewalls are treated as viscous (meaning that the

algebraic equilibrium wall model is active) and a boundary layer develops naturally. This approach led to a test section tunnel wall boundary layer that is $\approx 2\times$ thinner than measured in the experiment. The inability to precisely was a common theme amongst HLPW4 participants and is a current shortcoming of the in-tunnel simulations. Further efforts were made to address this discrepancy and are discussed in subsection 4.2.2.

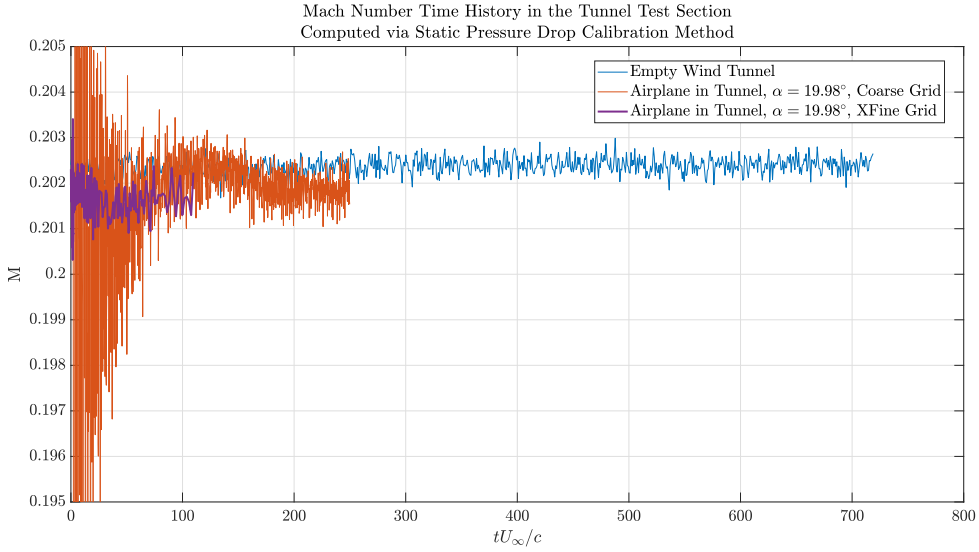

Figure 4.2: Achieved test section Mach number as calculated via the static pressure drop method of the experiment plotted against flow pass time.

For this configuration, the inflow boundary condition is prescribed as a total pressure and total temperature-based characteristic boundary condition. Total (stagnation) quantities are computed on the basis of the reference conditions in the test section, and losses in stagnation pressure and temperature between the inlet and test section are assumed to be small. This assumption is later assessed by means of evaluating the achieved test section quantities shown in Figure 4.2. At the outlet, we prescribe a mass flow rate-based Navier-Stokes characteristic boundary condition. The target mass flow rate is computed based on the test section cross-sectional area and reference density/velocity. The boundary condition modulates the back pressure until the target mass flow rate is achieved. A relaxation time scale based on the tunnel length and acoustic speed is set to minimize the effect of trapped acoustic waves, which a time history of instantaneous pressure slices through the wind tunnel showed to be absent in the simulations. In the experimental setup at the QinetiQ 5m wind-tunnel facility, reference conditions are controlled by measuring a static pressure drop upstream to downstream of the contraction and a control system modulates the fan speed accordingly. A calibration table then relates this pressure drop to an experimentally measured Mach number at

the test section center point. The approach to modeling the tunnel which uses the mass flow rate-based boundary condition eliminates the need to manually adjust the back pressure at every angle of attack to achieve the desired Mach number in the test section, which would vary based on angle of attack due to changes in the tunnel blockage. Verification that the test section Mach number is achieved to within ± 0.005 of the target value of 0.20 was successfully performed by monitoring the Mach number calculated by the pressure drop method throughout the computation, and is shown for a typical calculation in Figure 4.2. The figure shows that this wind tunnel startup method is robust with respect to blockage effects (empty vs. airplane in tunnel simulations) and grid resolution (coarse vs. fine simulations).

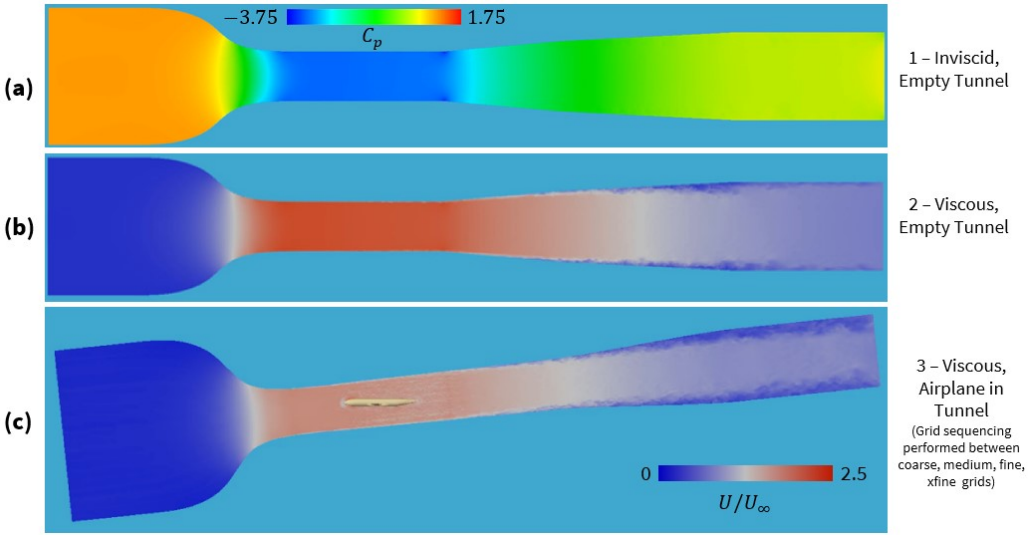

Figure 4.3: Side views of the wind tunnel showing the solution sequence for including variable cross-sectional area wind tunnels in the computational model of high-lift aircraft flows. The solution is initialized first with an inviscid empty tunnel calculation (a), after which a viscous empty tunnel calculation is conducted (b), and finally the aircraft configuration is included in the computational model (c). Verification that the appropriate reference conditions are achieved is done by computing test section Mach numbers using a static pressure drop method as was done in the experiment.

A further challenge with time-accurate in-tunnel CFD is the vastly different time scales required for the wind tunnel to achieve steady state compared to the airplane. While the airplane may require 10's of Mean Aerodynamic Chord flow passes to reach steady state, the wind tunnel may require 10's of domain length based flow passes to achieve its own steady state. To avoid the infeasible simulation length required by the full wind tunnel startup transient, a multi-stage startup method was utilized. First, a coarse inviscid case is used to stabilize the tunnel environment over a long time period. Next, resolution is added and tunnel wall viscosity is turned on. Finally, a grid staging approach is utilized to add the airplane to the stabilized domain. This allowed for solution time horizons comparable to free air cases of the same geometry (where the requisite solution length

was governed by the unsteadiness in the separated flowfield around the aircraft rather than the relaxation of acoustic waves in a confined wind tunnel environment). This sequencing procedure is shown graphically in Figure 4.3.

4.2.2 Attempts to match tunnel sidewall boundary layer profile

As mentioned in the previous subsection, an outstanding issue for in-tunnel simulations of high-lift aircraft remains the characterization of the wind tunnel sidewall boundary layers. For half-span models such as the one simulated here, this feature of the simulations is likely to play a more important role than for full-span sting or strut-mounted models because when the aircraft is installed on the tunnel floor it interacts directly with the sidewall boundary layer. In these situations, the aircraft is often installed with an offset/peniche to mitigate the influence of this complex interaction between the test apparatus and the boundary layer [84, 22], though it is not clear how effective this approach is. For this reason, wind tunnel simulations which include the offset/peniche and simulate the full tunnel facility, including the viscous tunnel sidewalls have been conducted. Initially, the simulations were conducted with smooth tunnel sidewalls that matched the available tunnel geometry specifications. The tunnel sidewall grid was sized so that approximately 10 points fit across its boundary layer at the experimental boundary layer probe location (which lies at the front end of the test section, about half of a fuselage length upstream of the fuselage nose).

As visible in Figure 4.5 from the "No Roughness" curve, this approach failed to appropriately capture the experimentally-measured boundary layer thickness and profile as viewed through the lens of the stagnation pressure deficit in the boundary layer. The boundary layer thickness associated with this smooth-wall LES simulation is denoted from here on out as δ_{LES}. Upon consultation with engineers involved in the experimental campaign and others performing experiments of the same configuration, none of which were able to match the quoted boundary layer profile from the experiment [64], the theory was developed that some realistic element of the geometry that is not included in the current computational model of the tunnel may be the culprit of the large differences observed in the profiles (e.g. an upstream protuberance, step, etc.). Because the characterization of the tunnel geometry is inadequate to directly represent these detailed features in the CAD model of the simulations, an artificial roughness strip was installed upstream of the test section to thicken the boundary layer.

This strip is shown in Figure 4.5; it is a wide region over which the tunnel floor in the simulation is perturbed to some maximum target height, varying from $1 - 4\delta_{LES}$. The height was parametrically varied because the exact nature of the upstream tunnel protuberance that caused the significant thickening of the boundary layer relative to the smooth wall simulations was not known, so a range of heights was used to assess the viability of this approach. Figure 4.6 shows a cross-section view through the grid at the location of the roughness strip. Because even the smooth wall simulation has 10 points per boundary layer thickness at this location, the minimum resolution per roughness

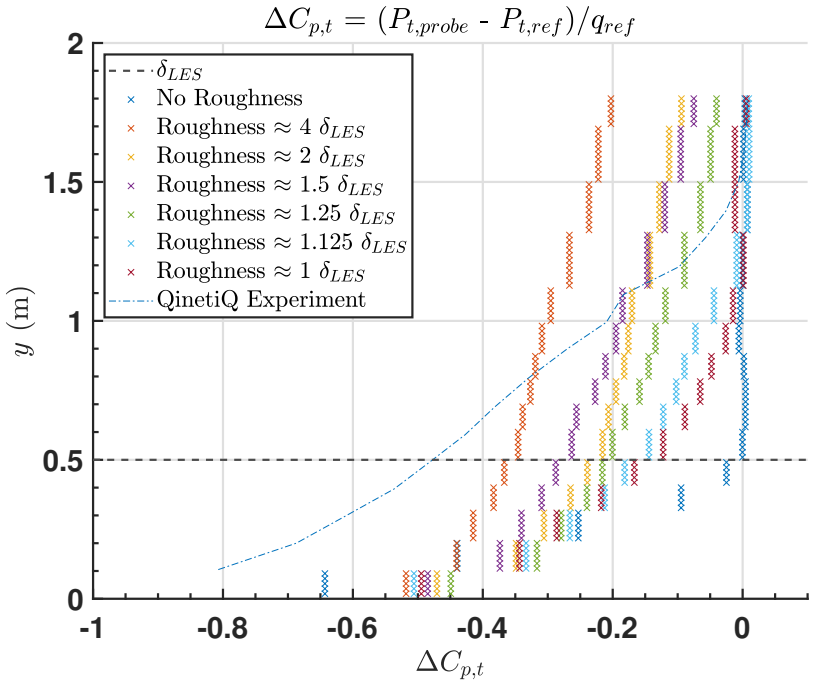

Figure 4.4: Boundary layer profiles plotted by means of stagnation pressure coefficient deficit from LES simulations using varying upstream roughness heights compared against experimental measurements [22].

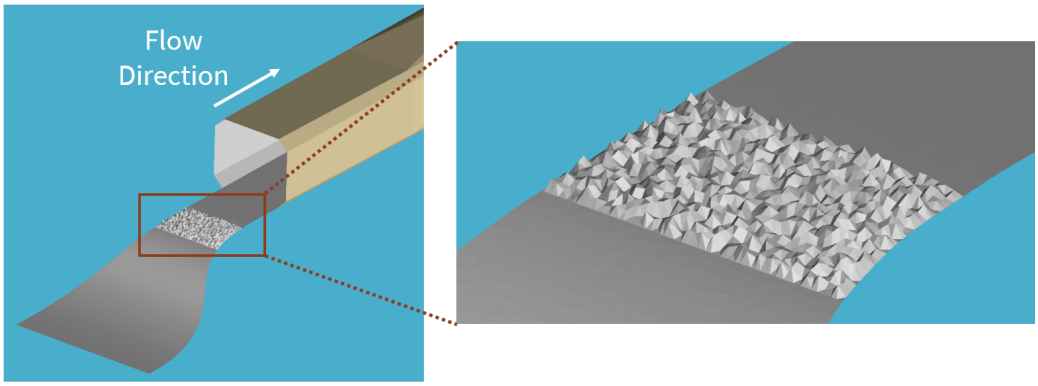

Figure 4.5: Geometry of the artificial roughness region. The strip of roughness is installed downstream of the contraction section and well upstream of the experimental boundary layer probe location, allowing the perturbed flow to develop naturally ahead of the probe location.

height for all of the simulations is 10 points (the smallest roughness height considered is of equal size to the local smooth wall boundary layer thickness). Unfortunately, figure 4.4 shows that getting

both the thickness and shape of the boundary layer profile correct was not achievable with this approach. An appropriate thickness alone was possible to achieve and occurs somewhere between $1.125 - 1.250\delta_{LES}$. However, because the roughness strip approach was not able to robustly account for discrepancies in the boundary layer profile between the experiment and clean wall simulations, it was not used in the wind tunnel simulations described in the following sections. The clean wall geometry remains the best known representation of the wind tunnel facility and the characterization of the oncoming boundary layer ahead of the aircraft model remains a known challenge associated with the calculations, as it was for all participants who attempted in the in-tunnel simulations at HLPW4 [64]. Boundary layer measurements at more locations ahead of the aircraft to assess its development and more detailed representation of the upstream tunnel geometry are needed to further diagnose this problem.

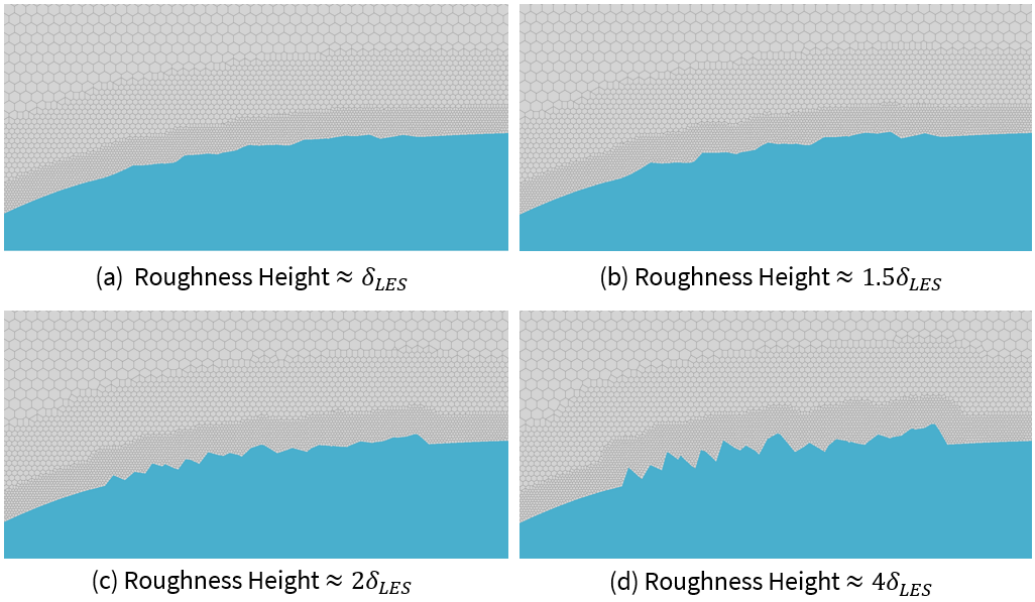

Figure 4.6: Four views of the grid sliced through the artificial roughness element. The roughness size varies from (a) one smooth wall boundary layer thickness (δ_{LES}) up to (d) four thickness heights.

4.3 Results

The following subsections will describe the quantitative impact associated with inclusion of the wind tunnel facility in the high-lift simulations of both the High-Lift Common Research Model and the JAXA Standard Model.

4.3.1 High-Lift Common Research Model

This subsection describes the results for the simulations carried out for the CRM-HL configuration in the wind tunnel. Quantitative results for the in-tunnel calculations are presented in Figures 4.7, 4.8, 4.9, and 4.10. These include the average lift, drag, pitching moment, and inboard/outboard pressures at the maximum uncorrected angle of attack of 19.98°. This angle corresponds to the corrected value of 21.46° which was the focus of the pressure and surface streamline comparisons in the free air section. The average skin friction surface streamlines are compared with experimental oil flow visualizations in Figure 4.11 while the raw lift force time histories are shown in Figure 4.12. The calculations in the wind tunnel environment were not impulsively started. Instead, a grid sequencing approach was adopted in which the solution on coarser grids was interpolated onto finer grids and used as a seed solution. This approach differed from the free air case because the confined environment of the tunnel had a propensity to trap acoustic waves. The use of a grid sequencing approach allowed these waves to be flushed out on the coarse grid, and for this reason they do not impact the temporal integration window needed on fine grid solutions. In free air, this does not pose an issue because of the massive domain extent and the relatively short window over which the statistics are averaged relative to the acoustic timescale associated with a 1000*MAC farfield. As in the free air cases, averaging of statistics was started after 20 convective flow passes (based on the MAC once the startup trapped tunnel acoustics had been cleared) had elapsed for the in-tunnel runs to mitigate the impact of transients associated with the interpolation. The time horizon of the simulation again depended on the angle of attack, with the high angles generally being run for longer due to increased unsteadiness associated with flow separation appearing. Figure 4.13 shows two key physical features that appear in the wind tunnel simulations that are absent from the free air runs: the wind tunnel sidewall boundary layer (distinguishable by the non-zero values of near-wall vorticity in sub-figure (b)) and the horseshoe vortex (distinguishable by the increased activity observed in the instantaneous q-criterion near the fuselage nose in sub-figure (b)). These features, along with the tunnel blockage are important phenomena arising for wind tunnel experiments involving half-span sidewall-mounted aircraft models and are absent from free air calculations of the same configuration. An effort was not made to isolate the impact of each feature independently, but we note that the aggregate effect of the inclusion of these features in the simulations did improved the prediction of the inboard separation pattern post-stall relative to free air simulations as shown in Figure 4.10. The effect of the tunnel installation is expected to be largest at the highest angles of attack [84].

To evaluate the quantitative results of the LES simulations, we first compare them with experimental measurements of integrated forces/moments. The key feature in the solutions after stall is the inboard separation and associated pitch break in Figure 4.9. This pitch break, however, now occurs approximately 1 degree early relative to the uncorrected experimental data in contrast to the free air results which showed this pitch break at an angle of attack that agreed with the corrected free air data. This observation suggests that some aspect of the way in which the wind tunnel is

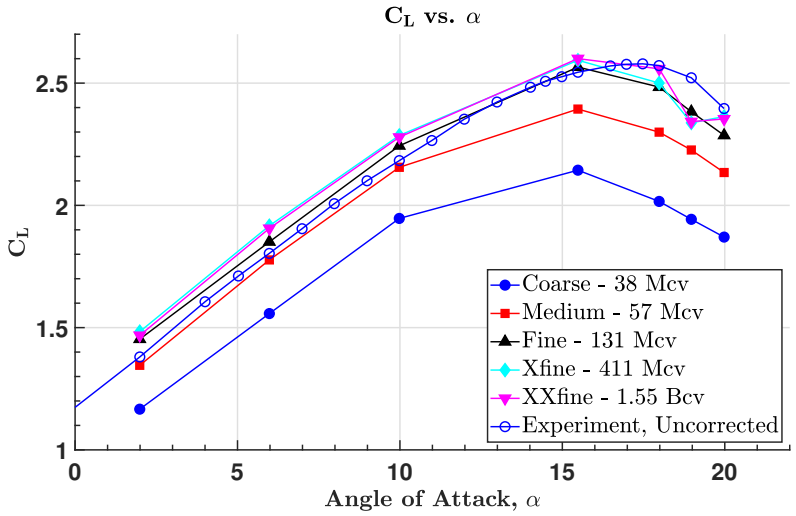

Figure 4.7: Uncorrected lift coefficient (C_L) plotted against aircraft angle of attack (α) in degrees for the CRM-HL in-tunnel configuration as compared with the predictions of LES using five different grid resolution levels. Uncorrected experimental data [22] are plotted for comparison because the LES results include the wind tunnel geometry.

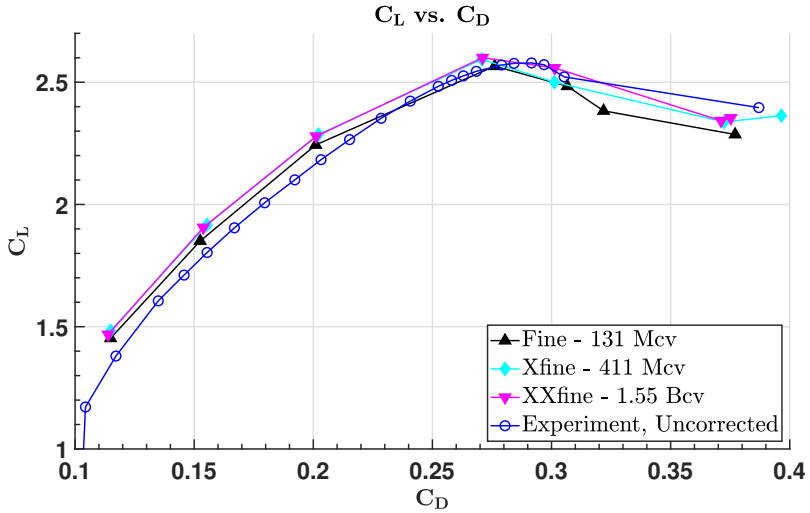

Figure 4.8: Uncorrected drag polar (C_L versus C_D) of the CRM-HL in tunnel configuration as compared against the predictions of LES using the three finest grid resolution levels.

modeled is not appropriate. For instance, it is known that the boundary layer thickness predicted by simulations is systematically low relative to the experimentally-measured boundary layer thickness [22, 64] for these test conditions and the simulations presented herein are no exception to this as

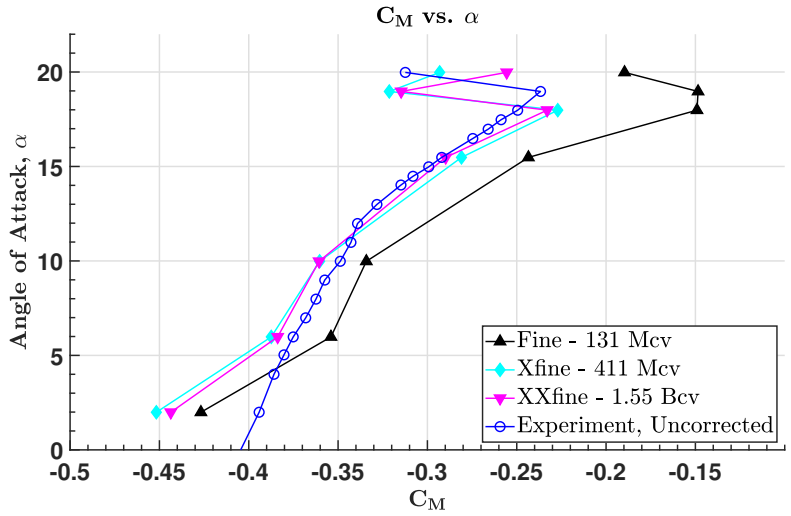

Figure 4.9: Aircraft angle of attack (α) in degrees plotted against the uncorrected aircraft pitching moment coefficient (C_M) for the CRM-HL in-tunnel configuration as compared with the predictions of LES using the three finest grid resolution levels.

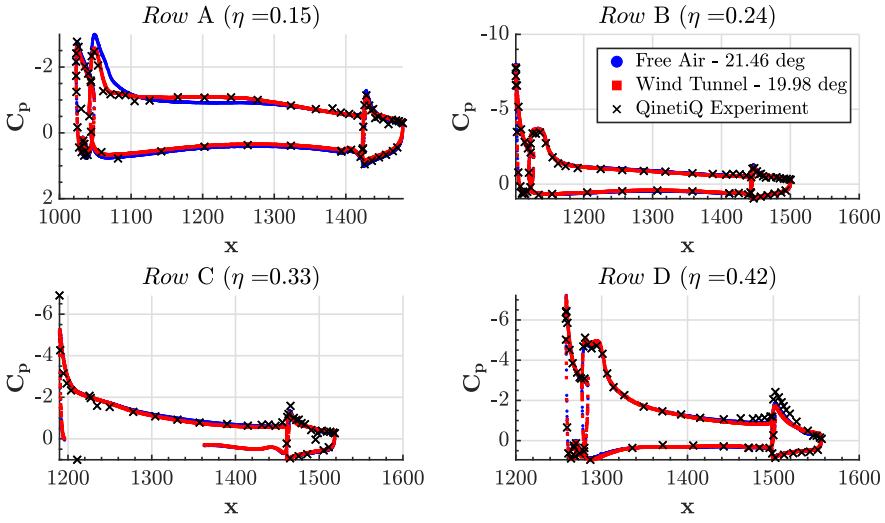

Figure 4.10: Pressure coefficient (C_p) plotted against chordwise location (x) for the CRM-HL configuration, showing a comparison of the predictions of LES when running in free air compared against the prediction when including the wind tunnel facility in the simulations on the XXfine grid. The four inboard pressure belts are shown.

was discussed in section 4.2.2. As discussed, efforts were made to artificially thicken the tunnel boundary layer via an artificial roughness element, but a boundary layer profile that matched both

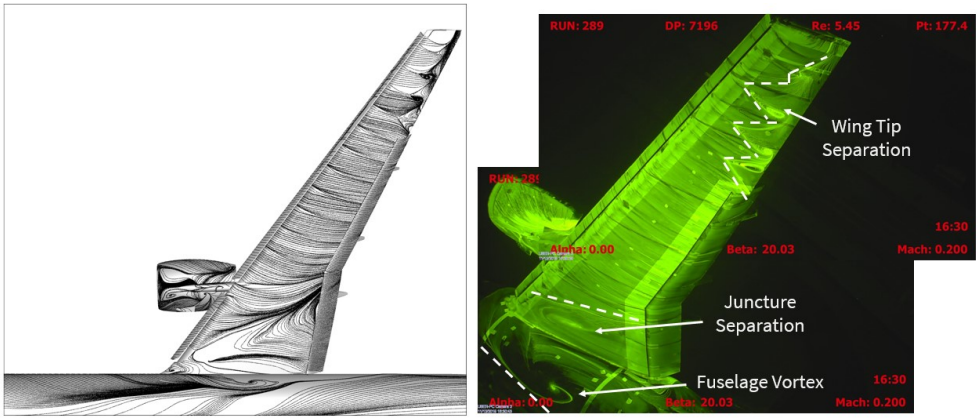

Figure 4.11: Average wall shear stress streamlines on the suction side of the CRM-HL at an angle of attack after $C_{L,max}$, 19.98° from the LES simulations that include the wind tunnel geometry (a) compared against oil flow visualizations obtained from the QinetiQ wind tunnel experiments (b).

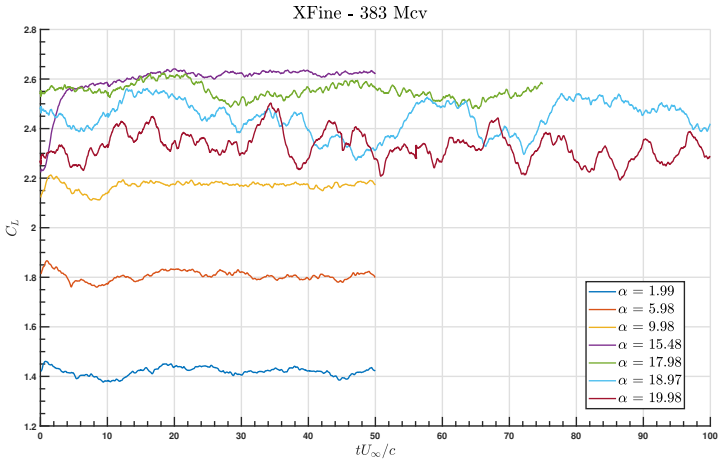

Figure 4.12: Lift force convergence time histories for the CRM-HL in tunnel calculations plotted against the non-dimensional flow pass time. A startup transient of 20 flow passes is excluded from the average statistics reported in this document.

the thickness and shape of the experimentally-provided profile was not achieved and this discrepancy remains an outstanding issue associated with the in-tunnel calculations that is not possible to rectify until more detailed experimental measurements of the upstream boundary layer growth are collected.

At the post-stall condition ($\alpha = 19.98°$), it does seem that there is some benefit to the tunnel installation in the prediction of the inboard separation patterns as the pressures at inboard station Row A in Figure 4.10 show an improvement in the characterization of loss in lift associated with the

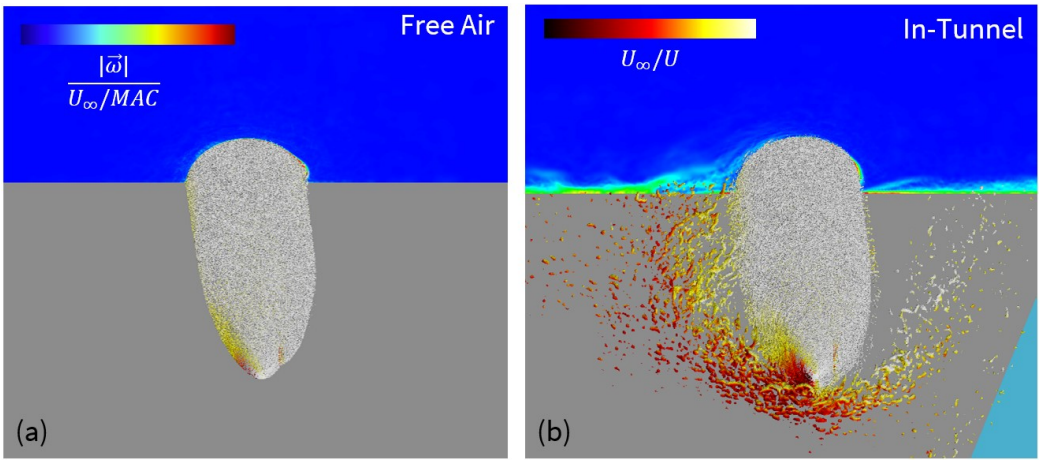

Figure 4.13: Isometric view of the forward fuselage on the CRM-HL at an angle of attack of 7.05° (sub-figure (a) in free air) and 5.98° (sub-figure (b) in the wind tunnel) with q-criterion colored by velocity magnitude and a planar slice colored by vorticity magnitude highlighting the tunnel boundary layer and fuselage horseshoe vortex .

inboard stall mechanism relative to the free air simulations. A qualitative view of this separation phenomenon is shown in the surface streamlines of Figure 4.11. Similarly to the free air simulations, we observe systematic improvement of the pressure prediction by LES with grid refinement, and in general, we also saw improvement of the solution with grid resolution at the level of the integrated forces/moments. The exception, again, is the lift over-prediction in the linear regime of the lift curve. The cause of this is again that the flap separation at low angle of attack is too small relative to the experimental observations. We note here that the grid resolutions used in the free air and in-tunnel simulations is identical in the vicinity of the aircraft. The grid counts differ between the two cases because the wind tunnel runs contain additional resolution of the tunnel sidewall boundary layers. To avoid confusion, we have quoted the grids as Coarse, Medium, Fine, Xfine, and XXfine for the in-tunnel runs. In the vicinity of the aircraft, these grids are identical to the Coarse 11 Mcv, Medium 30 Mcv, Fine 102 Mcv, Xfine 384 Mcv, and XXfine 1.50 Bcv grids quoted for the free air runs. As mentioned before, differences are expected between the free air and in-tunnel configurations due to the effects of tunnel blockage, tunnel sidewall boundary layers, as well as the formation of the horseshoe vortex at the fuselage nose. We find in this study that free air calculations are able to predict a pitching moment break and large inboard stall for this configuration, often thought to only appear in simulations that include the tunnel facility [36].

Nacelle Lip Separation

A key region of lingering discrepancy relative to the experimental oil flow visualizations was the nacelle lip. Specifically, the nacelle lip separates at the leading edge in the LES simulations prematurely relative to the experiment. For reasons of computational affordability, the nacelle leading edge trip dots from the experiment were not reproduced explicitly in the simulations (resolving this fine scale feature is prohibitively expensive for an solver which uses explicit time advancement scheme). Instead, a single test case at a post-stall angle of attack was chosen in which a synthetic suction boundary condition was applied at the location of the experimental trip dots in order to re-attach the nacelle lip boundary layer. This effort was successful when the suction strength was set to be at least 20% of the freestream velocity. This relatively strong forcing was applied on a narrow strip at the leading edge of the nacelle.

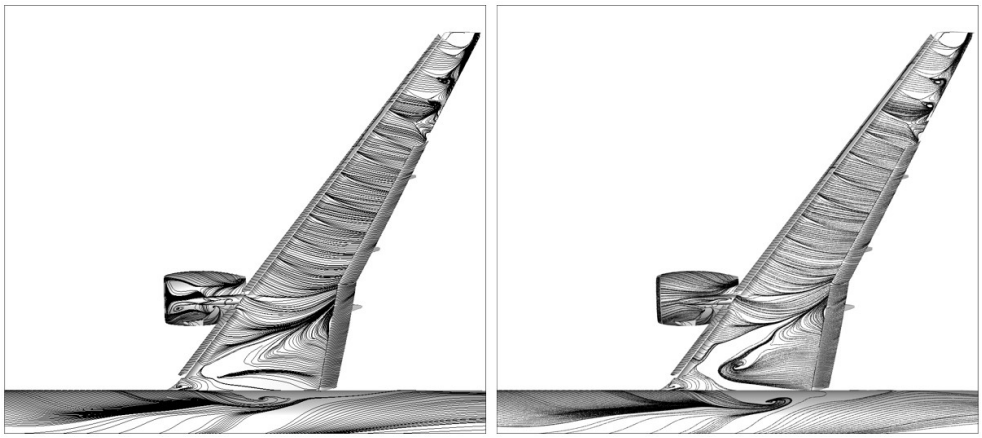

Figure 4.14: Average wall shear stress streamlines on the suction side of the CRM-HL at an angle of attack after $C_{L,max}$, 19.98° from LES simulations on the Xfine grid that include the wind tunnel geometry and include an unforced nacelle lip (a), which is stalled, and which include leading edge suction at the nacelle lip (b), which leads to an attached flow

The location and width of the suction region was chosen to match the location and diameter of the experimental trip dots. The impact of this forcing on the surface streamlines is shown in Figure 4.14. The visualizations show that the impact of the forcing is largely local, with other significant features of the flow (juncture separation, fuselage vortex, wingtip separations) mostly unaffected by the nacelle lip separation. Quantitatively, the impact of the nacelle forcing on the forces/moments was minimal and is not shown. When the nacelle lip separation was suppressed, the lift increased and drag decreased very slightly, while the pitching moment became more nose-down, which were subtle trends, though they did go in the direction of the experimental data. In general, we find that the influence of the nacelle lip separation is localized to the nacelle region and that while qualitative skin friction comparisons of the nacelle flowfield itself are significantly improved by the

leading edge suction, other features of the flow over the aircraft and the integrated forces/moments are not strongly impacted by this flow feature.

4.3.2 JAXA Standard Model

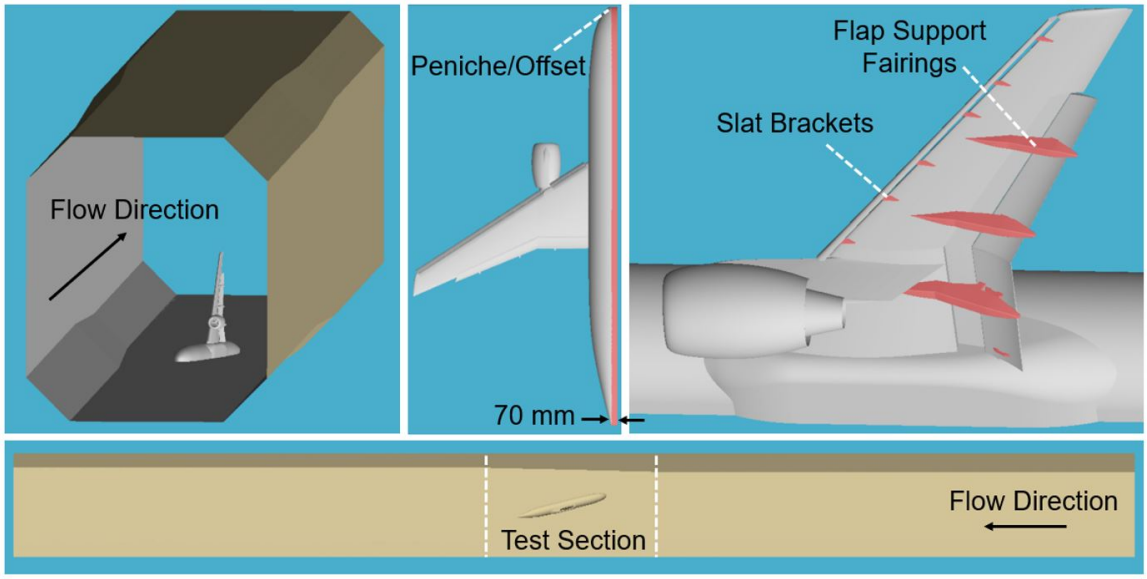

Figure 4.15: Schematic of the computational geometry for the JAXA Standard Model, nacelle on configuration, in the JAXA 6.5 m × 5.5 m Low-Speed Wind Tunnel.

The free air JAXA Standard Model simulations also showed encouraging comparisons against the experimental data, but there were two notable deficiencies in them: there was a systematic shift in the drag values obtained and a "kickback" of the pitching moment associated with the appearance of an inboard stall mechanism was 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 JAXA Low-Speed Wind Tunnel experimental facility. This configuration now includes a flow-through nacelle mounted on the underside of the wing and the test section walls (see Figure 4.15). 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 (contraction and diffuser geometries were not available for this tunnel, so only the test section is simulated). 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 modeled) 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 [36]. Figure 4.16 shows a the 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

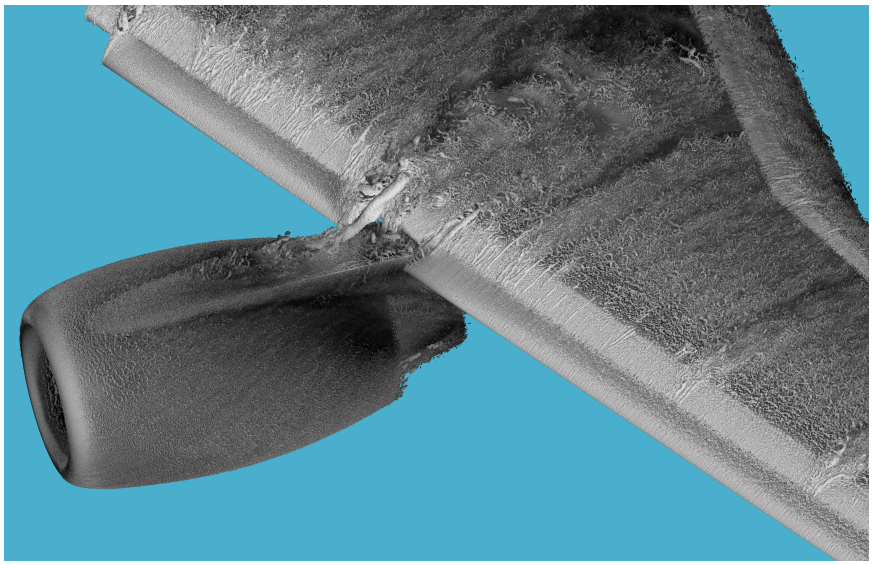

Figure 4.16: Iso-surfaces of Q-criterion colored by velocity magnitude showing the structure of the turbulent flowfield over the JAXA aircraft wing in landing configuration. Computed using LES at $\alpha = 18°$ including the engine nacelle and wind tunnel test section (not pictured) geometries.

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 4.18 shows the prediction of the lift, drag, and pitching moment in comparison to 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. Notably, 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 degree, as was the case for the CRM-HL simulations. Figure 4.17d shows surface streamlines of skin friction from LES compared against experimental oil flow visualizations at an angle of attack of 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 [63]. The inclusion of the wind tunnel effects appear to have remedied the highlighted shortcomings of the free air configuration predictions for the JAXA configuration.

Figure 4.17 shows some selected surface flow visualizations compared to 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

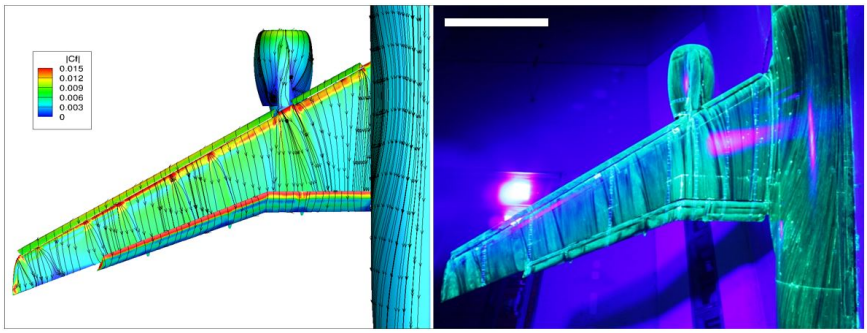

(a) $\alpha = 4°$

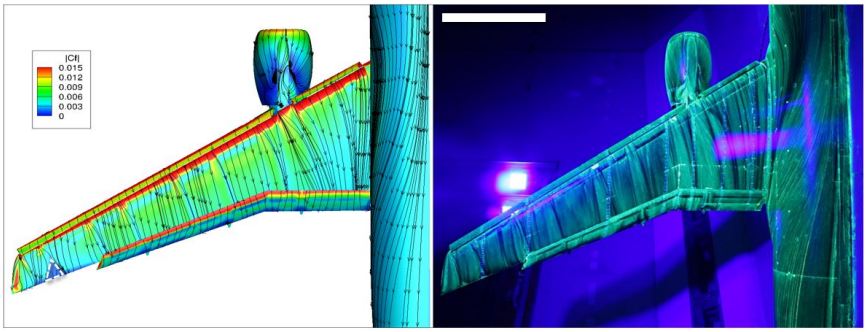

(b) $\alpha = 10°$

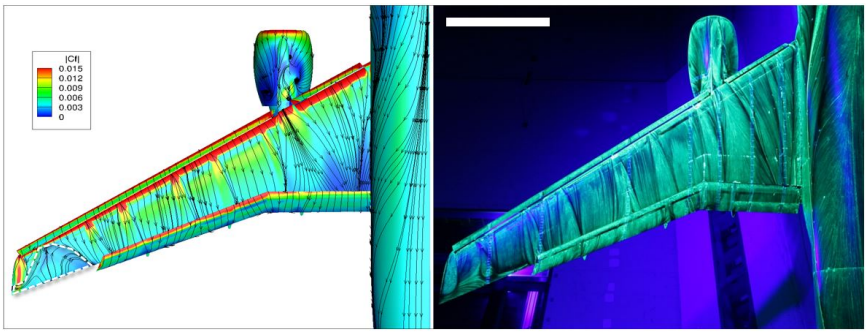

(c) $\alpha = 18°$

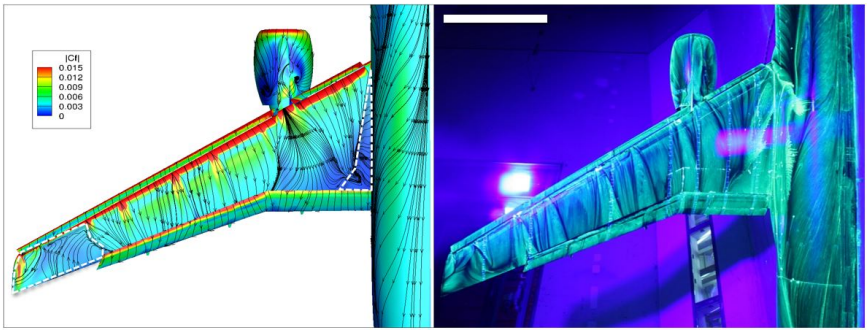

(d) $\alpha = 21°$

Figure 4.17: Average skin friction contours with skin friction streamlines (left column) from charLES fine grid simulations compared against experimental oil flow visualizations (right column), with white scale bar in the upper left corner of the experimental images corresponding to 1 mean aerodynamic chord length of ≈ 530 mm [83].

4.17a), leading to the aforementioned low angle of attack lift over-prediction. The impact of complex aircraft features 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 ($\alpha \approx 18^\circ$) show the separation of the boundary layer near the wing tip consistent with the experimental observations (Figure 4.17c). These visualizations provide some qualitative reinforcement that the flow mechanisms associated with the $C_{L,max}$ conditions are appropriately captured by the LES calculations.

4.4 Summary

In this chapter we have described simulations of the JSM and CRM-HL conducted using the compressible finite-volume solver charLES for cases which include the wind tunnel geometry. In general, the inclusion of wind tunnel facility effects tended to improve the characterization of the stall mechanism, which for both cases was a large separation originating at the wing/body juncture. The inclusion of wind tunnel facility effects adds several effects which the free air simulations miss, including the tunnel blockage/confinement and the tunnel sidewall boundary layers, which was at least minimally resolved in the simulations with 10 grid points. However, a known deficiency in the wind tunnel cases is the characterization of the sidewall boundary layer, which could have errors in its thickness of up to 3× based on available experimental data. Nevertheless, the integrated impact of these effects was shown to improve the prediction of the stall mechanism for both cases, though in the CRM-HL simulations the juncture stall and associated pitch break were also captured in the free air simulations. In the JSM case, these phenomena only appeared once the wind tunnel facility was modeled in the simulations. Additionally, an exploration of artificial numerical tripping to reproduce the effect of the nacelle lip trip dots was also undertaken and found to be successful at promoting an attached nacelle flow, but the impact on quantitative measures associated with the flow were minimal. Finally, we again emphasize as in the free air cases that the prediction of engineering quantities of interest (lift/drag/moment/sectional pressures) are nearly grid-insensitive between the finest two meshes considered and highly accurate LES solutions are achievable for the CRM-HL configuration within approximately 5 hours of computing wall time on 600 GPU cards using about 400 million control volumes.

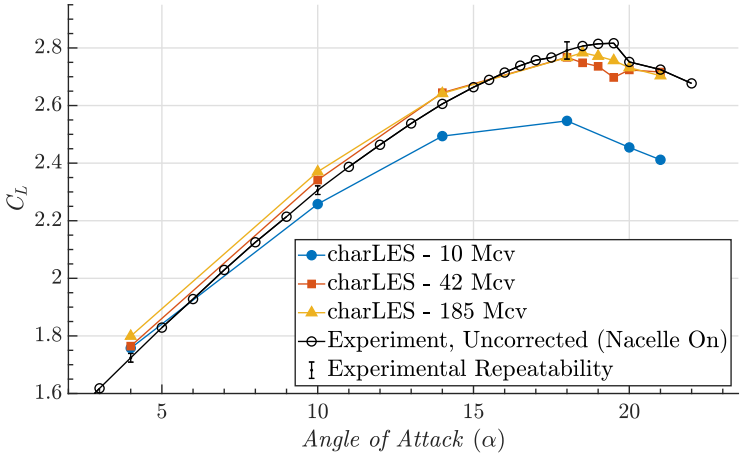

(a) C_L vs. α

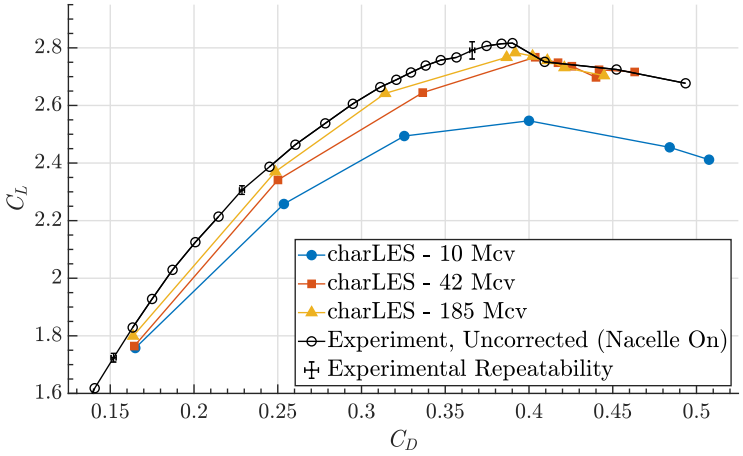

(b) Drag Polar

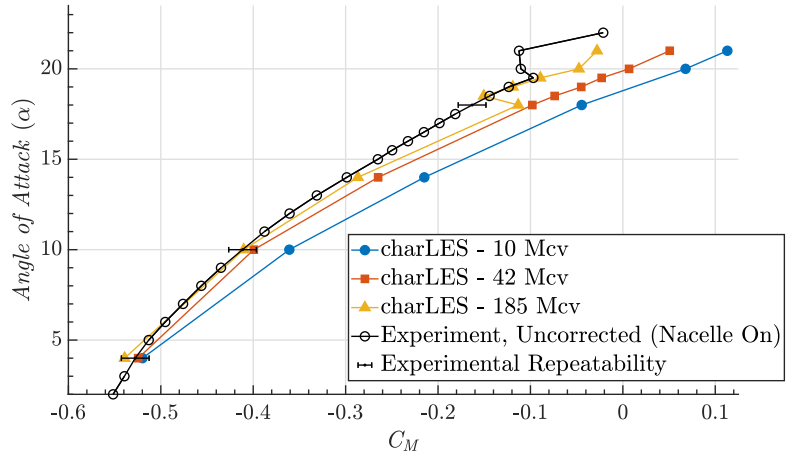

(c) α vs. C_M

Figure 4.18: (a) Lift curve, (b) drag polar, and (c) pitching moment coefficient plots for the JSM configuration including wind tunnel effects.

Chapter 5

Large-eddy simulation of the NASA transonic Common Research Model

5.1 Background

The accurate simulation of aircraft in transonic buffet conditions is a key pacing item towards Certification by Analysis (CbA) [53] whereby high-fidelity simulations would be leveraged in lieu or in addition to traditional wind tunnel and flight testing as part of the aircraft certification process. This chapter describes the use of Large Eddy Simulation (LES) in the context of transonic external aerodynamic flows over the transonic NASA Common Research Model (CRM), which has been the focus of several Drag Prediction Workshops (DPW). Sensitivities to laminar-to-turbulent transition, symmetry plane treatment, and grid topology are established and suggestions for best practices in these simulations are made. Simpler problems, such as a laminar Blasius boundary layer encountering a trip dot on a flat plate are studied in order to draw conclusions pertaining to forced transition behavior in a controlled environment. It is found that promoting transition to turbulence via an array of cylindrical trip dots, using the full span of the aircraft, including the sting mounting system in the simulations, and leveraging prismatic boundary layer grids all tend to improve the quality of the LES solutions. For the first time, novel non-Boussinesq subgrid-scale models and sensor-based wall models are applied to a complex external flow over an aircraft with promising results, particularly on coarse grids numbering about 100 million control volumes (100 Mcv). As has been observed in canonical smooth-body separation problems, non-monotonic grid convergence in wall-modeled LES continues to be an issue for problems of engineering interest, with finer grid solutions (up to 750 Mcv) performing poorer than their coarser counterparts. The details

of the boundary layer profiles both at the leading edge of the wing and within the shock-induced separation bubble are studied, with thicknesses and integral measures reported. Finally, an initial exploration of the buffet prediction capabilities of LES is made in the context of the NACA 0012 flow, with predictions showing reasonable agreement with available experimental data for alpha at initial buffet and shock oscillation frequency.

The chapter is organized as follows. In section 5.3.1 we describe validation efforts at modeling of experimental trip dots in LES calculations of transonic flows. In section 5.3.2 we describe the extension of these heuristics to a full transonic CRM aircraft configuration. Next, the sensitivities to the symmetry plane, grid topology, and subgrid/wall modeling are established in sections 5.4, 5.5, 5.6. Assessments of the boundary layer properties of the flow over a realistic aircraft configuration are made in section 5.7 while an extension of the LES methodology to the prediction of buffet onset is made in section 5.7.1. Finally, conclusions are drawn in section 5.8.

5.2 Results

5.2.1 Validation case: NASA transonic CRM

The reference conditions for the calculations presented herein are $Re = 5.0 \times 10^6$, M = 0.85 and $\alpha = 2.50° - 4.0°$ spaced at 0.25° increments. The Reynolds number is typical of wind tunnel conditions for transonic aircraft experiments. The reference data used is from the investigations of [60] at the NASA Langley National Transonic Facility (NTF), an image of which is shown in Figure 5.1. In all transonic CRM calculations conducted in this thesis, the aeroelastic deflection reported in [71] is used for the baseline geometries, resulting in new meshes and wing deformations for each angle of attack simulated, with up to $6.5''$ ($\approx$ 2.4% of MAC) of wingtip deflection (at full scale) having been observed between the lowest angle of attack, the $\alpha = 2.50°$ case, and the highest angle of attack, the $\alpha = 4.00°$ case. At full scale, the aircraft mean aerodynamic chord has a length of $275.8''$, while the wingspan takes on a value of $2313.5''$ ($\approx 8.4\times$ MAC). The aspect ratio of the wing is 9 and the leading edge is swept at an angle of 35° relative to the horizontal. The calculations described herein are run in a free air setting, because insufficient characterization of the NTF wind tunnel geometry is currently available in the open domain; however, the sting mounting system visible at the tail of the aircraft in Figure 5.1 is included in several of the simulations and improved shock locations have been observed upon inclusion of this feature, though it is not discussed in any more detail in this brief. Although this case has been a focus in the external aerodynamics community for many years, emerging paradigms such as LES have been used very sparsely to date on this configuration. The work of [46] was among the first of such efforts, which showed promising initial results on coarse grids at one angle of attack. More recent work by [25] studied the influence of the [76] subgrid-scale model coefficient choice on the shock location and found limited sensitivity to this choice. It was clear from these studies that further investigation into the treatment of turbulent

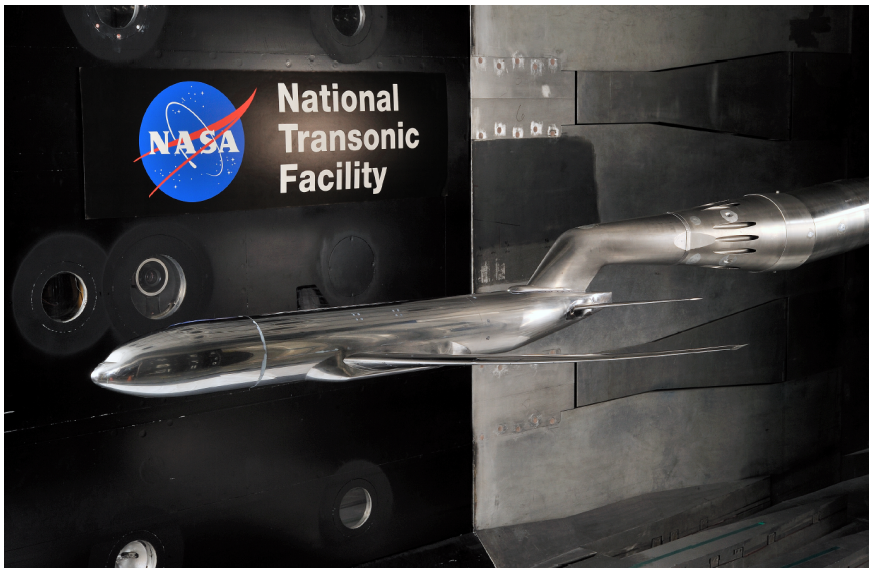

Figure 5.1: Image of the experimental apparatus from the NTF wind tunnel runs used as reference data in this brief. Reproduced from [60].

transition, grid resolution and subgrid-scale modeling approach are needed to build confidence in the predictive capabilities of LES in this flow regime. The aim of this chapter is to address these issues.

5.2.2 Description of grids used

This subsection is meant to provide clarity on the meshes used in the subsequent sections of the brief, explicating the geometric features represented on each mesh. Table 5.1 describes the details of the geometric features of the NASA CRM represented in each of the five meshes and the associated grid counts. The meshes used in each section are as follows: Section 5.3.2 uses meshes C and D, Section 5.4 uses meshes A and B, Section 5.5 uses meshes C and E, Section 5.6 uses meshes B and C, while Section 5.7 uses mesh E.

5.3 Importance of turbulent transition

5.3.1 Trip dot validation on a flat plate

It is fundamentally important to understand the role of trip dots in promoting transition in canonical problems such as the flow over a flat-plate boundary-layer. This analysis will anchor our investigations of the grid requirements for resolving trip dots in more complex flows such as full-scale aircraft. For this purpose, based on the experiment of [35], a single trip dot is placed in a laminar

Table 5.1: Details of the geometric features represented on each mesh and the associated grid counts. The meshes are given a label A-E.

	50 Mcv HCP Mesh (A)	**100 Mcv HCP Mesh (B)**	**750 Mcv HCP Mesh 1 (C)**	**750 Mcv HCP Mesh 2 (D)**	**270 Mcv Strand Mesh (E)**
Full/Half Span	Half	Full	Full	Full	Half
Trip Dots	Excluded	Excluded	Included	Excluded	Excluded
Sting Mount	Included	Included	Included	Included	Excluded
Aeroelastic Deflection	Included	Included	Included	Included	Included

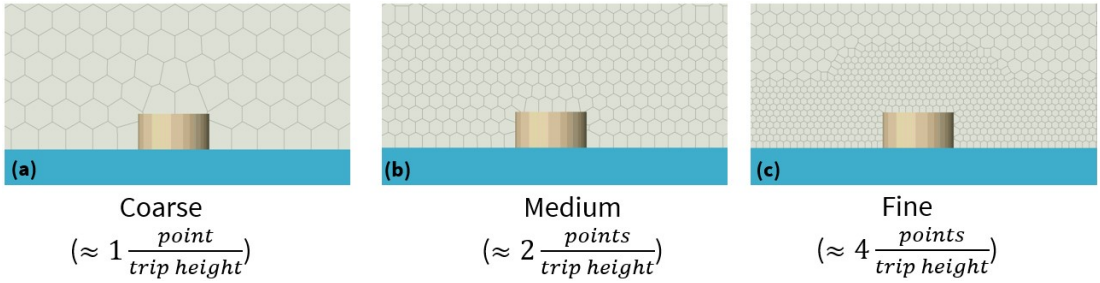

Figure 5.2: Spanwise slices of the three grids considered for the single trip dot calculations, ranging from a (a) coarse grid with just 1 point per trip dot height to (b) a medium grid with 2 points per trip height and (c) a fine grid with 4 points per trip dot height.

boundary-layer, and the wake spreading angle of the disturbed flow and the impact of the trip dot on skin friction (τ_w) are measured. [35] reported that the wake-spreading angle (measured as the angle between the symmetry plane and the edge of the wake) behind a single trip dot varied from $4.8° - 6.0°$ (depending on the height above the wall where the wake was measured) at a Reynolds number, Re_k, based on the trip dot height of 1000.

To replicate the laminar experimental inflow conditions, a Blasius boundary-layer solution is fed as the inflow to the LES simulation. The inlet boundary-layer thickness is adjusted such that the extrapolated boundary-layer thickness at the location of the trip exactly matches the experimental boundary-layer thickness, which was the same as the trip height (2 mm). These simulations are considered to be of an LES character even in the upstream laminar region because the DSM SGS model remains active everywhere in the domain and the proper laminar damping behavior, which is embedded into the DSM model [56], is relied upon in this limit. Three grid resolutions, as shown in Figure 5.2, were examined in this work. Note that the grids are isotropically refined by a factor of two to up to 4 points per trip dot height. The primary reason for not evaluating further grid sensitivities

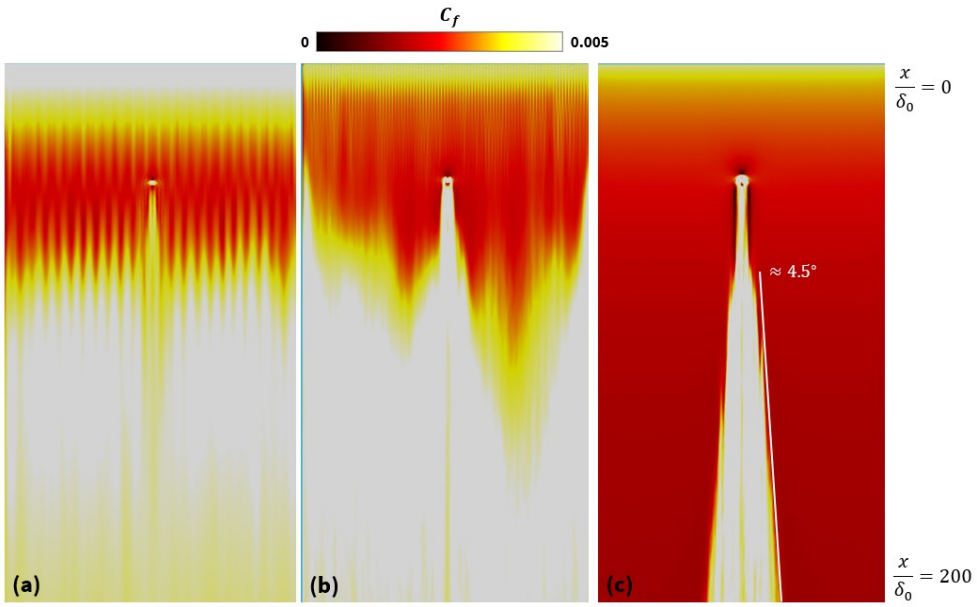

Figure 5.3: Surface contours of the average skin friction for the single trip dot flat-plate calculations ranging from (a) a coarse grid with just 1 point per trip dot height, (b) a medium grid with 2 points per trip dot height and (c) a fine grid with 4 points per trip dot height. The flow is from top to bottom and extends approximately 200 inlet boundary-layer thicknesses in this visualization.

is based on the maximum achievable resolution attainable on the full aircraft configuration. For example, even a fine grid with 4 points per trip dot height would approximately scale to 2 billion control volumes on the full aircraft simulation. The wake-spreading angles for the three considered grids are reported in Figure 5.3. It is clear that the coarsest grid fails to capture any meaningful impact associated with the trip dot, as patterned bands of transition appear at regularly spaced intervals that approximately match the grid spacing, meaning that the transition is almost entirely grid-induced in this case. On the medium mesh with 2 points per trip dot height, some impact of the trip dot is visible, with transition occurring earlier than in the untripped regions of the flow. Further, the interaction of the grid with the periodic boundary face leads to disturbances injected from the spanwise boundaries that contaminate the rest of the flowfield. Finally, with 4 points per trip dot height, a clean wake with a spreading angle of 4.5° is obtained, which is in reasonable agreement with the experimental results [35].

A further quantitative assessment of the influence of the trip dot on the flowfield was made by evaluating the streamwise skin friction evolution through the centerline of the domain, shown in Figure 5.4. Several comparisons are encapsulated in this plot, including the impact of the trip dot on a plug inflow and the impact on an inflow having a Blasius boundary-layer. An important conclusion is that regardless of the inflow chosen, the trip dot successfully promotes an elevated

level of skin friction in the wake of the trip dot. This elevated level persists far downstream of the trip dot. This study of the sensitivity to the influence of the inflow conditions is motivated from evidence that for a flow over an aircraft wing, the boundary-layer profile upstream of the trip dots is not necessarily well-known. In this sense, the plug inflow serves as an extreme case of an underresolved thin boundary-layer ahead of the trip dot. It is promising that the trip dots can promote a transition-like event even for this inflow. In addition, note that for the untripped case, the flow remains laminar throughout the simulation, which confirms that numerical noise or perturbations in the grid are not leading to transition in this setting.

Finally, we note that the untripped cases approach the same asymptotic skin friction value far downstream, while the tripped cases retain some memory of the inflow, with the plug inflow case retaining a larger skin friction value downstream of the trip dot than the case that had the Blasius boundary-layer at the inlet. This difference in the turbulent state between the tripped cases with differing inlet conditions was at most 10% and was not explored further in this study, as this level of control over the oncoming state of the boundary-layer ahead of the trip dot is not practical in a full aircraft configuration. This sensitivity does, however, provide some indication of the amount of error that could propagate into the wake of the trip dots due to underresolved leading-edge boundary-layers. Based on these findings, we selected 4 grid points as an entry point into trip dot resolving simulations in external aerodynamics at Re_k of order 1000, which is typical in wind tunnel experiments of aircraft [60, 22].

5.3.2 Impact of trip dots on the NASA transonic CRM

With the single trip dot validation case complete, the influence of the full array of trip dots was explored on the transonic CRM configuration. In transonic external aerodynamics problems, experiments conducted at wind tunnel Reynolds numbers require physical tripping either via trip dots, active blowing/suction or surface roughness to achieve comparable boundary-layer characteristics to realistic flight Reynolds number conditions [60]. The wind tunnel data used as a reference in this study were collected at the NASA Langley National Transonic Facility (NTF) and used cylindrical trip dots to promote transition in the flow over the CRM. These trip dots were placed at 10% of the local chord on the suction side of the wing along the entirety of the span as well as near the fuselage nose. The trip dots were made of a vinyl adhesive material and measured just $0.003'' - 0.0035''$ at model scale [61] (alternatively, the trip dot heights could be viewed as a percentage of the mean aerodynamic chord length of $7.447''$, in which case they would be between $\approx 0.04\% - 0.045\%$ of a chord length in their height). They measured $0.05''$ in diameter and were spaced at a distance of $0.1''$ apart, resulting in a total of ≈ 315 total trip dots along the span of the wing, which was $63.47''$ in length at model scale. These trip dots were added to the baseline CRM geometry using a CAD tool, and they match the experimental dimensions, spacing and location exactly, as provided in [60]. Figure 5.5 shows the qualitative impact of this trip dot array. We can see from this image

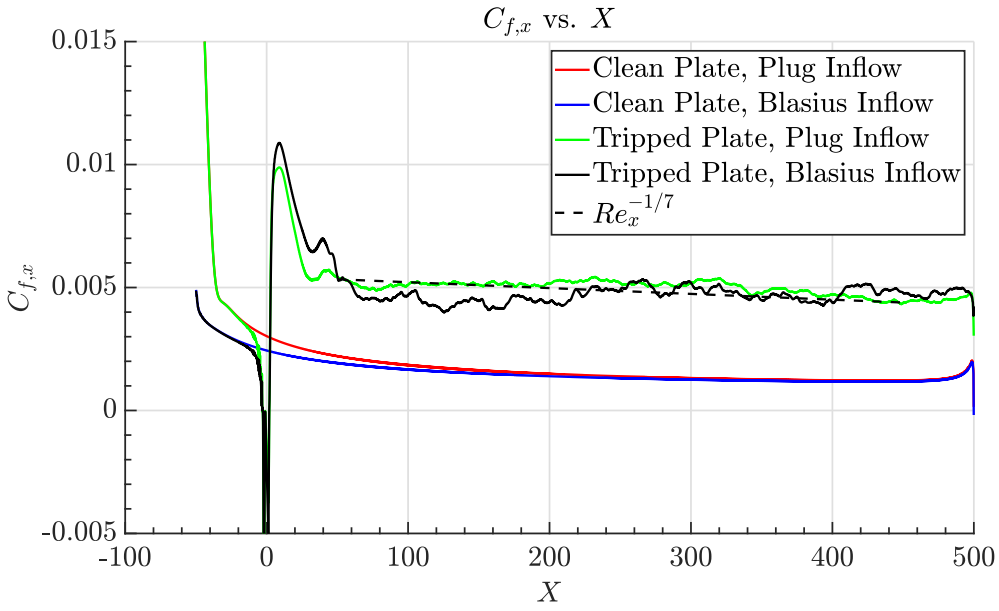

Figure 5.4: Centerline slice of time-averaged streamwise component of skin friction plotted against the streamwise direction normalized by the inlet boundary-layer thickness, ($X = \frac{x}{\delta_0}$) on the fine grid, which has 4 points per trip dot height. The location $X = 0$ is at the point of the trip dot.

that the trip dots are effective at promoting an earlier and more consistent transition front. In these cases, the wake spreading observed in the single trip dot case is not reproduced, potentially owing to the aggregate effect of the trip dot array suppressing this mechanism. The location at which the projected first cell velocity abruptly changes from light to dark is the shock location. At this point, the strong adverse pressure gradient of the shock strongly decelerates the flow and the turbulent structures grow rapidly in size.

A quantitative assessment of the influence of the trip dots is shown in Figure 5.6 by means of the surface pressure coefficient. Reference experimental data from the NTF [61]. The grids for the trip dot–resolving and clean wing cases are identical except in the immediate vicinity of the trip dots and are both number about 750 Mcv (Meshes C & D) . Due to computational cost considerations, a full 4 points per trip dot–height resolution was not possible, as the flat-plate calculations indicated was necessary. Instead, approximately 3 points were used to resolve the trip dot height. Despite this marginal resolution level, the earlier transition associated with the inclusion of trip dots on the CRM wing led to improved shock locations relative to the experimental data on the suction side of the wing. The perturbation in the C_p plot around 10% x/c reflects the location at which the trip dots are installed in the simulations. Overall, these numerical experiments highlight the importance of laminar-to-turbulent transition in transonic external aerodynamic flows over aircraft, suggesting that careful representation and grid resolution around these protuberances are necessary to accurately

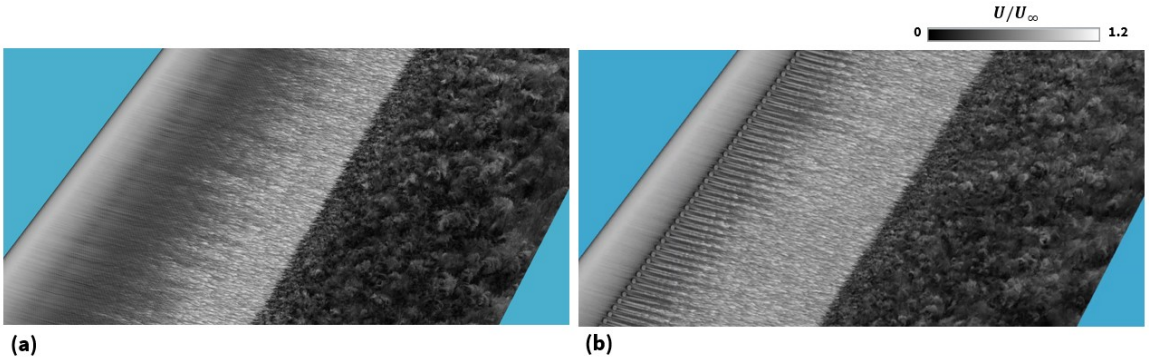

Figure 5.5: Instantaneous velocity magnitude projection on the suction side of the Transonic CRM wing showing (a) an untripped and (b) a tripped case where trip dots have been placed along the entire span of the wing at a constant spanline of 10%. Flow is from left to right.

predict engineering quantities of interest such as pressure coefficients in transonic aircraft flows.

5.4 Impact of symmetry plane boundary condition

It is the standard practice of the external aerodynamics community to perform aircraft simulations on a half-span domain using a symmetry plane boundary condition on the center plane to save on computational cost. This has been the baseline approach for the duration of the High-Lift [65, 62, 63, 64] and Drag Prediction [48, 43, 75, 74, 47, 71] workshop series. These practices, however, are established for steady-state Reynolds-Averaged Navier Stokes (RANS) simulations in which instantaneous turbulent fluctuations are not explicitly resolved. In LES, a study of the appropriate symmetry plane treatment needs to be conducted, as the component of turbulent fluctuations normal to the center plane violates instantaneous symmetry even if the time-averaged flow is symmetrical in the span. Constraining these fluctuations to be bounded by a no-penetration symmetry boundary condition is not a true representation of the instantaneous flowfield even if the time-averaged flow does obey no-penetration in the mean. For this reason, an exploration on relatively coarse grids (Meshes A & B) was performed to establish sensitivities to the symmetry plane boundary condition. Figure 5.7 shows the differences between these where the full-span case is simply the half-span case that includes the mirrored part of the domain.

A sensitivity of the time-averaged flowfield to the treatment of the symmetry plane is found to be present for these calculations, with the full-span simulations showing improved characterization of the shock structure relative to the half-span cases, as shown in Figure 5.8. Surprisingly, the largest differences between the sectional pressure measurements are concentrated on the outboard part of the wing ($\eta \approx 0.5$ and beyond). One might expect that the strongest differences between the cases would be observed near the symmetry plane; however, we find the opposite: that large changes in

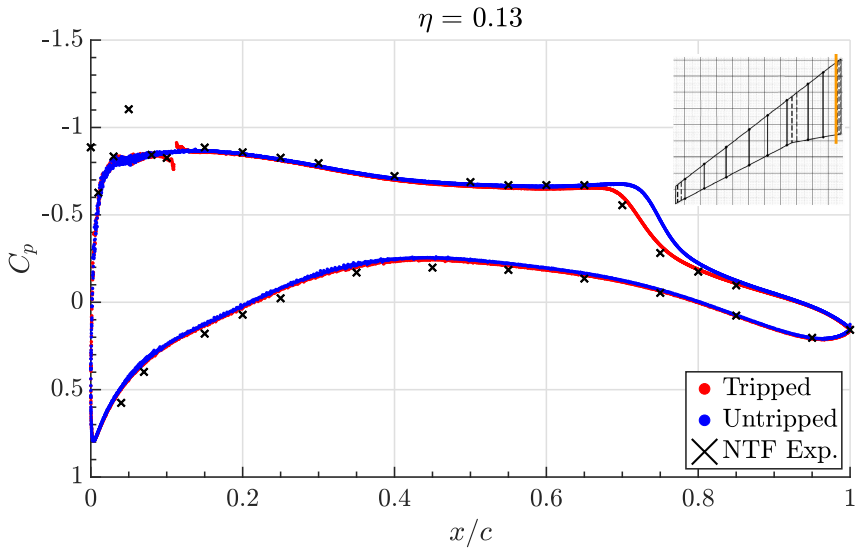

Figure 5.6: Average pressure coefficient slice through an inboard wing station at $\alpha = 4°$ showing the impact of the wing trip dots. The earlier shock location observed in the tripped case is consistent with the experimental measurements. The vertical orange line in the inset graphic shows the location of the pressure cut along the span of the wing. The effect of tripping is most pronounced at this inboard pressure belt station.

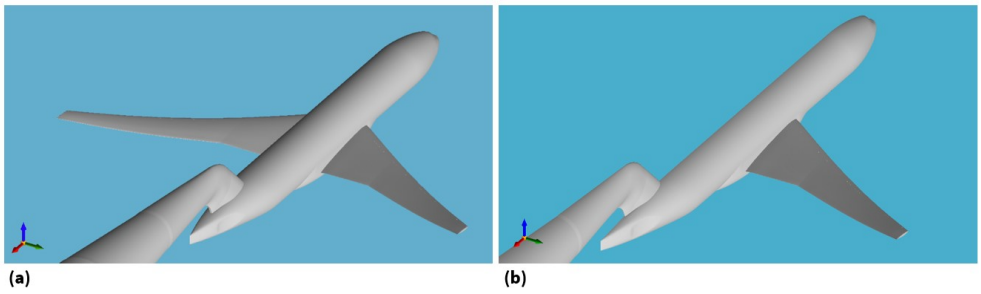

Figure 5.7: Geometry of the Transonic CRM including (a) the full span and including (b) a symmetry plane. This configuration will be used to evaluate the sensitivity of the LES calculations to the center plane boundary condition.

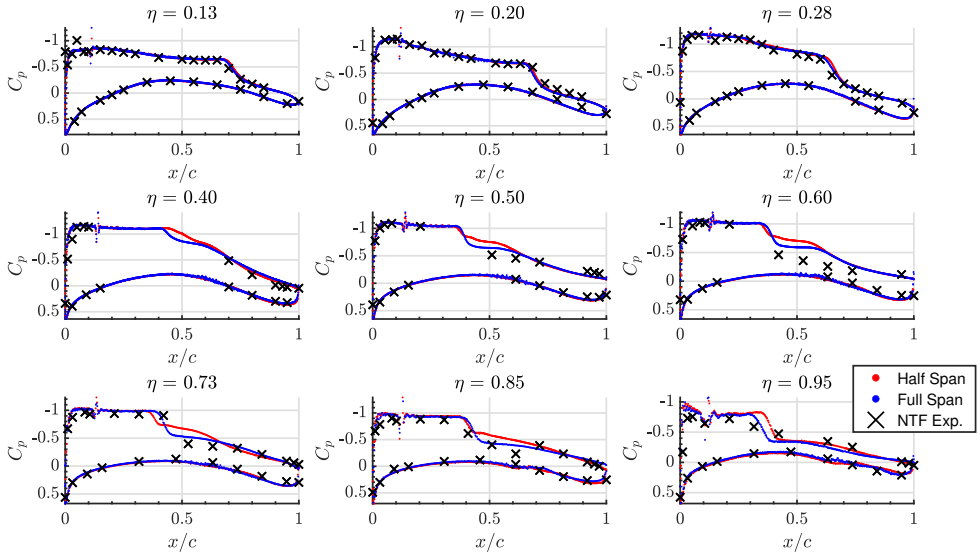

Figure 5.8: Spanwise pressure measurements for the half-span versus full-span Transonic CRM calculations. The largest differences between the half-span and full-span calculations are observed starting mid-span and moving outboard.

the shock structure are concentrated near the wingtip, potentially owing the the 2D-type versus 3D-type shock-relieving effect often studied in supersonic flows over cylinders and spheres [10].

Significant instantaneous asymmetries in the rolling moment (computed about the x-axis passing through the fuselage nose) do arise in this flow despite the time-averaged rolling moment being approximately zero. Figure 5.9 shows a plot of rolling and pitching moments over the course of a simulation horizon of 100 flow passes based on the mean aerodynamic chord. The instantaneous oscillations in the rolling moment are larger than those in the pitching moment. We find on the basis of these simulations that while the time-averaged value of the rolling moment is nearly zero, it is inadequate to perform the calculations half span and use a symmetry plane boundary condition due to the presence of significant instantaneous asymmetries in the turbulent flowfield.

Finally, the integrated forces and moments are compared in Figure 5.10. It is clear from this plot that for a given grid, the full-span configuration is in better agreement with experiment than the half-span case. This result is found on a relatively coarse grid that is sized to fit 512 grid points across the mean aerodynamic chord. Further studies related to grid refinement and grid topology sensitivities will be established later in the report. However, we conclude on the basis of this experiment that symmetry plane effects are important for LES simulations of transonic aircraft flows and recommend the use of full-span domains.

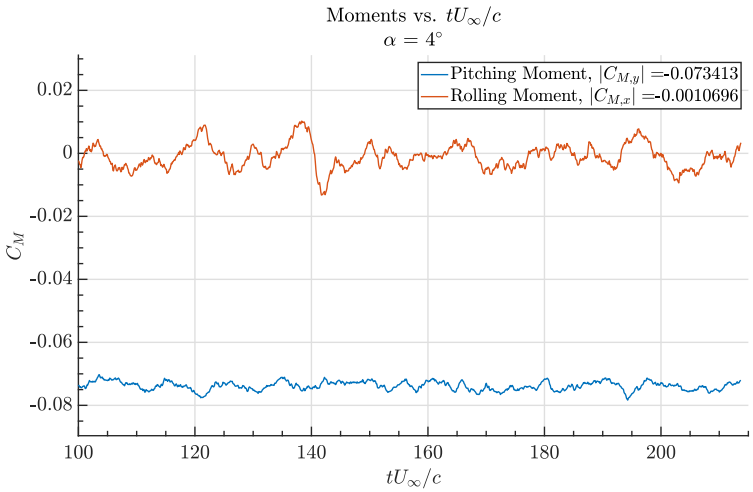

Figure 5.9: Rolling and pitching moments on the Transonic CRM for the full-span case. Instantaneous asymmetries in the rolling moment are observed despite the time-averaged value being near zero.

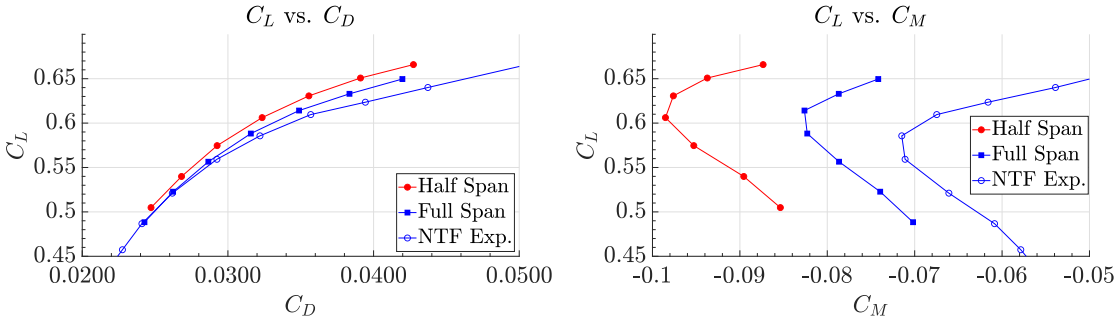

Figure 5.10: Non-dimensional (a) drag polar, and (b) lift versus pitching moment coefficient plots for the half-span and full-span LES cases compared against experimental data from the NTF tunnel [61]. The simulations are run on a $\approx$ 50-Mcv mesh in the half-span case (Mesh A) that is mirrored about the symmetry plane when the case is run in full-span mode (Mesh B). In each metric, the full-span simulations outperform the half-span ones.

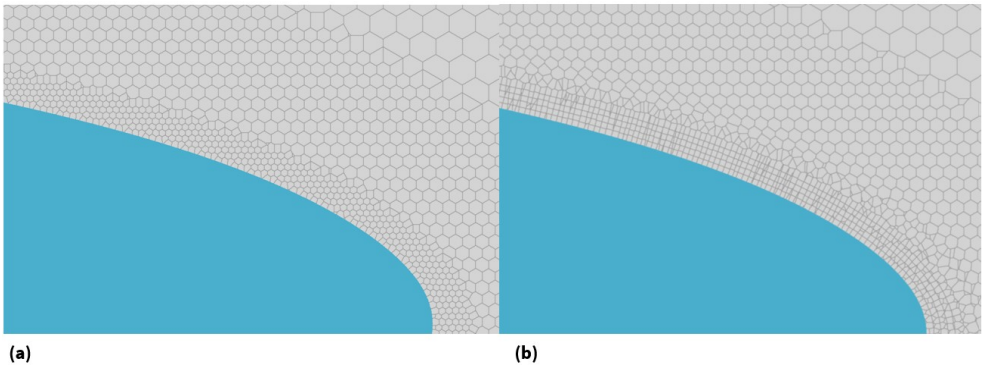

Figure 5.11: Slice through the leading edge of the CRM wing showing (a) the default HCP grid topology and (b) a modified stranded mesh in the wall-adjacent region.

5.5 Sensitivity to grid topology

The shock location in transonic flows is sensitive to the details of the boundary-layer because the shock location is set by the inviscid acceleration (dictated by the effective body shape in which viscous effects are important). For this reason, a grid that used prismatic boundary-layer elements in which additional refinement in the wall-normal direction was explored (see Figure 5.11 for a schematic). Prismatic-type boundary-layer elements (hereinafter referred to as stranded grids) have the advantage over isotropic hexagonally close-packed (HCP) meshes in viscous regions because it is possible with this grid topology to stretch the mesh in the wall-normal direction in order to resolve the vertical length scale of the wall-bounded turbulence, which shrinks at a faster rate than the streamwise or transverse length scales in the boundary-layer as the wall is approached [52].

Figure 5.11 shows a comparison between an isotropic HCP and an isotropic strand-mesh. Further explorations of appropriate aniostropy ratios for stranded mesh calculations were conducted, and it was found that superior performance of the solution on the stranded mesh was achieved for an anisotropy ratio of 4:1 and a maximum wall-normal stretching ratio of 1.15. Also concluded from this study was the importance of the wall-model LES exchange location on stranded meshes, and all stranded cases presented in this report use a second cell centroid exchange height. This same sensitivity was not observed on HCP meshes, and all HCP calculations therefore used the first cell centroid as the exchange height. Figure 5.12 shows the result of this sensitivity exploration to grid topology. The HCP and strand cases were not chosen to be analogous in terms of point count/cell size; instead, they should be viewed as the best achievable result with each topology at the bleeding edge of computing affordability on modern GPU computing resources. The stranded-mesh (Mesh E) result reported in Figure 5.12 has 2× finer wall-normal resolution, while being 2× coarser in the stream/spanwise directions than the reference HCP solution (Mesh C). Additionally, because of the difficulty associated with creating a high-quality prismatic mesh around the leading-edge trip dots,

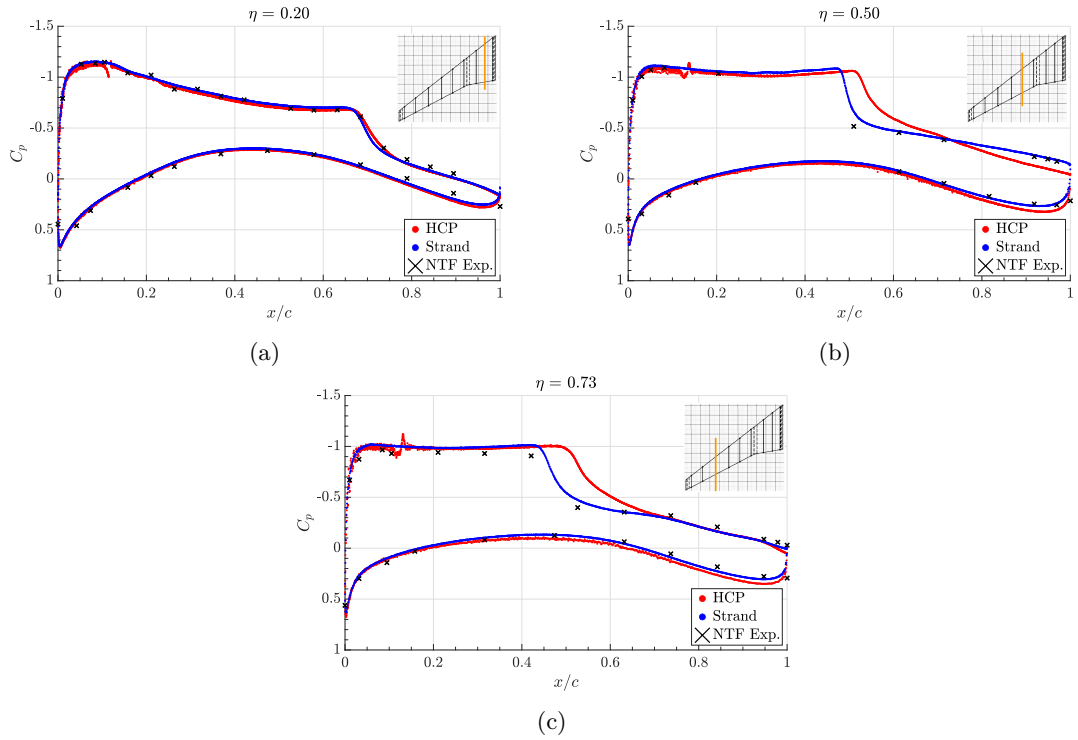

Figure 5.12: Average static pressure measurements for the Transonic CRM at an angle of attack of $\alpha = 4.0°$ at three different stations along the span of the wing, ranging from (a) inboard to (b) mid-span to (c) outboard. The comparison includes results computed on a fine mesh of isotropic HCP grid elements ($\approx$ 750 Mcv, Mesh C) and a strand mesh solution that is 2$\times$ finer than the HCP mesh in the wall-normal direction and 2$\times$ coarser in the stream/spanwise directions ($\approx$ 270 Mcv, Mesh E). The glitch in the HCP results is due to the presence of geometric trip dots in these calculations which are absent from the stranded cases.

these have been removed from the stranded calculations.

Further comparison was made to the surface flow patterns predicted by each mesh topology, shown in Figure 5.13. Experimental reference data for the surface flow was not available, so OVERFLOW RANS data are shown as a benchmark. The RANS data represent a best practice result from the Sixth Drag Prediction Workshop (DPW6) contribution of the Boeing Southern California group using the OVERFLOW solver. The Spalart-Allmaras turbulence model with Quadratic Constitutive Relation (SA-QCR) was used for the RANS simulations, a modeling practice known to perform well in this flow, particularly in the wing juncture region [71]. The RANS calculations were run in fully turbulent mode with no explicit representation of the trip dots. The shock in this case is strong enough to induce a pocket of separation in the mid-span region, visible in the surface streamline patterns by the regions in which the streamlines become parallel to the trailing-edge sweep. The extent of the separation region from the HCP solution is smaller than the RANS solution (as was observed in the sectional pressure measurements that showed that the shock lies too far aft). In

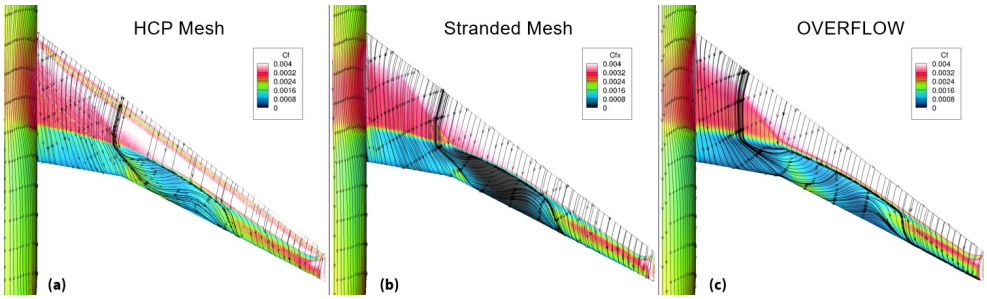

Figure 5.13: Average surface skin friction magnitude and surface streamlines from three simulations, including LES simulations with best-achievable solutions on (a) HCP meshes, (b) stranded Meshes and from (c) RANS simulations using the OVERFLOW solver.

the stranded solution, the extent of the separation is approximately in agreement with the RANS solution, while the strength of the separation is more intense (lower values of skin friction magnitude in the bubble).

Finally, the integrated forces/moments are compared against reference data in Figure 5.14. On the basis of these results, it is clear that stranded meshing offers advantages over HCP grids in transonic flows, particularly at higher angles of attack at which significant shock-induced separation is observed. However, the results from the best-practice strand mesh do still exhibit some degree of delay in the pitch break (where the C_M curve abruptly becomes less negative) in Figure 5.14(c) associated with a delay in the onset of shock-induced separation. Still, significant benefit is achieved from the use of stranded meshes in the prediction of lift, drag and pitching moment for this flow.

5.6 Novel modeling approaches

Novel subgrid-scale (SGS) and wall-modeling approaches for LES such as those discussed in [2] and [21] have recently been developed and deployed on the transonic CRM to assess their applicability beyond canonical flows (including turbulent channel flows, flow over the Boeing speed bump and Sandia transonic hump). A key issue arising in wall-modeled LES calculations of external aerodynamic flows is the problem of non-monotonic convergence of the quantities of interest. Most notably, on the Boeing Gaussian Bump flow, this has been a persistent issue with standard modeling approaches [78]. Non-Boussinesq SGS models such as the Dynamic Tensor Coefficient Smagorinsky Model (DTCSM) developed by [2] have been shown to rectify this issue. Figure 5.15, however, shows a mixed result on this more complex flow. The prediction on a given grid is improved by the DTCSM SGS + Sensor Wall Model, when compared to its Dynamic Smagorinsky + Equilibrium Wall Model (EQWM) counterpart. However, the deficiency of the coarse-grid result being more accurate than the fine-grid result is still observed. The inclusion of a recently developed sensor-aided equilibrium wall model [1] is aimed at identifying regions of strong adverse pressure gradients. In these regions, the model

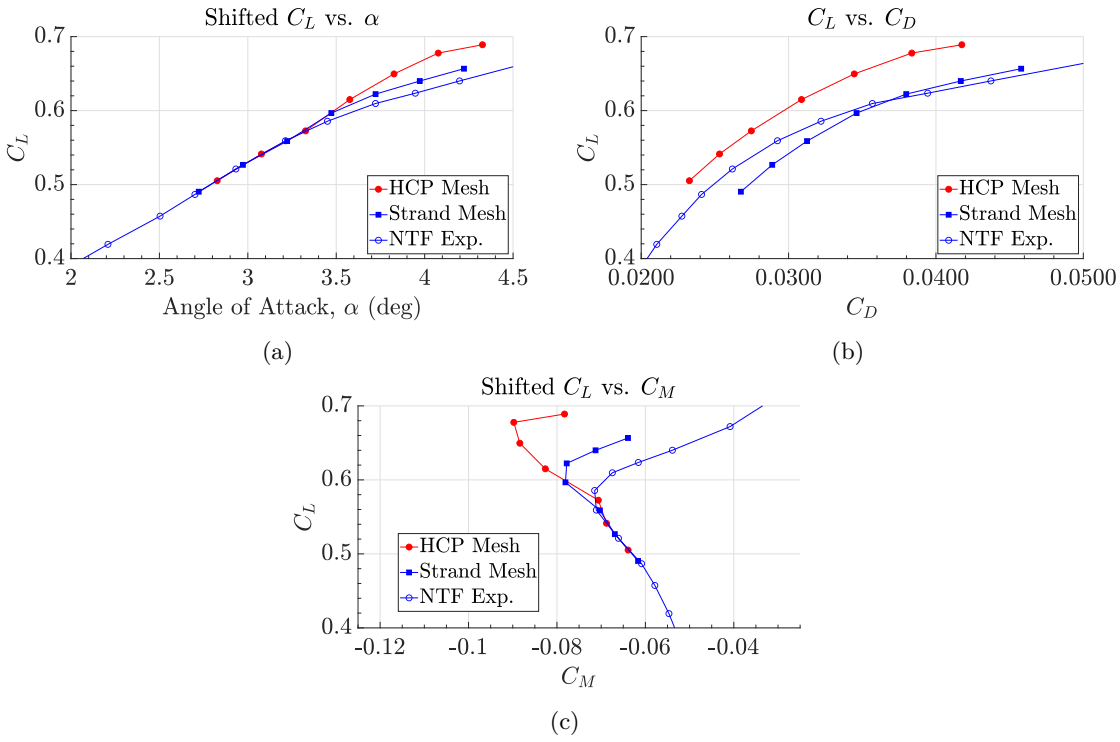

Figure 5.14: Average forces/moments from the LES simulations of the Transonic CRM, including (a) the shifted lift, (b) drag polar and (c) shifted pitching moment. The shifting of the lift and pitching moment plot is performed as suggested in [70].

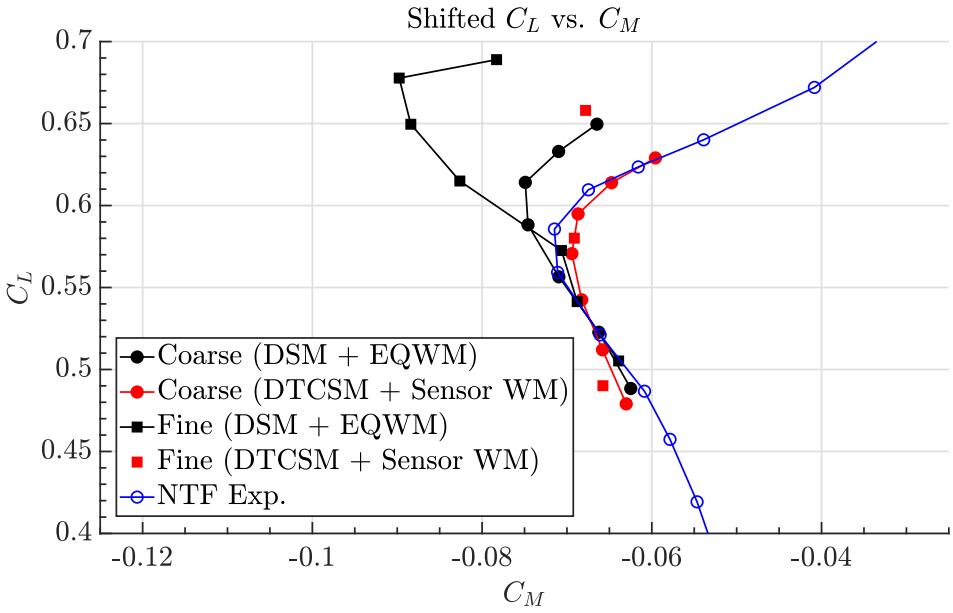

Figure 5.15: Average lift versus moment plot from the LES simulations of the Transonic CRM with results from novel SGS and Wall Models included in the plot. The shifting of the pitching moment plot is performed as suggested in [70].

includes the subgrid-stress in addition to the equilibrium wall-modeled stress when prescribing the total wall stress. This method is developed as a proxy to account for the non-zero resolved stresses on the wall and has been successfully applied to the case of the Boeing speed bump. It is clear that while at this juncture novel modeling approaches such as those developed in [2] provide tangible benefits over standard modeling choices, the issue of non-monotonic grid convergence is not fully rectified by these models on full-scale aircraft. Further study is needed to identify the reason behind this deficiency but may include the presence of compressibility effects, aircraft three-dimensionality effects (finite span wing, spanwise pressure gradient) or Reynolds number differences.

The pressure measurements at the 4-degree condition are shown in Figure 5.16 from the coarse-grid and fine-grid DTCSM + Sensor WM cases (which are considered the best-practice modeling approach based on experience in canonical flow validations) and show that in general, the prediction of the lift and moment achieved on the coarse grid is approximately for the right reasons because the shock location is in good agreement with available experimental data on the coarse grid, while on the finer mesh, the shock moves aft, which is consistent with the collapse of the shock-induced separation bubble and leads to higher lift levels and more nose-down pitching moment tendencies.

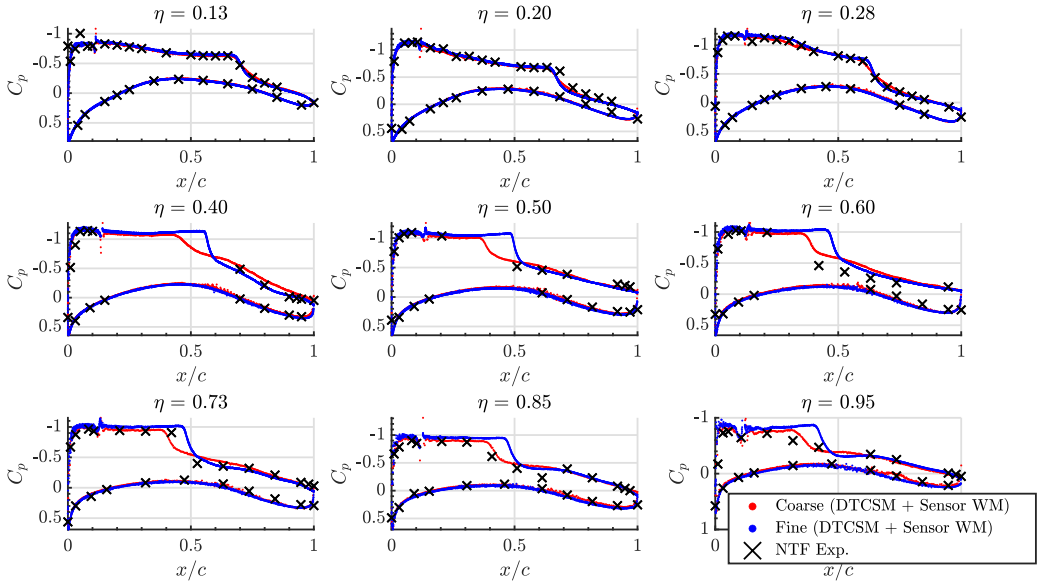

Figure 5.16: Sectional pressure measurements from the Sensor Wall Model + Dynamic Tensor Coefficient Smagorinsky Model calculations for the Transonic CRM at angle of attack of 4 degrees.

5.7 Boundary-layer details

Although the engineering quantities of interest for the flow over an aircraft are integrated forces and moments, it is important to quantify the effect of the Reynolds numbers and pressure gradients in the development of the boundary-layer over the complex flow for informing modeling efforts. In this work, multiple wall-normal velocity profiles are collected and reported from one of the more accurate and refined calculations at the 4-degree angle of attack. These profiles are sampled at the leading edge of an inboard wing station (in order to understand the nature of the flow ahead of the shock and away from the separation bubble and wingtip effects) and within the separation bubble at mid-span. The wall-parallel velocity profiles ($U_{||}$) are shown in Figure 5.17. The generalized, compressibility effects–inclusive, method developed in [32] is used to evaluate the boundary-layer edge (δ_{99}). The point of maximum stagnation enthalpy ($h_{0,ref}$) is used to establish the far-field conditions used in the method for computing the boundary-layer edge.

Table 5.2 shows the quantitative boundary-layer measures associated with Figure 5.17. Note that these numbers apply to full-scale aircraft at wind tunnel–type Reynolds numbers (typically an order of magnitude smaller than the true flight Reynolds number). Although the experiment is at model scale, the simulations are run at full scale, which explains why the measures of boundary-layer parameters apply to a somewhat synthetic case of a full aircraft at wind tunnel Reynolds number.

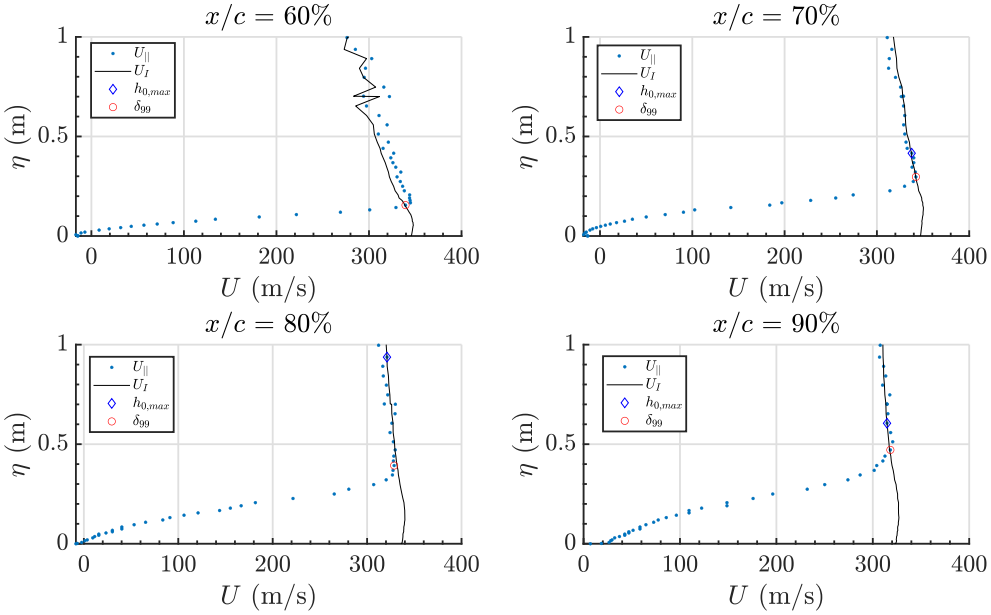

Figure 5.17: Wall-parallel velocity profiles at mid-span (50% of the semi-span) within the separation bubble. The boundary-layer edge is determined using the compressible method of [32] for flows in non-equilibrium.

Still, the details of the boundary-layer characteristics will be informative from the perspective of understanding the character of the boundary-layer and for the selection of pertinent canonical cases for future reduced-order modeling efforts. It is apparent that the flow over an aircraft wing experiences a wide range of Reynolds numbers based on momentum thickness (Re_θ) and shape factors (H), and suitable canonical cases must be chosen that are representative of these nondimensional measures. A different case must likely be chosen to represent the thin leading-edge flow, and a different case would be suitable for the thick separated flow near the trailing edge. Finally, even for the finest grids considered in this work, the thin leading-edge boundary-layers are marginally resolved with at most 5 grid points within the boundary-layer. This points to another important necessity of future LES modeling efforts: the need to predict quantities of interest with $\mathcal{O}(5)$ points within the leading-edge boundary-layer.

5.7.1 Initial Buffet Prediction

Transonic buffet is a corner of the flight envelope flow condition in which, either by increasing the angle of attack or Mach number, flow unsteadiness associated with shock oscillations begins to play an increasing role in governing the flow around an aircraft wing. The phenomenon has been studied extensively using Unsteady RANS (URANS) and analytical techniques [19] in the context

Table 5.2: Boundary-layer characteristics for the Transonic CRM flow both near the leading edge (spanwise location 0.2) and within the separation bubble (spanwise location 0.5). A wide range of Reynolds numbers and shape factors are observed as the boundary-layer thickens and experiences shock-induced separation.

Spanwise Location	**Chordwise Location (%)**	δ_{99} **(cm)**	**Shape Factor (H)**	Re_θ	$\frac{\delta}{\Delta}$
0.2	1	0.76	1.86	472	2
0.2	5	1.24	2.50	657	3
0.2	10	1.25	2.44	729	4
0.2	20	2.62	2.34	1,505	5
0.5	60	15.5	6.33	10,558	20
0.5	70	29.7	6.67	18,006	28
0.5	80	39.3	4.88	28,648	29
0.5	90	47.1	3.80	41,382	30

of a canonical NACA 0012 flow. In the context of full aircraft flows, limited to no data concerning buffet onset is available. Still, it is interesting to explore the predictive potential of LES in this flow regime which is very important for certification by analysis. Figure 5.18 shows the instantaneous force and moment time histories for the stranded mesh calculations.

It is important to note that the onset of buffet in aircraft design is typically identified as the point at which an accelerometer placed in the cockpit experiences a minimum of 10% oscillations about the mean in the normal direction. This can be interpreted in the present context as the point at which a 10% oscillation about the mean lift value is observed. Clearly over the range of angles of attack considered in this work, that amount of force oscillations was not observed. In the pitching moment, which shows more sensitivity to flow separation that integrated lift, we do begin to see meaningful oscillations about the mean even at an alpha of 3.75°, which can be interpreted as a precursor to buffet. This data suggests that while buffet onset may be within the predictive scope of LES, a more controlled test case is needed to explore this phenomenon further. The chosen case for this purpose is the NACA 0012, which contains a wealth of experimental data pertaining to buffeting flow conditions [54]. The buffet boundary was identified by this experiment and is reproduced in Figure 5.19. The M = 0.75 flow condition was selected for this study. First, a static $\alpha = 2°$ case was chosen to establish confidence in the LES predictions ahead of buffet onset. The static pressure prediction from this case is shown in Figure 5.20. The agreement with the experiment in regards to the shock location and strength is excellent and validates that the static pre-buffet case is predicted appropriately, a necessary precursor to buffet prediction.

Next, a rotating mesh version of the charLES code was used to identify the onset of buffet in the NACA 0012 flow at M = 0.75. A contour of the Mach number through a center plane of the flow is shown in Figure 5.21. In this image, a shock wave and the associated shock-induced separation bubble are clearly visible as the region in which the Mach number abruptly drops. For the rotating

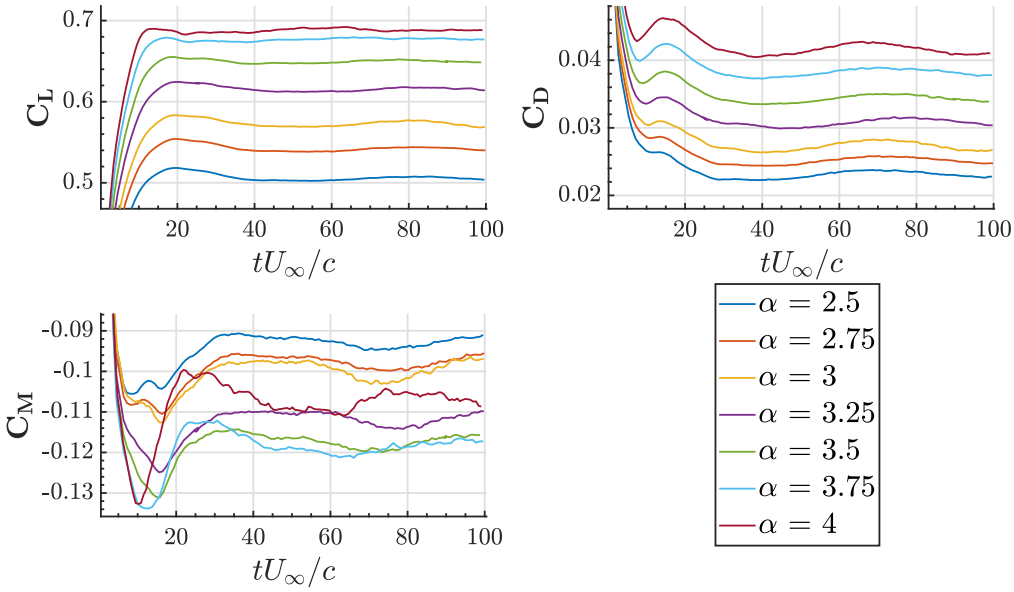

Figure 5.18: Time history of the forces/moments from the stranded mesh LES simulations of the Transonic CRM, including the lift, drag, and pitching moment. The calculations are run for 100 flow pass times based on the mean aerodynamic chord and averaging takes place over the final 50 flow passes.

cases, the airfoil was started from a low angle of attack and rotated slowly through a range of alpha's of interest until significant force oscillations were observed. The time history of the normal force is shown in Figure 5.22. Two rotation rates 1° per second and 5° per second were chosen to understand the sensitivity of the predictions to rotation rate. In experiments, the rotation rates are typically significantly lower than in simulations (usually up to 0.1° per second), but this is not achievable presently with LES because of the prohibitively long temporal integration window needed to match the low experimental rotation rates. Despite this, we have established low sensitivity to the rotation rate in the present simulations, with both rates agreeing well with one another and with the static alpha case (the same case which showed excellent agreement with static pressure measurements from the experiment). We can conclude based on this that rotation does not play an important role in the onset of buffet in this problem and static angle of attack calculations can be used to identify the onset of unsteadiness.

Still, the rotating alpha cases were useful in identifying the location at which buffet onset was likely. We find on the basis of the rotating airfoil cases that significant force oscillations appear between 4 and 5 degrees. Further investigation into these set points was carried out at fixed alpha and it was found that 5 degrees was the minimum alpha needed to achieve at least 10% oscillations

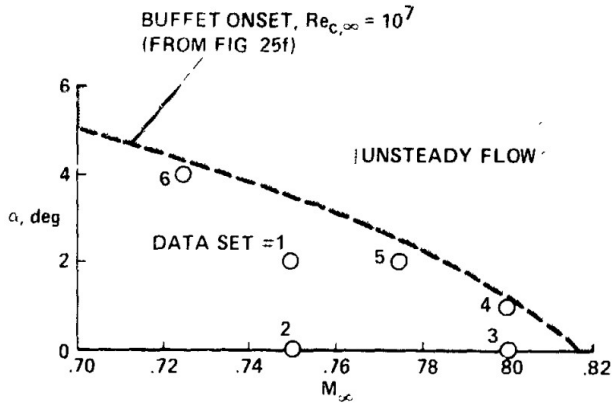

Figure 5.19: Buffet boundary diagram reproduced from [54]. Buffet can be achieved either by increasing the angle of attack or by increasing the Mach number.

about the mean lift coefficient value. This was 1 degree delayed relative to the experimental measurements. Figure 5.23 shows the amplitude spectrum associated with the lift force time history of the transonic flow over the NACA 0012 airfoil at 5 degrees. We can see that the lift force oscillations exist predominantly at relatively low frequencies (≈ 17.5 Hz). Table 5.3 shows the non-dimensional shock oscillation frequencies associated with buffet. LES predicts shock oscillations that are slightly less frequent than those observed in reality (with uncertainty due to the number of statistical samples of the oscillation phenomenon being limited from the simulation due to computational cost considerations) and the oscillations begin about 1 degree later than in the experiments. The delay in the shock oscillation onset is consistent with the conclusions of Figure 5.14 in which subfigure (c) showed a delay in the pitching moment break which is associated with shock-induced separation and is a precursor to buffet onset. Overall, we observe that state-of-the-art LES methods do provide value in the prediction of buffet onset due to their time-accurate nature, but quantitative measures such as the alpha of initial buffet or the precise shock oscillation frequency are subject to $10-20\%$ errors despite accurate prediction of shock location ahead of buffet onset in more benign flow regimes.

Table 5.3: Non-dimensional shock oscillation frequencies of the transonic 0012. Experimental frequencies are from [54]. The LES uncertainty is a statistical measure of uncertainty associated with the bin width used in the FFT. The non-dimenstional frequency is defined as $\bar{f} = 2\pi f c/U_\infty$.

Case	**Alpha**	f
Experiment	4	0.47
LES	5	0.43±0.025

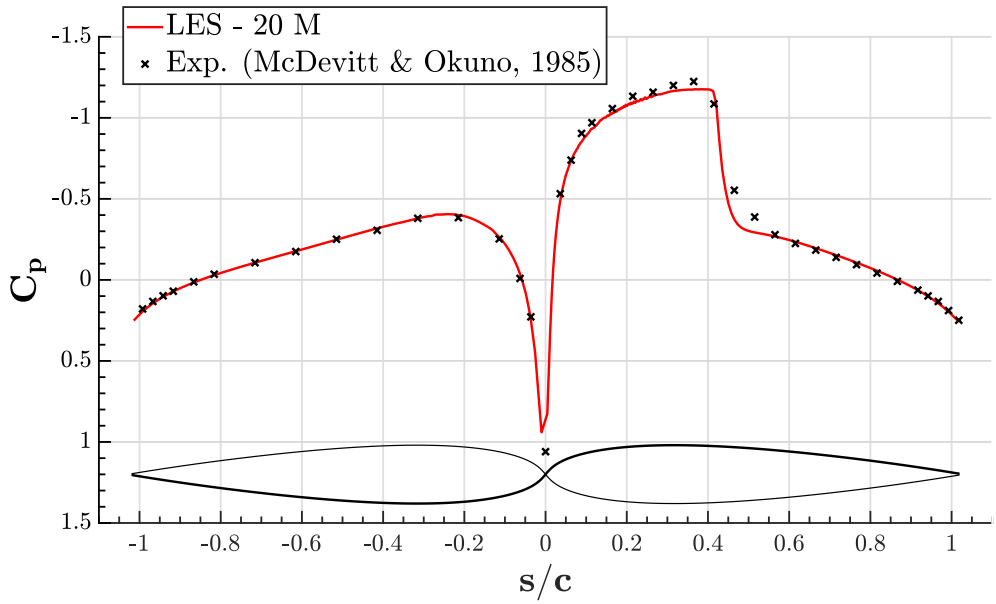

Figure 5.20: Centerline pressure prediction by LES at M = 0.75, Re = 10M, and $\alpha = 2°$, ahead of buffet onset.

5.8 Summary

In this chapter we have described simulations of the transonic version of the NASA CRM, a benchmark flow for drag prediction and CFD validation in the compressible flow regime, representing a commercial aircraft in cruise configuration. These calculations are among the first forays into use of LES for this kind of flow and initial heuristics were developed around best practices, including the use of an array of small cylindrical trip dots lining the leading edge of the wing at a constant spanline of 10% (which replicates the way that turbulent transition was achieved in the experiment), simulation of the full span configuration (which again matches the experimental setup), inclusion of the sting mounting apparatus at the tail of the aircraft, and the use of prismatic boundary layer grids of maximum anisotropy of 4 to 1 in the wall-normal direction relative to the streamwise and spanwise directions. Also explored were the combination of a novel non-Boussinesq subgrid model, the Dynamic Tensor Coefficient Smagorinsky Model, and a sensor-based wall model, which identifies regions of flow separation and turns of the equilibrium wall model in those regions. The integrated effect of these models resulted in an improved characterization of the pitching moment break, particularly on coarse grids. The solutions were plagued by the problem of non-monotonic grid convergence observed in canonical smooth body separation flows, though the novel modeling approaches did tend to improve the quality of the solutions on a given grid relative to traditional models. The boundary

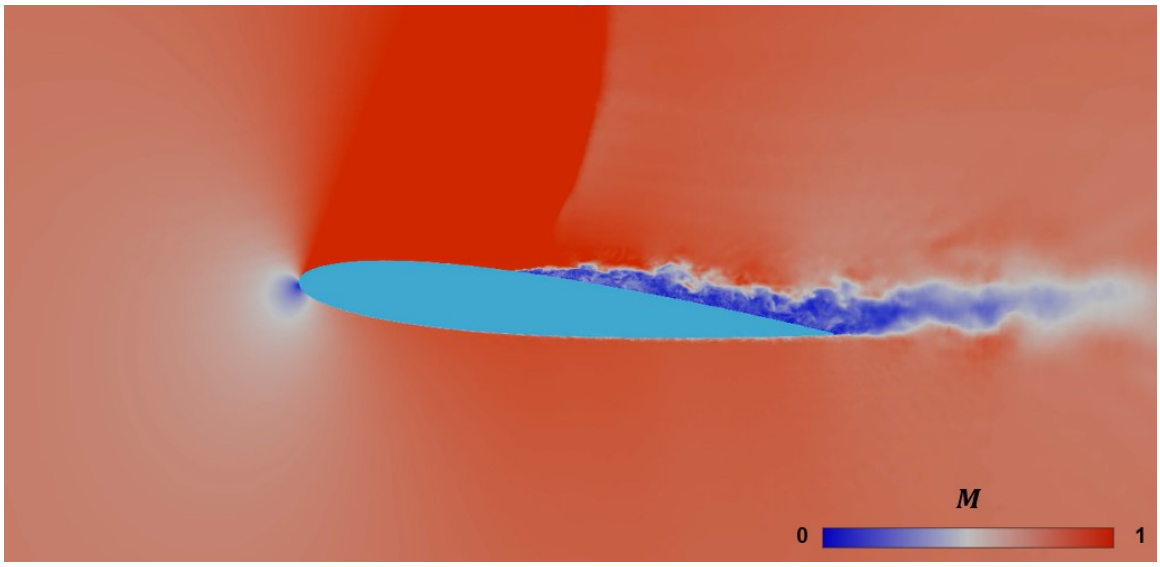

Figure 5.21: Center plane Mach number contours of the rotating mesh NACA 0012 simulation. The airfoil was dynamically rotated with time to identify the angle at which significant force oscillations were observed.

layer details were computed based on profiles of the wall-parallel velocity component and it was found that a wide range of Reynolds numbers and boundary layer thicknesses is encountered in this flow, which poses a difficult wall-modeling task due to the diversity of flow phenomena and the wide range of pressure gradients experienced in aircraft flows. The stark contrast between thin leading edge boundary layers which are marginally resolved with up to 5 points on even the finest grids currently achievable up to thick viscous regions of shock-induced flow separation encountering strong adverse and spanwise pressure gradients highlights this challenge. Finally, an exploration into the predictive capabilities of LES with regards to initial buffet was made with promising qualitative results, though further investigation is needed regarding the precise angle of attack at which buffet begins and the shock oscillation frequency once buffet is sustained.

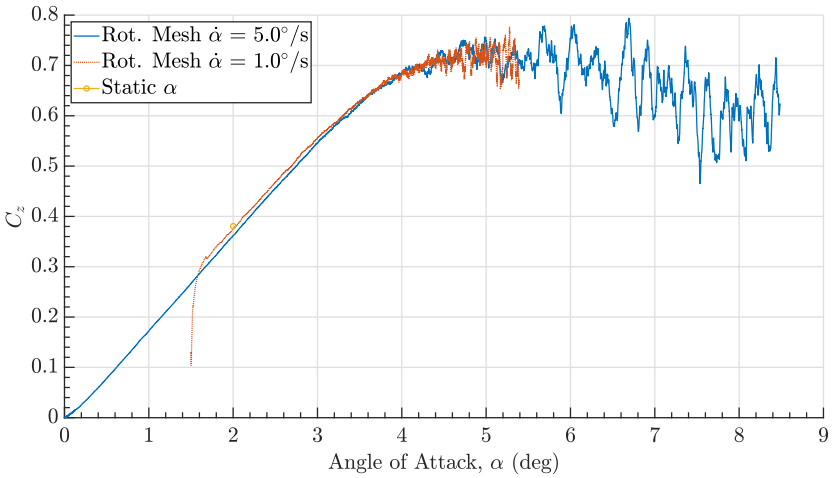

Figure 5.22: Time history of the normal force component for the rotating airfoil cases used to identify the onset of buffet.

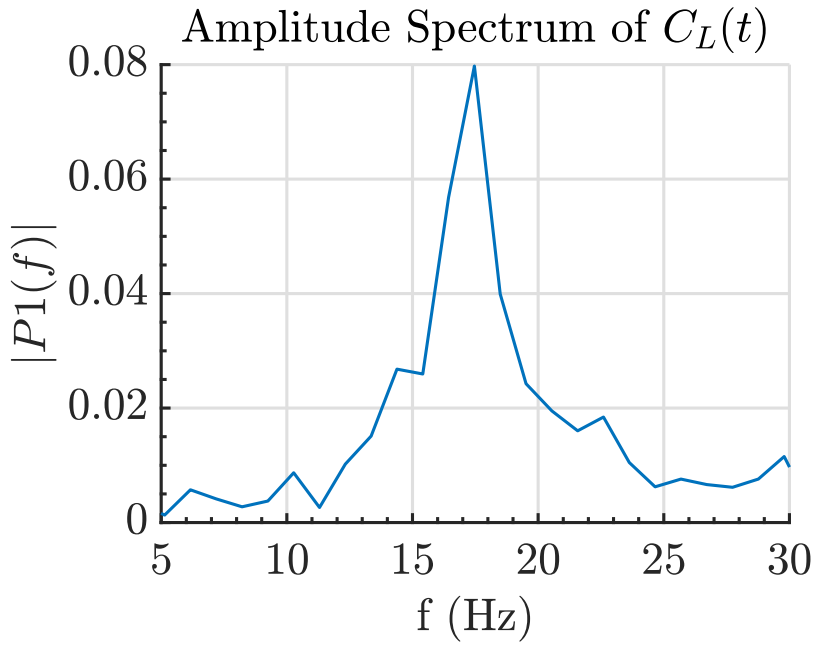

Figure 5.23: Amplitude spectrum of the lift force oscillations at $\alpha = 5°$ for the spanwise-periodic NACA 0012 simulations. Buffet is identified at the point where unsteady oscillations in the lift force are at least 10% about the mean. Mean lift at this condition is ≈ 0.6. f is the dimensional frequency used to construct the non-dimensional frequency, $\bar{f}$.

Chapter 6

Conclusion

This dissertation has described the studies which have been carried out involving the use of Large Eddy Simulation for commercial aircraft applications at the corners of the flight envelope, including high-lift and transonic regimes. Significant progress has been made since the seminal calculations at the 2018 CTR Summer Program to better understand the predictive capabilities of high-fidelity LES in the context of maximum lift prediction. It has been demonstrated that LES can accurately, affordably, and consistently predict both the stall condition of an aircraft in landing configuration and the shock location of a cruise wing configuration. However, several issues remain; including the prediction of low angle of attack flow separation over the flap in high-lift flows and the precise angle of attack at which buffet onset occurs in the transonic regime. Novel non-Boussinesq modeling approaches offer some promise to remedy these issues, but require more detailed study in relevant configurations of engineering significance such as aircraft flows.

In the context of high-lift flows, it was found that at approximate resolutions of 20 points per boundary layer thickness, grid-independent predictions of integrated lift, drag, and pitching moment are achieved, though it was not clear whether viscous length scales are setting the grid resolution requirements as the grid sensitivities are preferentially concentrated towards the wingtip. GPU acceleration enabled calculations to be run within several hours on grids numbering nearly half a billion grid points. It was also found that inclusion of the wind tunnel facility in the simulations improves the prediction of the stall mechanism, particularly for half-span sidewall-mounted aircraft models. Robustness of LES methods with respect to Reynolds number and aircraft configuration were also established by means of studying both the JSM and CRM-HL configurations, which differ in Reynolds number by nearly a factor of 3. Importantly, the superiority of dynamic subgrid-scale models was established, as the ability of the dynamic Smagorinsky subgrid-scale model to predict both smooth body separation over the trailing edge flap and the juncture flow was superior to the static Vreman model. This likely owes to the inherent ability of dynamic models to adapt the dissipation in space and time based on the finest resolved turbulence scales as compared with static

coefficient models which require some calibration of their model coefficients and are therefore less universal. The target set forth by industry to reliably predict maximum lift coefficient to within 3 lift counts and to predict stall to within 1 degree angle of attack was achieved for both of the high-lift configurations considered. LES is therefore poised to be a key enabling technology towards the objectives set forth in the CFD 2030 Vision Report.

In the transonic regime, the prediction of shock location at a cruise condition ahead of buffet onset to within 2% of the experimentally observed location across the wingspan was achieved, but only after the inclusion of wing aeroelastic deflection, the sting mounting system, and by running the simulations in full span mode. Boundary layer transition was found to be critically important in these flows, potentially due to the strong sensitivity of the shock location to the effective body shape of the aircraft which will change significantly depending on the state of the wing boundary layer. For this reason, an array of 300+ experimental trip dots were included in the simulations and their impact on the transition front was resolved by the grid. The more streamlined nature of the flow in this cruise-type configuration made the solutions more amenable to the use of prismatic boundary layer grids of maximum $4:1$ anisotropy in the wall-normal direction, which were shown to improve the quality of solution, particularly as it relates to shock location when appropriate second cell off-wall LES/wall model exchange heights were used. Finally, a tractable workflow for identification of buffet onset was proposed in which rotating mesh capabilities were leveraged to identify the beginnings of unsteady shock motion in the solution.

The results discussed in this thesis represent the state-of-the-art in high-fidelity LES (with physics-based equilibrium wall models and dynamic subgrid-scale models) applications for commercial aircraft flows. The ultimate goal of certifying commercial aircraft by analysis-based means (both computational and experimental) instead of traditional flight testing-based ones remains a lofty goal, but one that high-fidelity LES simulations will play an integral role in making a reality.

Appendix A

Consistency of the Algebraic and ODE Equilibrium Wall Models

The aim of this appendix is to show by derivation and by *a-posteriori* assessment of an LES calculation, that in the low Mach limit in the absence of heat flux at the wall, the simplified algebraic equilibrium wall model in Equation A.1, returns a nearly identical velocity profile (and therefore skin friction) to the equilibrium wall model of [38]. The algebraic version of the model is used in all wall-modeled LES calculations described in this dissertation. Both models neglect non-equilibrium effects such as flow unsteadiness, nonlinear convection, and pressure gradients in the wall modeled layer. The equilibrium wall model of [38] differs from the algebraic formulation in that it is coupled to the energy equation and solves an ordinary differential equation (ODE) for the wall modeled velocity profile and wall friction assuming a damped eddy viscosity profile (Eq. A.2), for instance that of Van Driest [73], shown in Eq. A.3. In the present work, a significant speedup is achieved in the time integration of the governing equations by neglecting non-equilibrium effects and by assuming a form of the damping of the eddy viscosity that allows for the EQWM equations to be solved analytically. In this case, the wall model equations reduce to algebraic equations for the wall stresses which can be solved efficiently with a Newton-Raphson root finding algorithm. The algebraic wall model implementation shown in Eq. (A.1), recovers the linearity of the velocity profile in the viscous sublayer and the logarithmic character of the profile off the wall.

$$u^+(y^+) = \begin{cases} y^+ + a_1(y^+)^2 & \text{for } y < y^* \\ \frac{1}{\kappa} ln(y^+) + B & \text{otherwise,} \end{cases} \tag{A.1}$$

In Eq. (A.1), $\kappa = 0.41$, B = 5.2, $y^* \approx 23$, and $a_1 = \frac{1}{2y^*}(\frac{1}{\kappa y^*} - 1)$. The value of a_1 is chosen to enforce the C^1 continuity of the velocity profile.

$$\mu_t = \rho\kappa y\sqrt{\frac{\tau_w}{\rho}}D(y) \tag{A.2}$$

$$D(y) = (1 - exp(-y^+/A^+))^2 \tag{A.3}$$

Considering the momentum equation, simplified from the thin boundary layer equations of [4] by invoking the aforementioned equilibrium assumptions, we arrive at Eq. A.4, which is a two point boundary value problem for the streamwise velocity profile.

$$\frac{d\tau}{dy} = \frac{d}{dy}\left((\mu + \mu_t)\frac{du}{dy}\right) = 0 \tag{A.4}$$

Eq. A.4 is subject to the boundary conditions $u(y^*) = U$ and $u(0) = 0$ where the farfield Dirichlet condition is applied at the LES/wall model exchange location (1st cell in charLES, 3rd cell in Alya) and enforces continuity between the inner and outer layer velocity profiles. We can integrate Eq. A.4 (the form of the eddy viscosity and its damping not being strictly necessary) to arrive at Eq. A.6.

$$\int_0^y \left[\frac{d}{dy}\left((\mu + \mu_t)\frac{du}{dy}\right) = 0\right] \tag{A.5}$$

$$\frac{du}{dy} - \frac{1}{\mu_{tot}}\tau_w = 0 \tag{A.6}$$

Where $\mu_{tot} = \mu + \mu_t$ and A.6 is valid because μ_{tot} is strictly positive for all points in space. We can integrate this equation from the wall to the matching location.

$$U = \tau_w\left(\int_0^{y^*} \frac{dy}{\mu_{tot}(y)}\right) \tag{A.7}$$

Eq. A.7 was derived by making use of the Dirichlet boundary conditions mentioned earlier. If we restrict ourselves to flows without heat flux at the wall, i.e. adiabatic settings where $\frac{dT}{dy}|_w = 0$, and thin boundary layers where the wall-normal momentum equation simplifies to $\frac{dp}{dy} = 0$ then the ideal gas equation of state implies that $\frac{d\rho}{dy}|_w = 0$. If we further admit the approximation that the density and molecular viscosity are constant over the inner layer and use the eddy viscosity profile of Eq. A.2, Eq. A.7 becomes the following:

$$U = \frac{\rho u_\tau^2}{\mu}\int_0^{y^*} \frac{dy}{1 + \frac{\kappa\rho u_\tau y}{\mu}D(y)} \tag{A.8}$$

$$U = u_\tau \int_0^{y^{*,+}} \frac{dy^+}{1 + \kappa y^+ D(y^+)} \equiv u_\tau \Phi(y^{*,+}) \tag{A.9}$$

We can rearrange Eq. A.9 to solve for the velocity profile in viscous units

$$\frac{U(y^*)}{u_\tau} = U^{*+} = \Phi(y^{*,+}) \tag{A.10}$$

From Eq. A.10, we can observe the equivalence of the algebraic model and the equilibrium ODE solution of the thin boundary layer equations. In the former (Eq. A.1), the mean velocity profile ($\Phi(y^+)$) is prescribed, resulting in a nonlinear equation for the wall stress. In the latter, specification of the eddy viscosity implies a velocity profile based on the definition of Φ. An equivalent damping function can be derived for the two models by supposing that ρ and μ are both approximately constant over the region of integration. Eq. A.4 and A.2 combine to yield

$$\frac{d}{dy^+}\left((1+\kappa y^+ D(y^+))\frac{du^+}{dy^+}\right) = 0 \tag{A.11}$$

Integration of Eq. A.11 from 0 to y^+ immediately gives the expression for the equivalent damping function

$$D(y^+) = \frac{1-\frac{du^+}{dy^+}}{\kappa y^+ \frac{du^+}{dy^+}} \tag{A.12}$$

The limiting behavior of Eq. A.12 is that $D(y^+) \to 1$ for $y^+ \gg 1$ if $\frac{du^+}{dy^+} = \frac{1}{\kappa y^+}$. Eq. A.12 gives the damping behavior necessary to recover a particular inner layer velocity profile when solving the equilibrium ODE of [38]. In this light, the algebraic wall model in Eq. A.1 can be regarded as a solution to the equilibrium ODE for a particular choice of damping function for adiabatic, constant property flow.

Finally we show an *a-posteriori* validation of the derivation described above in the context of an LES simulation of a flat plate zero pressure gradient turbulent boundary layer. Figure A.1 shows the velocity profile computed at $Re_\theta = 6500$ in a zero pressure gradient turbulent boundary layer using LES with the two equilibrium wall models of interest (ODE and algebraic) compared with the prediction of a direct numerical simulation (DNS) at the same Reynolds number. Both models agree very well in the viscous sublayer, buffer layer, and logarithmic layer with the DNS and, importantly, give predictions that are consistent with one another. The ODE model slightly outperforms the algebraic model in the buffer layer. We use this test case to justify our choice of the algebraic model, noting that differences between the models are negligible in the flow regime that we are simulating. We choose the algebraic model over the ODE model due to its $\approx 5\times$ cheaper cost, owing to the simplification achieved when we solve a nonlinear algebraic equation for the wall stress instead of an ODE.

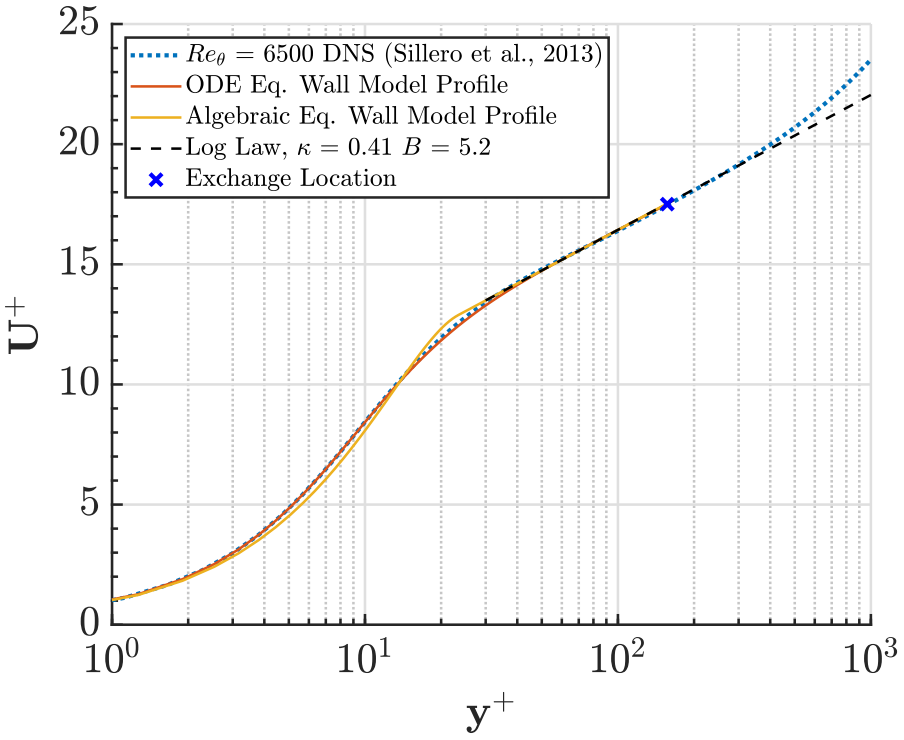

Figure A.1: Semi-log plot of the near-wall velocity profile in viscous units. Dotted lines are from the DNS of [66] while solid lines are from LES calculations at $Re_\theta = 6500$, $Mach = 0.2$ of a turbulent flat plate boundary layer using charLES with two equilibrium wall modeling approaches. The exchange location is in the logarithmic layer and there is excellent consistency between the wall models and the DNS.

Appendix B

Sectional Pressure Coefficient Measurements for the High-Lift CRM

Appendix B contains sectional pressure measurements at all available experimental locations at the corrected angles of attack of 2.78°, 7.05°, 11.29°, 17.05°. 19.57°. 20.55°, and 21.46° from the best achieved LES simulation of the High-Lift Common Research Model. In this case, that calculation was performed on a 1.5 Bcv grid that used isotropic hexagonal close packed elements throughout the domain and a first cell centroid LES/wall model exchange height. These calculations were carried out in a free air setting.

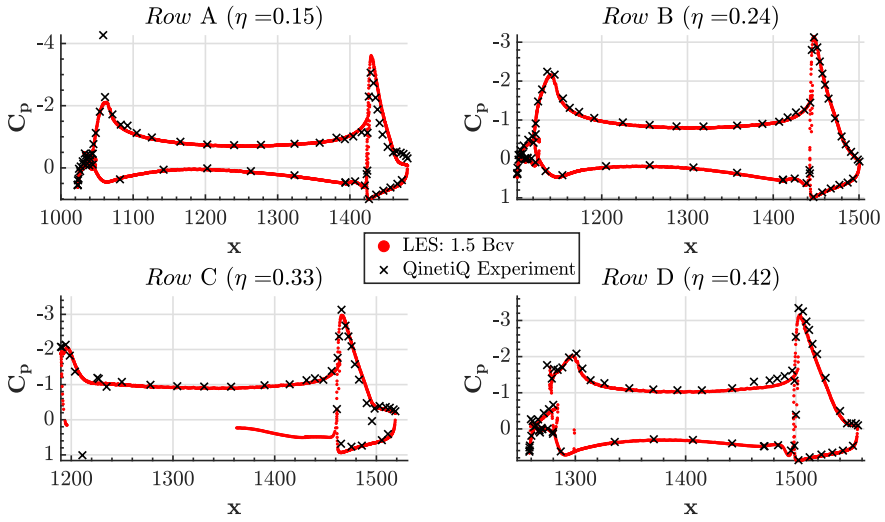

Figure B.1: Average sectional pressure measurements at 4 inboard spanwise locations from best-practice LES calculations compared against experimental data from the QinetiQ Facility for the flow over the High-Lift Common Research Model. The bulk flow is at M=0.20, Re=5.49e6 (based on mean aerodynamic chord), and $\alpha = 2.78°$.

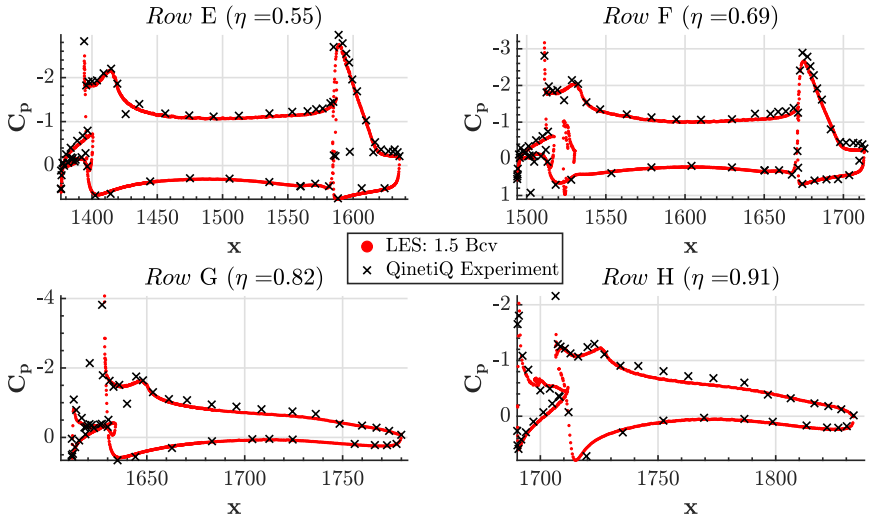

Figure B.2: Average sectional pressure measurements at 4 outboard spanwise locations from best-practice LES calculations compared against experimental data from the QinetiQ Facility for the flow over the High-Lift Common Research Model. The bulk flow is at M=0.20, Re=5.49e6 (based on mean aerodynamic chord), and $\alpha = 2.78°$.

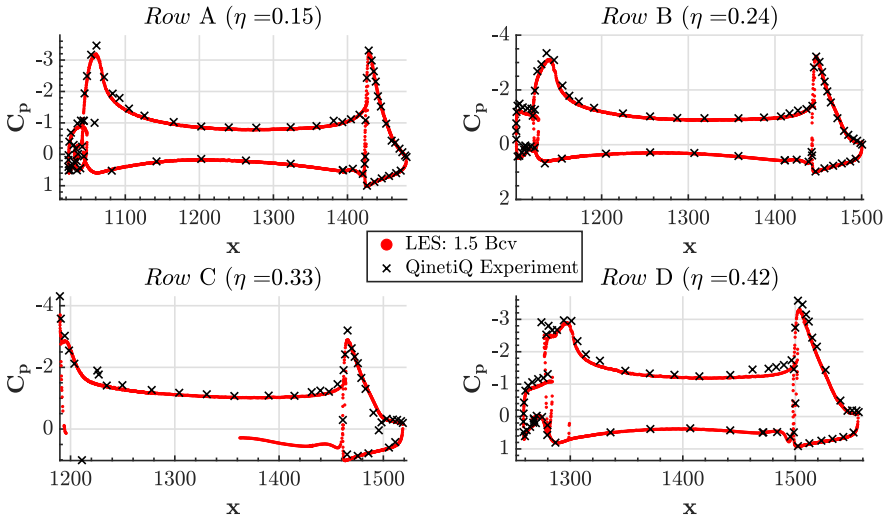

Figure B.3: Average sectional pressure measurements at 4 inboard spanwise locations from best-practice LES calculations compared against experimental data from the QinetiQ Facility for the flow over the High-Lift Common Research Model. The bulk flow is at M=0.20, Re=5.49e6 (based on mean aerodynamic chord), and $\alpha = 7.05°$.

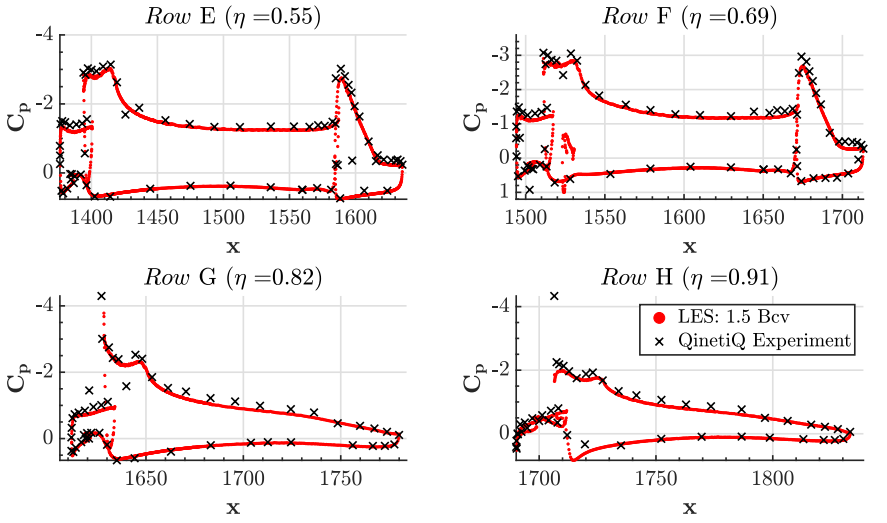

Figure B.4: Average sectional pressure measurements at 4 outboard spanwise locations from best-practice LES calculations compared against experimental data from the QinetiQ Facility for the flow over the High-Lift Common Research Model. The bulk flow is at M=0.20, Re=5.49e6 (based on mean aerodynamic chord), and $\alpha = 7.05°$.

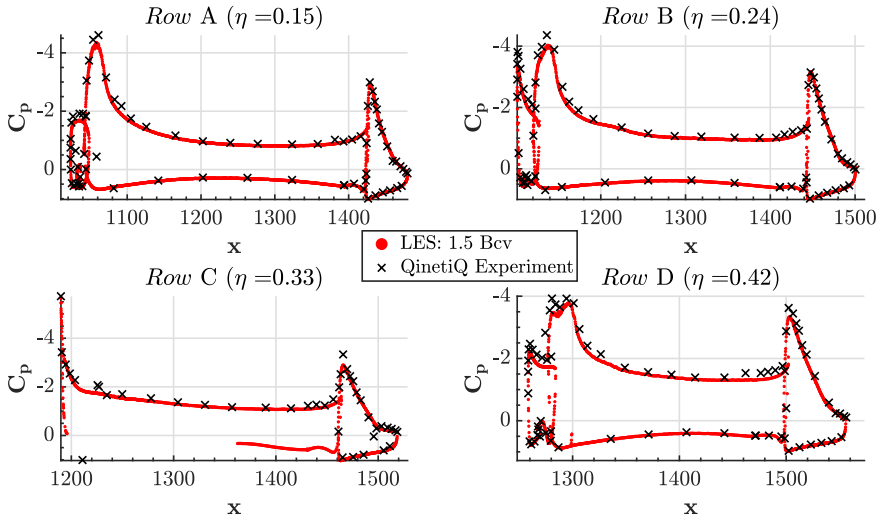

Figure B.5: Average sectional pressure measurements at 4 inboard spanwise locations from best-practice LES calculations compared against experimental data from the QinetiQ Facility for the flow over the High-Lift Common Research Model. The bulk flow is at M=0.20, Re=5.49e6 (based on mean aerodynamic chord), and $\alpha = 11.29°$.

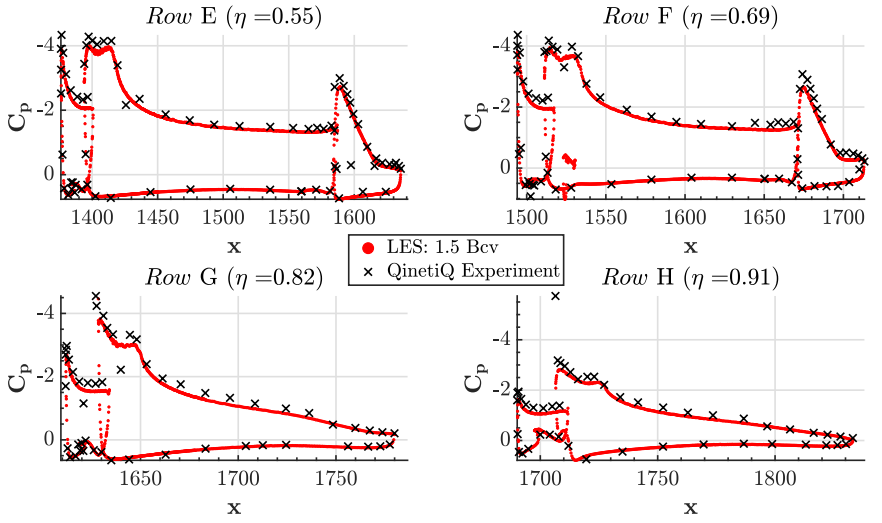

Figure B.6: Average sectional pressure measurements at 4 outboard spanwise locations from best-practice LES calculations compared against experimental data from the QinetiQ Facility for the flow over the High-Lift Common Research Model. The bulk flow is at M=0.20, Re=5.49e6 (based on mean aerodynamic chord), and $\alpha = 11.29°$.

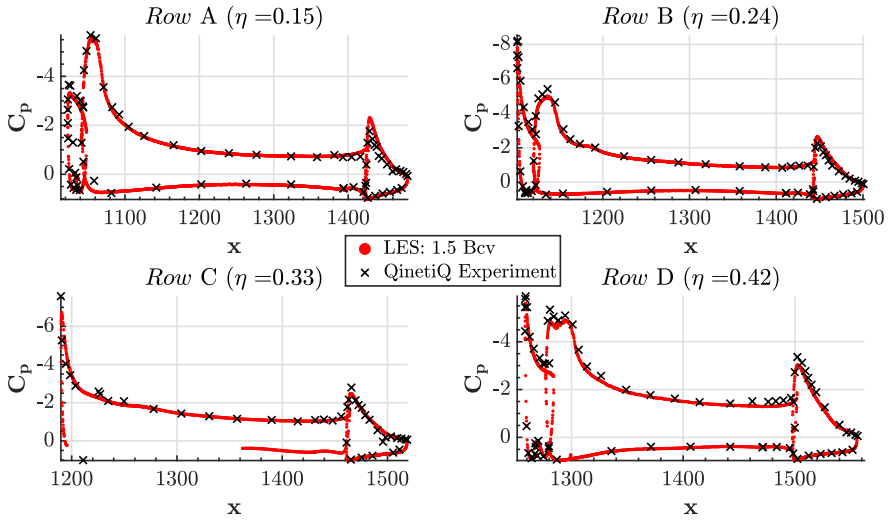

Figure B.7: Average sectional pressure measurements at 4 inboard spanwise locations from best-practice LES calculations compared against experimental data from the QinetiQ Facility for the flow over the High-Lift Common Research Model. The bulk flow is at M=0.20, Re=5.49e6 (based on mean aerodynamic chord), and $\alpha = 17.05°$.

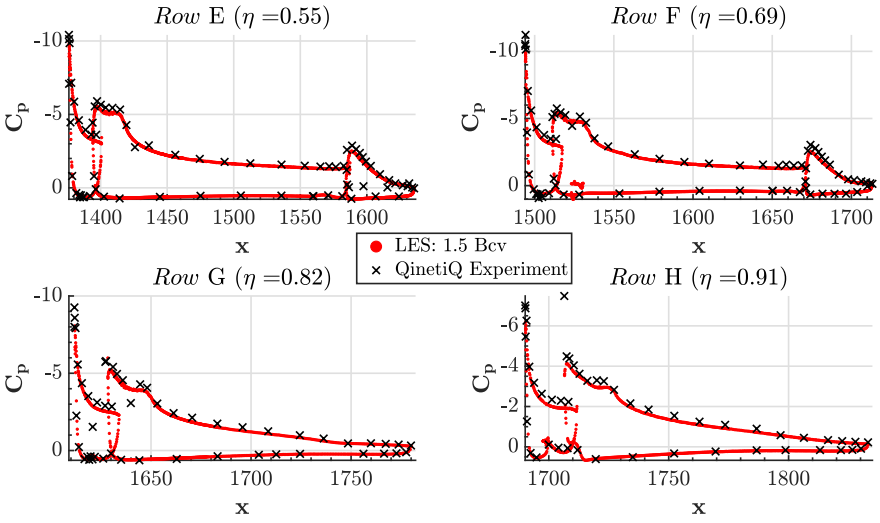

Figure B.8: Average sectional pressure measurements at 4 outboard spanwise locations from best-practice LES calculations compared against experimental data from the QinetiQ Facility for the flow over the High-Lift Common Research Model. The bulk flow is at M=0.20, Re=5.49e6 (based on mean aerodynamic chord), and $\alpha = 17.05°$.

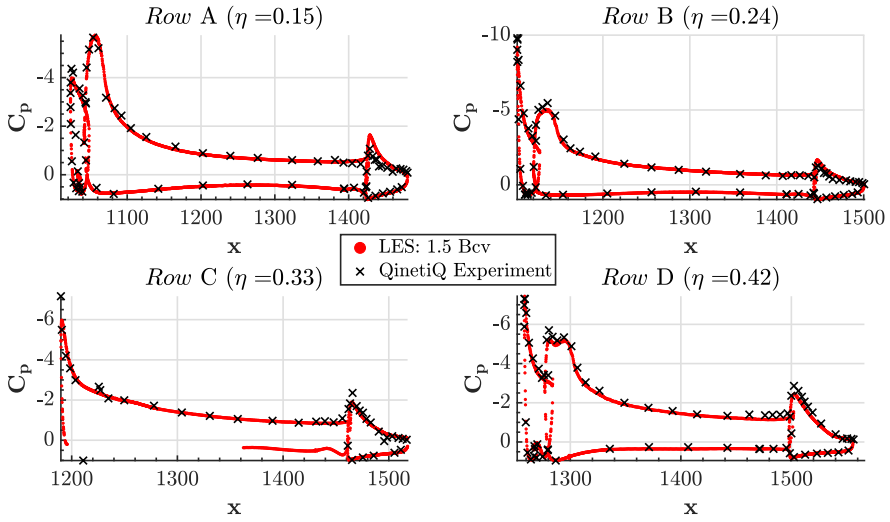

Figure B.9: Average sectional pressure measurements at 4 inboard spanwise locations from best-practice LES calculations compared against experimental data from the QinetiQ Facility for the flow over the High-Lift Common Research Model. The bulk flow is at M=0.20, Re=5.49e6 (based on mean aerodynamic chord), and $\alpha = 19.57°$.

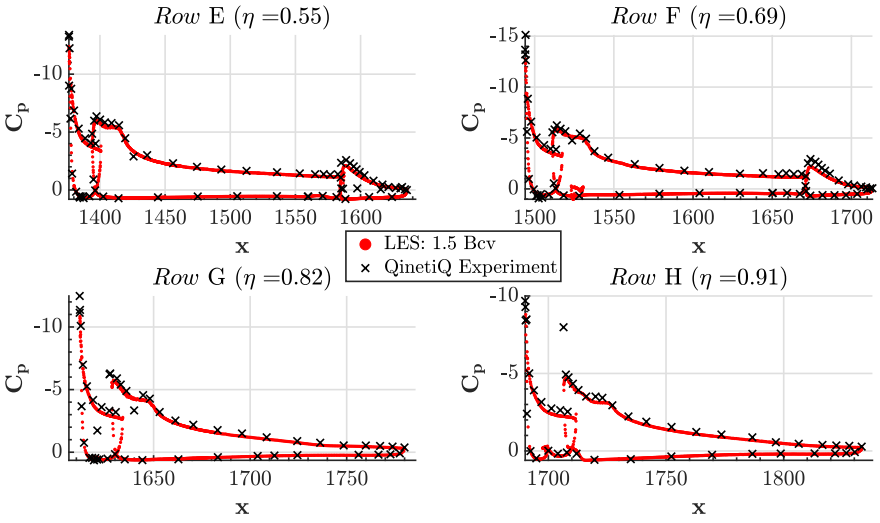

Figure B.10: Average sectional pressure measurements at 4 outboard spanwise locations from best-practice LES calculations compared against experimental data from the QinetiQ Facility for the flow over the High-Lift Common Research Model. The bulk flow is at M=0.20, Re=5.49e6 (based on mean aerodynamic chord), and $\alpha = 19.57°$.

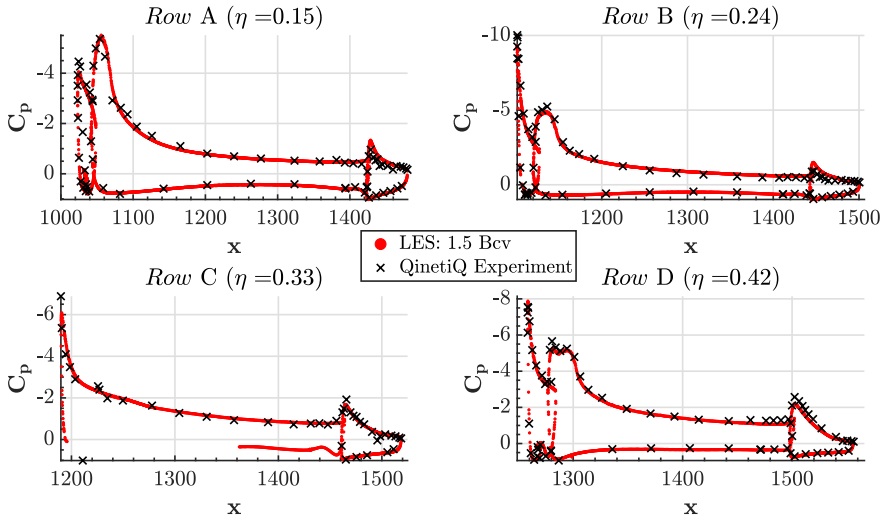

Figure B.11: Average sectional pressure measurements at 4 inboard spanwise locations from best-practice LES calculations compared against experimental data from the QinetiQ Facility for the flow over the High-Lift Common Research Model. The bulk flow is at M=0.20, Re=5.49e6 (based on mean aerodynamic chord), and $\alpha = 20.55°$.

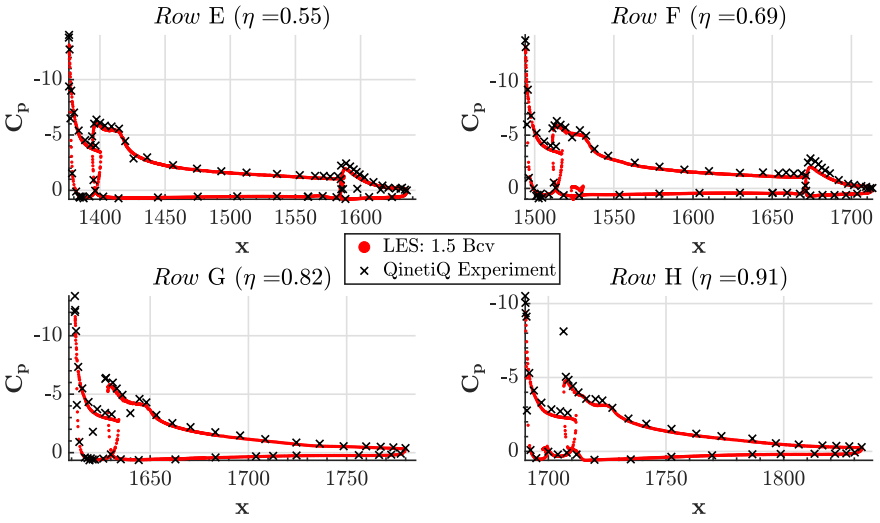

Figure B.12: Average sectional pressure measurements at 4 outboard spanwise locations from best-practice LES calculations compared against experimental data from the QinetiQ Facility for the flow over the High-Lift Common Research Model. The bulk flow is at M=0.20, Re=5.49e6 (based on mean aerodynamic chord), and $\alpha = 20.55°$.

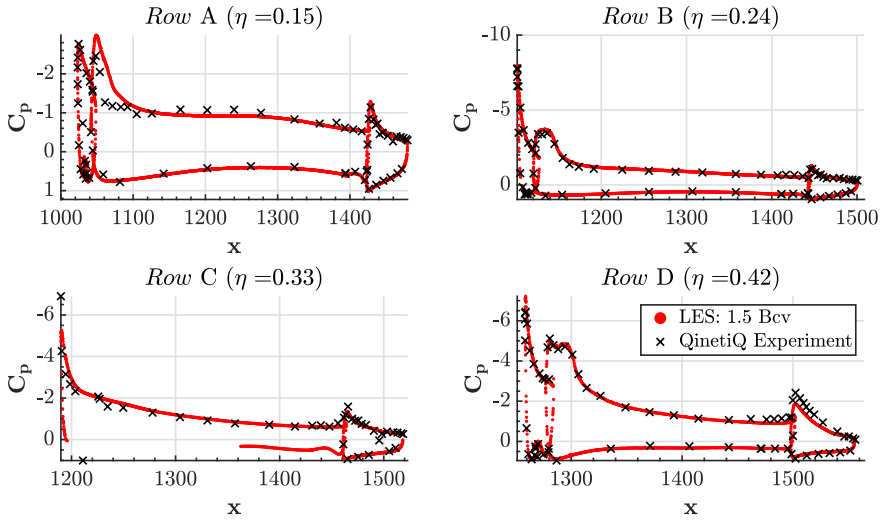

Figure B.13: Average sectional pressure measurements at 4 inboard spanwise locations from best-practice LES calculations compared against experimental data from the QinetiQ Facility for the flow over the High-Lift Common Research Model. The bulk flow is at M=0.20, Re=5.49e6 (based on mean aerodynamic chord), and $\alpha = 21.46°$.

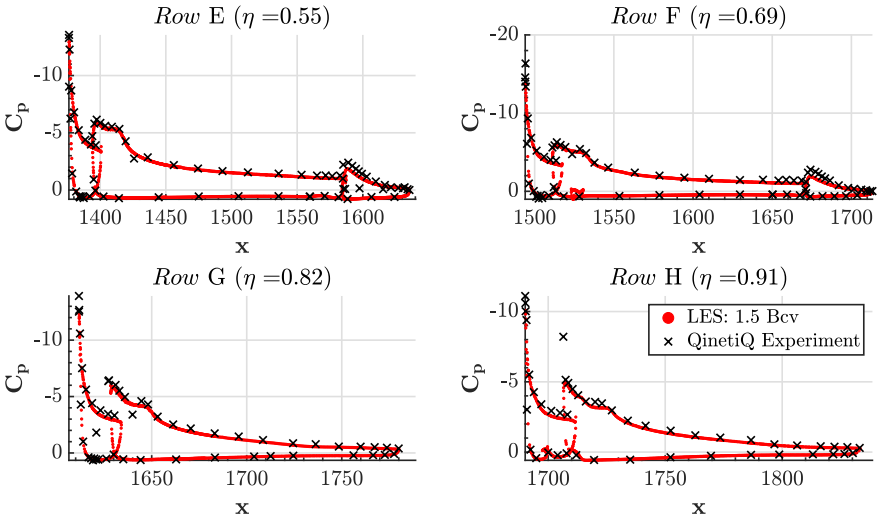

Figure B.14: Average sectional pressure measurements at 4 outboard spanwise locations from best-practice LES calculations compared against experimental data from the QinetiQ Facility for the flow over the High-Lift Common Research Model. The bulk flow is at M=0.20, Re=5.49e6 (based on mean aerodynamic chord), and $\alpha = 21.46°$.

Appendix C

Sectional Pressure Coefficient Measurements for the Transonic CRM

Appendix C contains sectional pressure measurements at all available experimental locations at the angles of attack of $2.50^\circ - 4.00^\circ$ spaced at 0.25° increments from the best achieved LES simulation of the transonic Common Research Model. In this case, that calculation was a $\approx$ 270 Mcv grid that used prismatic boundary layer elements with maximum anisotropy of 4:1 with a wall-normal stretching ratio of 1.15 and a second cell centroid LES/wall model exchange height.

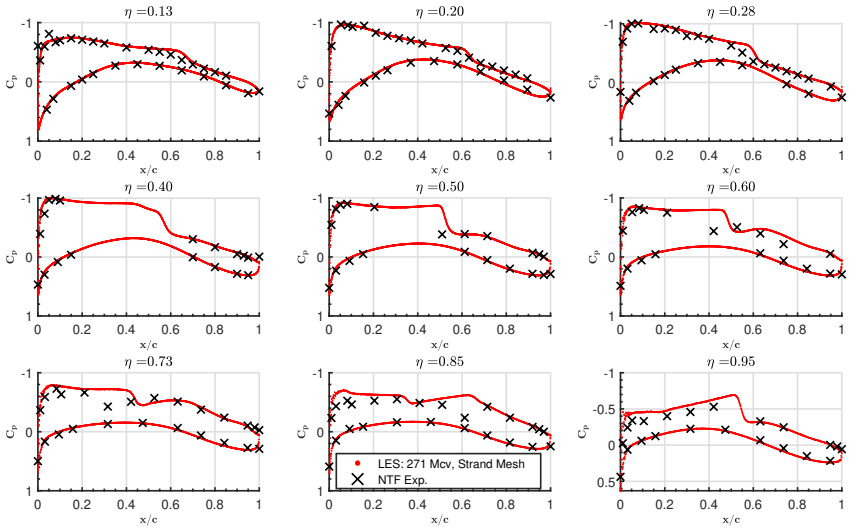

Figure C.1: Average sectional pressure measurements at 9 spanwise locations from best-practice LES calculations compared against experimental data from the National Transonic Facility for the flow over the tranonic Common Research Model. The bulk flow is at M=0.85, Re=5e6 (based on mean aerodynamic chord), and $\alpha = 2.50°$.

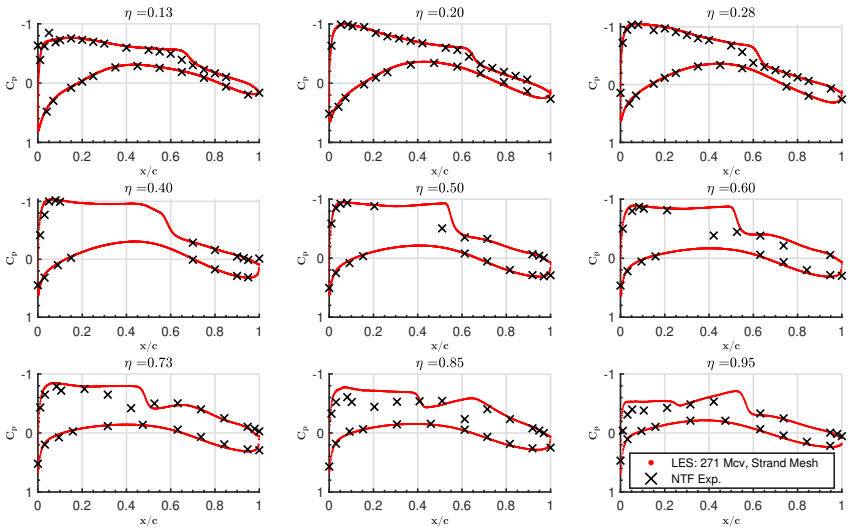

Figure C.2: Average sectional pressure measurements at 9 spanwise locations from best-practice LES calculations compared against experimental data from the National Transonic Facility for the flow over the tranonic Common Research Model. The bulk flow is at M=0.85, Re=5e6 (based on mean aerodynamic chord), and $\alpha = 2.75°$.

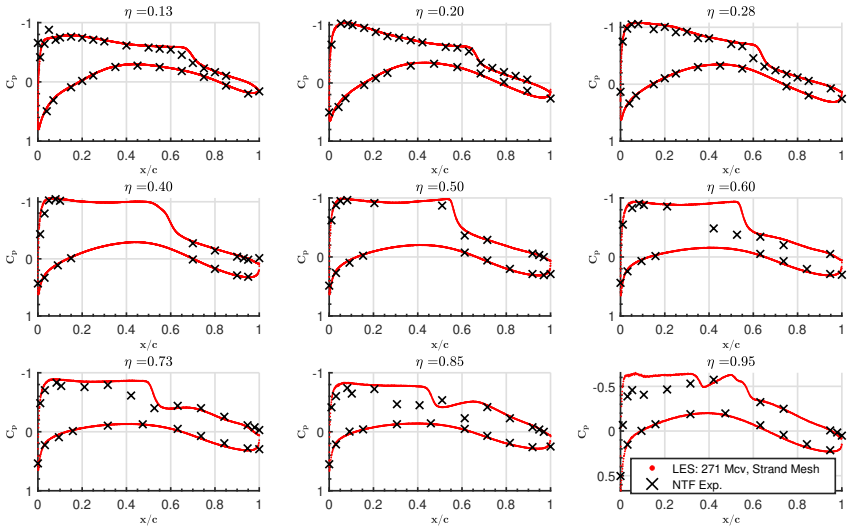

Figure C.3: Average sectional pressure measurements at 9 spanwise locations from best-practice LES calculations compared against experimental data from the National Transonic Facility for the flow over the tranonic Common Research Model. The bulk flow is at M=0.85, Re=5e6 (based on mean aerodynamic chord), and $\alpha = 3.00°$.

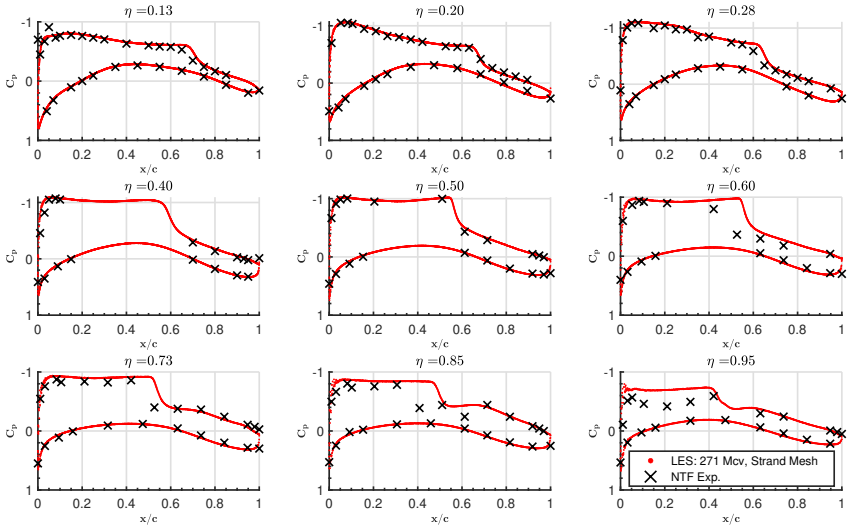

Figure C.4: Average sectional pressure measurements at 9 spanwise locations from best-practice LES calculations compared against experimental data from the National Transonic Facility for the flow over the tranonic Common Research Model. The bulk flow is at M=0.85, Re=5e6 (based on mean aerodynamic chord), and $\alpha = 3.25°$.

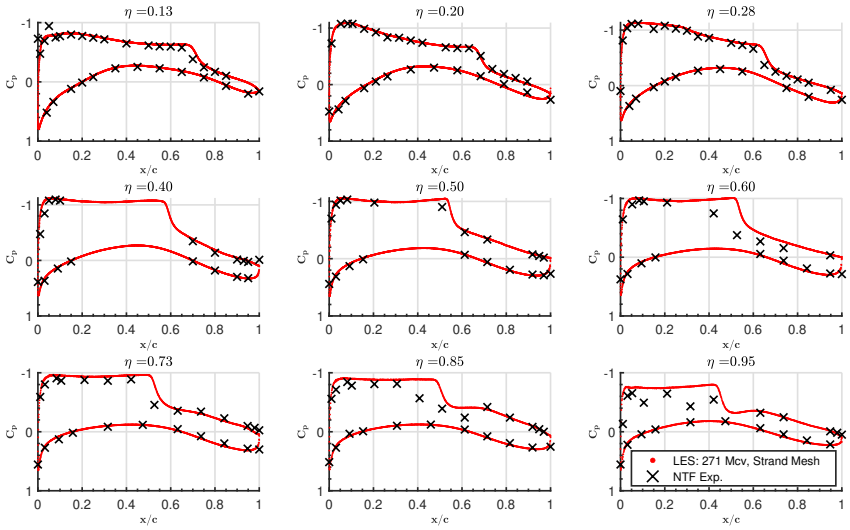

Figure C.5: Average sectional pressure measurements at 9 spanwise locations from best-practice LES calculations compared against experimental data from the National Transonic Facility for the flow over the tranonic Common Research Model. The bulk flow is at M=0.85, Re=5e6 (based on mean aerodynamic chord), and $\alpha = 3.50°$.

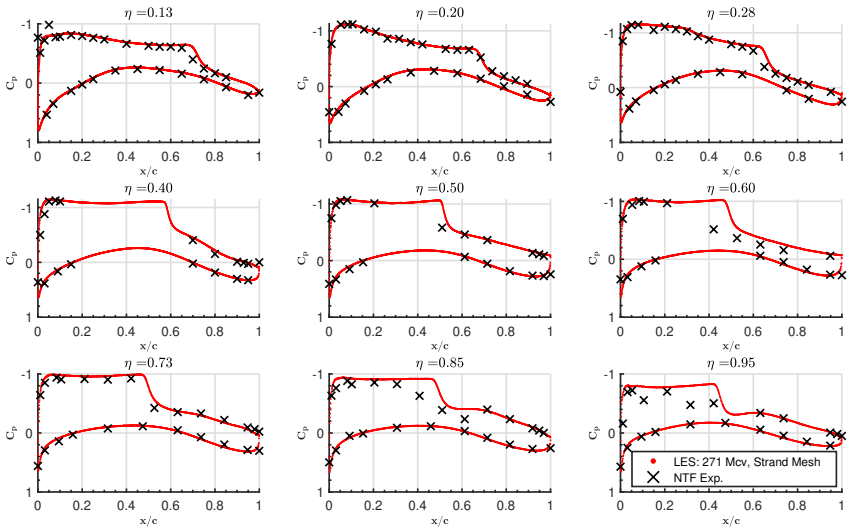

Figure C.6: Average sectional pressure measurements at 9 spanwise locations from best-practice LES calculations compared against experimental data from the National Transonic Facility for the flow over the tranonic Common Research Model. The bulk flow is at M=0.85, Re=5e6 (based on mean aerodynamic chord), and $\alpha = 3.75°$.

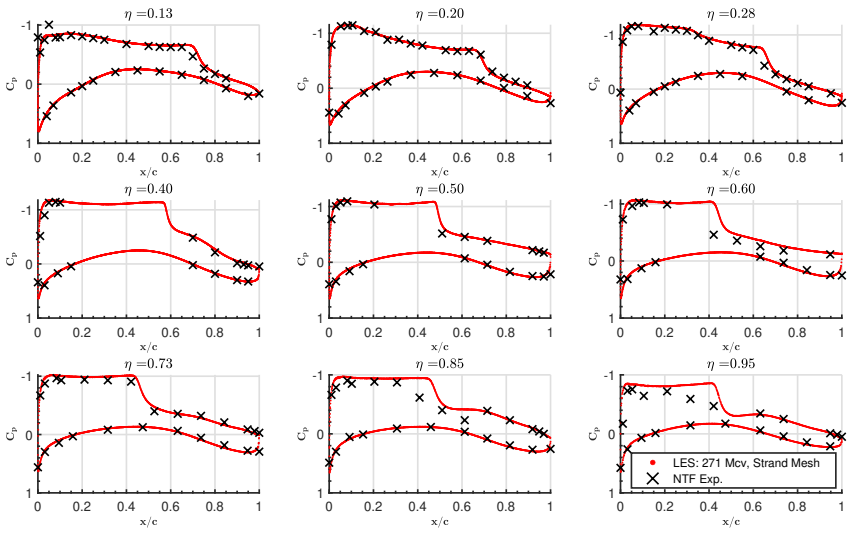

Figure C.7: Average sectional pressure measurements at 9 spanwise locations from best-practice LES calculations compared against experimental data from the National Transonic Facility for the flow over the tranonic Common Research Model. The bulk flow is at M=0.85, Re=5e6 (based on mean aerodynamic chord), and $\alpha = 4.00°$.

Appendix D

Validations of LES in canonical flow environments

This appendix describes validations performed using LES in canonical flow settings, including bump flows such as the NASA Wall-Mounted Hump as well as a simplified airfoil section, the MD30P30N. For recent studies of LES in other canonical flow settings, including smooth body separations, and shock-induced separations, the reader is referred to: [2, 1, 78, 21, 23].

D.1 NASA Wall-Mounted Hump

The NASA Wall-Mounted Hump is a well-studied canonical low-Mach bump flow which admits mild-to-strong favorable/adverse pressure gradients and leads to a geometrically-imposed separation on the leeward side of the bump. Of interest in this case is the prediction of separation/reattachment locations (the latter being more challenging owing to strong geometric curvature at the separation point) and surface pressure/skin friction measurements. The reference data are from [31]. Figure D.1 shows the surface pressure prediction of LES from a few different sources, including the present Algebraic EQWM investigation as well as the EQWM/NEQWM simulations of [57] and the Wall-Resolved LES (WRLES) calculations of [72]. It is clear that the surface pressure prediction for this case is largely an inviscid phenomenon and is relatively well predicted by all approaches. Figure D.2 shows the skin friction prediction from the same simulations. Here, we see differences between the calculations. For instance, WRLES is the only approach capable of predicting the re-laminarization on the windward-facing side of the bump ($x/c \approx 0 - 0.3$), which WMLES (owing to the built-in assumption within the wall-modeled equations that the boundary layer is turbulent) is unable to predict. Despite the differences in skin friction behavior ahead of the separation, all approaches are able to predict an accurate separation and reattachment location, with the EQWM result of [57] showing too weak a reversal of the skin friction within the separation bubble. This result

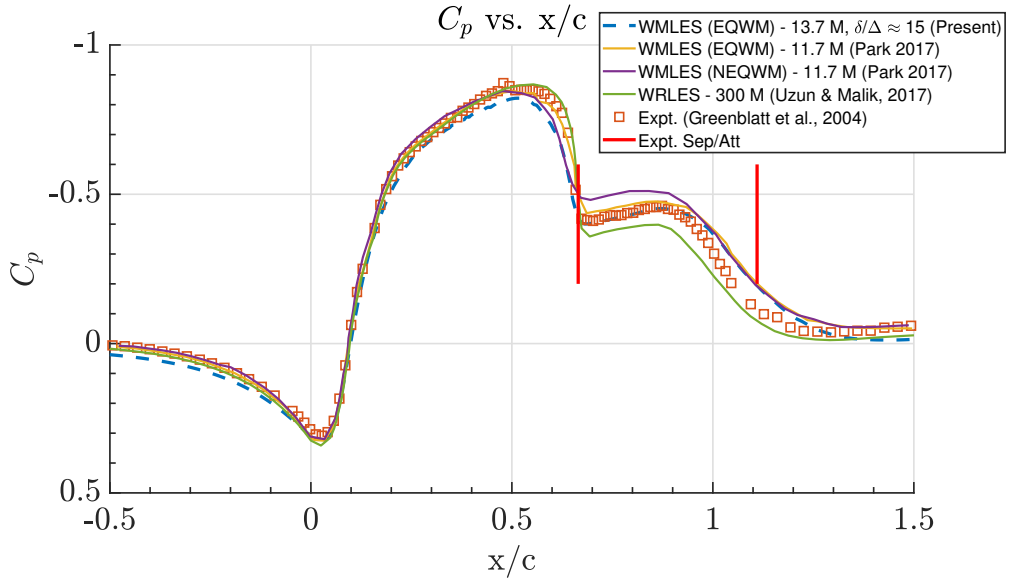

Figure D.1: Average surface pressure measurements along the centerline of the NASA Wall-Mounted Hump canonical bump flow. The bulk flow is at M=0.10, Re=936,000 (based on the bump chord). Experimental data are from [31].

indicates that the flow "loses memory" of its history upstream of the separation as each approach predicts an accurate separation/reattachment location even if errors in the attached region were present. Further details of the flow state were investigated by means of vertical velocity probes. Figure D.3 shows the locations of these probes just ahead of and within, the separation bubble. The present EQWM calculations are shown in Figures D.4 and D.5 to make reasonable estimates of the streamwise and vertical velocity components within the separation region, further lending confidence to this modeling approach for separated flows. Overall, this benchmark case was meant to build confidence in the predictive capability of WMLES in a simple separated flow environment ahead of its deployment on high-lift separated aircraft flows. It is clear that the present EQWM approach is on par with previous investigations of LES for this flow such as that of [57] and is able to accurately predict surface pressures and separation/reattachment locations of flows experiencing mild-to-strong favorable/adverse pressure gradients, both of which are key quantities of interest tied to the prediction of lift/drag on separated aircraft flows. A key feature that was shown to be outside of the predictive scope of WMLES is that of flow re-laminarization.

D.2 The McDonnell Douglas 30P30N

The MD30P30N is an important validation case for high lift flows. It consists of a three-element airfoil with a leading edge slat, main element, and trailing edge flap. The slat and flap are deployed

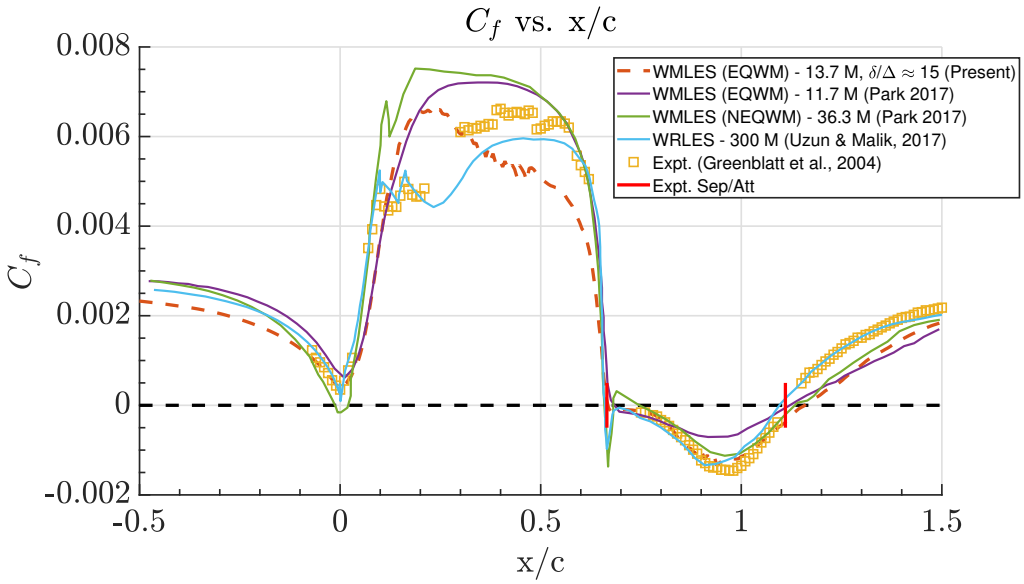

Figure D.2: Average surface skin friction measurements along the centerline of the NASA Wall-Mounted Hump canonical bump flow. The bulk flow is at M=0.10, Re=936,000 (based on the bump chord). Experimental data are from [31].

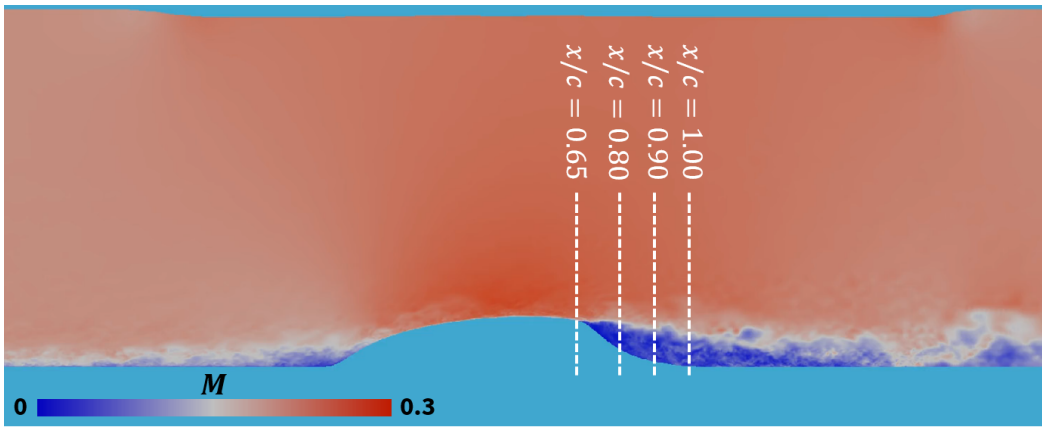

Figure D.3: Center plane slice of instantaneous Mach number in the domain of the NASA Wall-Mounted Hump simulation. Included for reference are the locations within the separation bubble of the vertical velocity profiles.

in order to be representative of a commercial aircraft high-lift system for landing configuration [14]. This case was previously simulated using WMLES at CTR by [6, 7] with a transition sensor used to identify regions of laminar flow. The authors also explored the use of no-slip boundary conditions (as opposed to wall-modeled) in separated regions of the flow (such as the flap cove). The cases in this report are meant to be representative of the full aircraft simulations performed on the JSM and

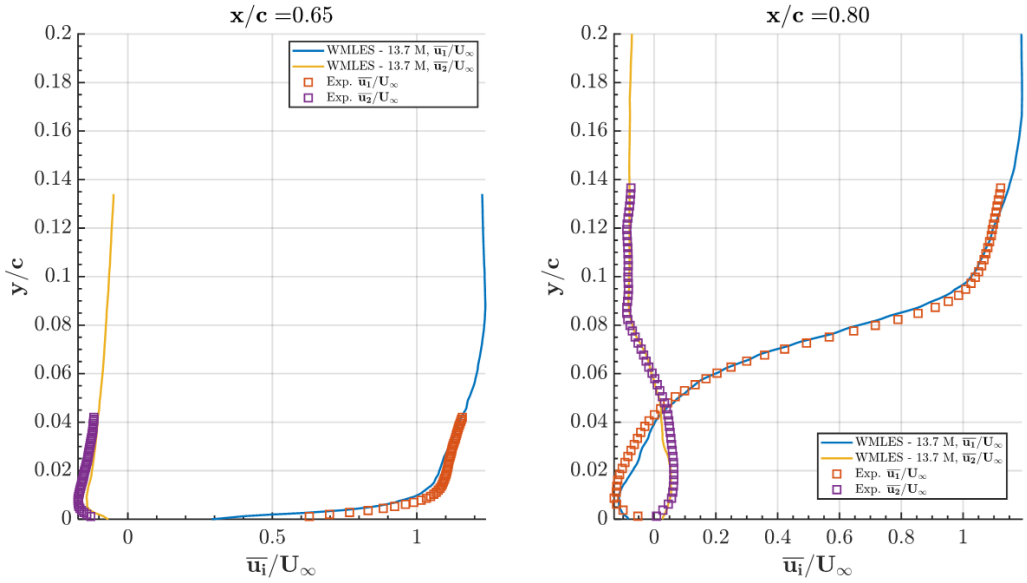

Figure D.4: Streamwise and vertical velocity profiles at the start of the separation bubble ($x/c = 0.65, 0.80$). Experimental data are for [31].

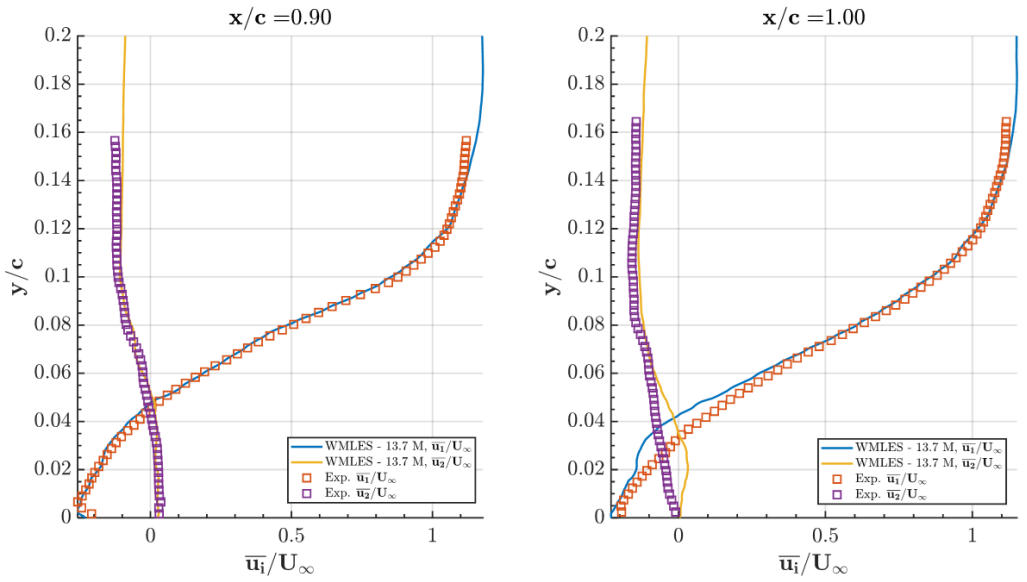

Figure D.5: Streamwise and vertical velocity profiles within the separation bubble ($x/c = 0.90, 1.00$). Experimental data are for [31].

CRM-HL cases and so use identical gridding and modeling practices (EQWM is active everywhere with DSM SGS model). The naming convention of the grid resolutions presented in Figures D.6, D.7,

and D.8 corresponds exactly to the grid resolution levels used in the CRM-HL cases. A one-chord-wide spanwise extent is used to allow a sufficient extent for de-correlation of turbulent structures in the wake. The cases are run using a periodic boundary treatment in the spanwise direction. It is clear from the sectional pressure traces in Figures D.6 and D.8 that strong grid sensitivity over the resolution levels considered here is not observed for either the low or high angle of attack cases. This is in contrast to the full aircraft configurations, which did show grid sensitivities, particularly on the flap at low angles of attack and at the slat suction peak at high angles of attack. This lack of grid sensitivity at the pertinent flow conditions is a deficiency of this simplified two-dimensional airfoil section compared to the full aircraft configuration. A key difference visible in the pressure plots is the strength of the leading edge suction at high angles of attack, which is more than $5\times$ stronger at 21° compared with 8°. As mentioned before, this phenomenon dominated the grid counts at high angle of attack for the full aircraft simulations, but was not sensitive to grid here. Figure D.7 shows that the average skin friction magnitude was sensitive to grid resolution across the refinement sweep. As the grid is refined, additional grid support for the transition mechanism is achieved and the transition moves towards the leading edge of the wing. Although experimental skin friction data is not available, [14] quotes that at the $9M$ Reynolds number simulated here, transition is expected within the first $5 - 10\%$ of the leading edge of each wing element. From this perspective, the XXFine calculation behavior is most consistent with experiment, though the skin friction prediction of LES is clearly not grid-converged for this case and poses stricter resolution requirements than the surface pressure, which was gird-converged even for the Medium/Fine meshes. This validation case has highlighted that while the full aircraft simulation practices are appropriate enough to capture average surface pressure distributions in a simplified flow, a two-dimensional airfoil section is inadequate to study the full complexity of an aircraft flow, including the challenges experienced in the CRM-HL/JSM cases such as flap separation collapse and wingtip/leading edge flow acceleration grid sensitivities. Furthermore, it is clear that grid convergence of LES solutions depends on the quantity of interest.

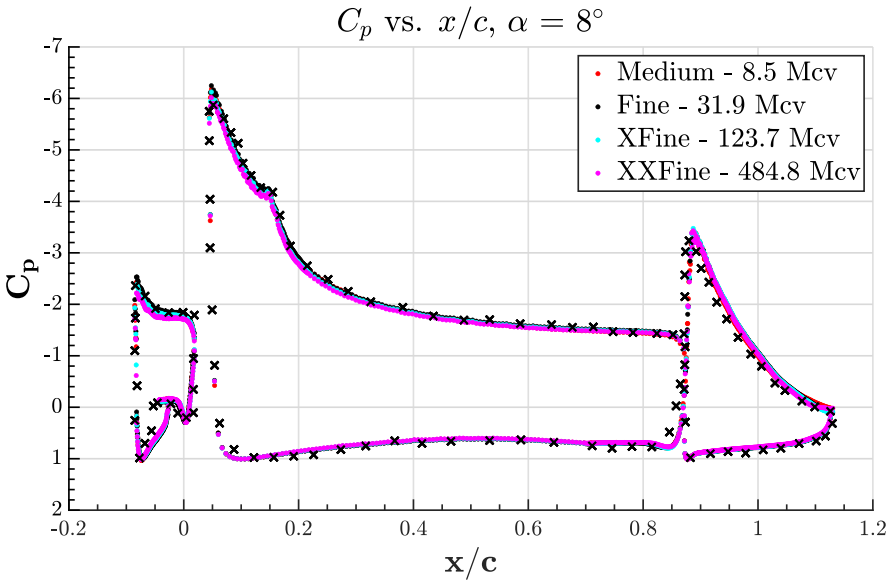

Figure D.6: Average sectional pressure measurements along the centerline of the MD30P30N high-lift validation case at four grid resolutions. Experimental data are plotted with a black x. The bulk flow is at M=0.20, Re=9.0×10^6 (based on mean aerodynamic chord), and $\alpha = 8°$.

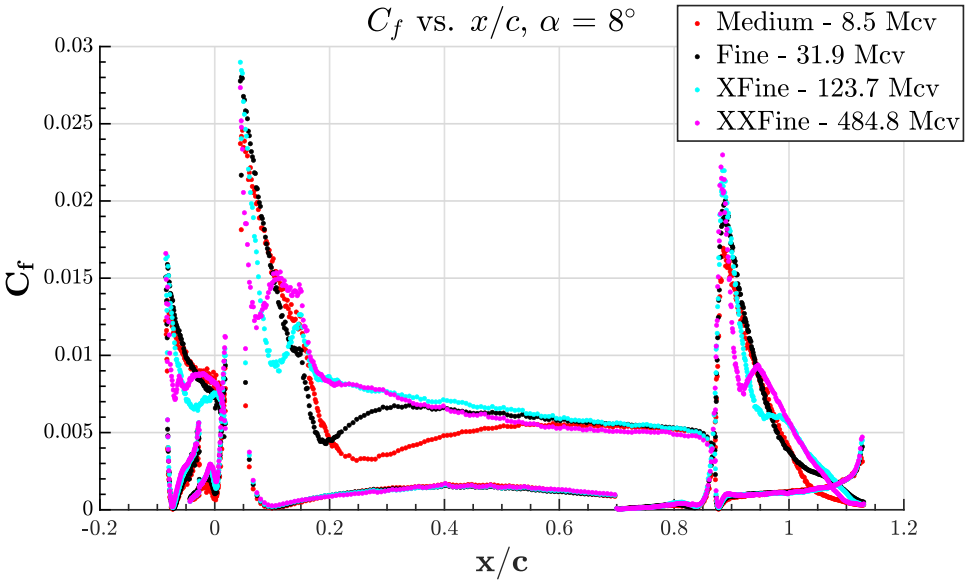

Figure D.7: Average surface skin friction magnitude measurements along the centerline of the MD30P30N high-lift validation case at four grid resolutions. The bulk flow is at M=0.20, Re=9.0×10^6 (based on mean aerodynamic chord), and $\alpha = 8°$.

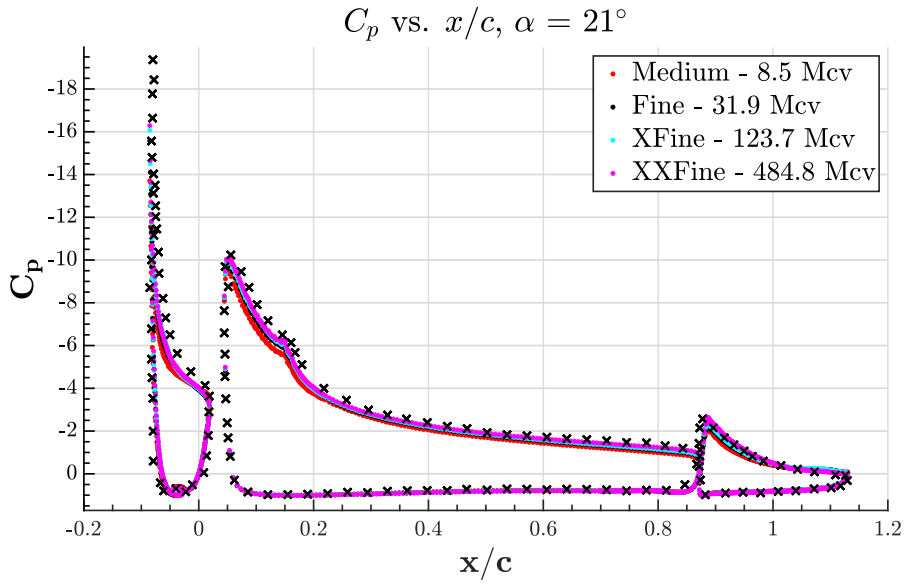

Figure D.8: Average sectional pressure measurements along the centerline of the MD30P30N high-lift validation case at four grid resolutions. Experimental data are plotted with a black x. The bulk flow is at M=0.20, Re=9.0×10^6 (based on mean aerodynamic chord), and $\alpha = 21°$.

Bibliography

[1] Rahul Agrawal, Sanjeeb T. Bose, and Parviz Moin. Wall modeled LES of the Boeing speed bump using a non-Boussinesq modeling framework. *Annual Research Briefs*, Center for Turbulence Research, Stanford University. pages 43–58, 2022.

[2] Rahul Agrawal, Michael P Whitmore, Kevin P Griffin, Sanjeeb T Bose, and Parviz Moin. Non-boussinesq subgrid-scale model with dynamic tensorial coefficients. *Physical Review Fluids*, 7(7):074602, 2022.

[3] Neil Ashton, Marie Denison, and Marian Zastawny. 3rd High-Lift Workshop summary paper-OpenFOAM, STAR-CCM+ & LAVA simulations on unstructured grids. In *2018 AIAA Aerospace Sciences Meeting*, page 1253, 2018.

[4] Elias Balaras, Carlo Benocci, and Ugo Piomelli. Two-layer approximate boundary conditions for large-eddy simulations. *AIAA Journal*, 34(6):1111–1119, 1996.

[5] Riccardo Balin and Kenneth E Jansen. A comparison of rans, urans, and ddes for high lift systems from hiliftpw-3. In *2018 AIAA Aerospace Sciences Meeting*, page 1254, 2018.

[6] Julien Bodart and Johan Larsson. Wall-modeled large eddy simulation of the McDonnell-Douglas 30P/30N high-lift airfoil in near-stall conditions. In *30th AIAA Applied Aerodynamics Conference*, page 3022, 2012.

[7] Julien Bodart, Johan Larsson, and Parviz Moin. Large eddy simulation of high-lift devices. In *21st AIAA Computational Fluid Dynamics Conference*, page 2724, 2013.

[8] Sanjeeb T Bose and George Ilhwan Park. Wall-modeled large-eddy simulation for complex turbulent flows. *Annual Review of Fluid Mechanics*, 50:535–561, 2018.

[9] Guillaume A Bres, Sanjeeb T Bose, Michael Emory, Frank E Ham, Oliver T Schmidt, Georgios Rigas, and Tim Colonius. Large-eddy simulations of co-annular turbulent jet using a Voronoi-based mesh generation framework. In *2018 AIAA/CEAS Aeroacoustics Conference*, page 3302, 2018.

[10] AE Bryson and RWF Gross. Diffraction of strong shocks by cones, cylinders, and spheres. *Journal of Fluid Mechanics*, 10(1):1–16, 1961.

[11] W Cabot and Parviz Moin. Approximate wall boundary conditions in the large-eddy simulation of high reynolds number flow. *Flow, Turbulence and Combustion*, 63(1):269–291, 2000.

[12] Andrew W Cary, Mohamed Yousuf, Pei Li, and Mortaza Mani. Current practice unstructured grid CFD results for 3rd AIAA High Lift Prediction Workshop. In *2018 AIAA Aerospace Sciences Meeting*, page 1037, 2018.

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

[14] Vincent Chin, David Peters, Frank Spaid, and Robert McGhee. Flowfield measurements about a multi-element airfoil at high Reynolds numbers. *AIAA Paper 1993-3137.*, 1993.

[15] Haecheon Choi and Parviz Moin. Effects of the computational time step on numerical solutions of turbulent flow. *Journal of Computational Physics*, 113(1):1–4, 1994.

[16] Haecheon Choi and Parviz Moin. Grid-point requirements for large eddy simulation: Chapman's estimates revisited. *Physics of Fluids*, 24(1):011702, 2012.

[17] Adam M Clark, Jeffrey P Slotnick, Nigel J Taylor, and Christopher L Rumsey. Requirements and challenges for CFD validation within the High-Lift Common Research Model ecosystem. In *AIAA AVIATION 2020 FORUM*, page 2772, 2020.

[18] James G Coder, Thomas H Pulliam, and James C Jensen. High-lift simulations of the japan aerospace exploration agency standard model. *Journal of Aircraft*, 56(5):1822–1832, 2019.

[19] JD Crouch, A Garbaruk, and M Strelets. Global instability in the onset of transonic-wing buffet. *Journal of Fluid Mechanics*, 881:3–22, 2019.

[20] Qiang Du, Maria Emelianenko, and Lili Ju. Convergence of the Lloyd algorithm for computing centroidal Voronoi tessellations. *SIAM J. Numer. Anal.*, 44:102–119, 2006.

[21] A. Elnahhas, R. Agrawal, and P. Moin. Wall-modeled large-eddy simulation of the Sandia Axisymmetric Transonic Bump. *Annual Research Briefs*, Center for Turbulence Research, Stanford University. pages 31–42, 2022.

[22] Ashley N Evans, Doug S Lacy, Ian Smith, and Melissa B Rivers. Test summary of the NASA high-lift common research model half-span at QinetiQ 5-metre pressurized low-speed wind tunnel. *AIAA Paper 2020-2770.*, 2020.

[23] Lin Fu, Michael Karp, Sanjeeb T Bose, Parviz Moin, and Javier Urzay. Shock-induced heating and transition to turbulence in a hypersonic boundary layer. *Journal of Fluid Mechanics*, 909, 2021.

[24] Massimo Germano, Ugo Piomelli, Parviz Moin, and William H Cabot. A dynamic subgrid-scale eddy viscosity model. *Phys. of Fluids A-Fluid Dynamics*, 3:1760–1765, 1991.

[25] Aditya S Ghate, Gaetan K Kenway, Gerrit-Daniel Stich, Oliver M Browne, Jeffrey A Housman, and Cetin C Kiris. Transonic lift and drag predictions using wall-modelled large eddy simulations. In *AIAA Scitech 2021 Forum*, page 1439, 2021.

[26] Konrad Goc, Sanjeeb Bose, and Parviz Moin. Wall-modeled large-eddy simulation of aircraft in landing configuration. *Annual Research Briefs*, Center for Turbulence Research, Stanford University,. pages 3–14, 2019.

[27] Konrad Goc, Sanjeeb Bose, and Parviz Moin. Subgrid-scale modeling sensitivities in wall-modeled large-eddy simulations of a high-lift aircraft configuration. *Annual Research Briefs*, Center for Turbulence Research, Stanford University,. pages 49–58, 2020.

[28] Konrad Goc, Sanjeeb Bose, and Parviz Moin. Wall-modeled large eddy simulation of an aircraft in landing configuration. *AIAA Paper 2020-3002.*, 2020.

[29] Konrad A Goc, Oriol Lehmkuhl, George Ilhwan Park, Sanjeeb T Bose, and Parviz Moin. Large eddy simulation of aircraft at affordable cost: a milestone in computational fluid dynamics. *Flow*, 1:E14, 2021.

[30] Sigal Gottlieb, Chi-Wang Shu, and Eitan Tadmor. Strong stability-preserving high-order time discretization methods. *SIAM review*, 43(1):89–112, 2001.

[31] David Greenblatt, Keith Paschal, Chungsheng Yao, Jerome Harris, Norman Schaeffler, and Anthony Washburn. A separation control CFD validation test case. part 1: Baseline & steady suction. In *2nd AIAA Flow Control Conference*, page 2220, 2004.

[32] Kevin Patrick Griffin, Lin Fu, and Parviz Moin. General method for determining the boundary layer thickness in nonequilibrium flows. *Physical Review Fluids*, 6(2):024608, 2021.

[33] Albert E Honein and Parviz Moin. Higher entropy conservation and numerical stability of compressible turbulence simulations. *Journal of Computational Physics*, 201(2):531–545, 2004.

[34] John Hunt, AA Wray, and Parviz Moin. Eddies, streams and convergence zones in turbulent flows. In *Proceedings of the Summer Program*. Center for Turbulence Research, Stanford University, 1988.

[35] Masashi Ichimiya, Yoshiyuki Nakase, and Junichiro Fukutomi. Structure of a turbulence wedge developed from a single roughness element on a flat plate. In *Engineering Turbulence Modelling and Experiments*, pages 613–622. Elsevier, 1993.

[36] Yasushi Ito, Mitsuhiro Murayama, Yuzuru Yokokawa, Kazuomi Yamamoto, Kentaro Tanaka, and Tohru Hirai. Wind tunnel interference effects on Japan Aerospace Exploration Agency's Standard Model. *AIAA Paper 2019-2178.*, 2019.

[37] Javier Jiménez. Cascades in wall-bounded turbulence. *Annual Review of Fluid Mechanics*, 44:27–45, 2012.

[38] Soshi Kawai and Johan Larsson. Wall-modeling in large eddy simulation: Length scales, grid resolution, and accuracy. *Physics of Fluids*, 24:015105, 2012.

[39] Dokyun Kim, Christopher B Ivey, Frank E Ham, and Luis G Bravo. An efficient high-resolution volume-of-fluid method with low numerical diffusion on unstructured grids. *Journal of Computational Physics*, 446:110606, 2021.

[40] Cetin C Kiris, Aditya S Ghate, Jared C Duensing, Oliver M Browne, Jeffrey A Housman, Gerrit-Daniel Stich, Gaetan Kenway, Luis S Fernandes, and Leonardo M Machado. High-lift common research model: Rans, hrles, and wmles perspectives for clmax prediction using lava. In *AIAA SCITECH 2022 Forum*, page 1554, 2022.

[41] Benedikt Konig, Ehab Fares, Mitsuhiro Murayama, and Yasushi Ito. PowerFLOW simulations for the third AIAA high-lift prediction workshop. In *2018 AIAA Aerospace Sciences Meeting*, page 1255, 2018.

[42] Doug S Lacy and Anthony J Sclafani. Development of the high lift common research model (HL-CRM): A representative high lift configuration for transonic transports. *AIAA Paper 2016-0308.*, 2016.

[43] Kelly R Laflin, Steven M Klausmeyer, Thomas Zickuhr, John C Vassberg, Richard A Wahls, Joseph H Morrison, Olaf P Brodersen, Mark E Rakowitz, Edward N Tinoco, and Jean-Luc Godard. Data summary from second aiaa computational fluid dynamics drag prediction workshop. *Journal of Aircraft*, 42(5):1165–1178, 2005.

[44] Johan Larsson, Soshi Kawai, Julien Bodart, and Ivan Bermejo-Moreno. Large eddy simulation with modeled wall-stress: recent progress and future directions. *Mechanical Engineering Reviews*, 3(1):15–00418, 2016.

[45] O Lehmkuhl, GI Park, ST Bose, and P Moin. Large-eddy simulation of practical aeronautical flows at stall conditions. *Proceedings of the 2018 Summer Program*, Center for Turbulence Research, Stanford University. pages 87–96, 2018.

[46] O Lehmkuhl, GI Park, and P Moin. LES of flow over the NASA Common Research Model with near-wall modeling. *Proceedings of the 2016 Summer Program*, Center for Turbulence Research, Stanford University. pages 335–341, 2016.

[47] David Levy, Kelly Laflin, John Vassberg, Edward Tinoco, Mortaza Mani, Ben Rider, Olaf Brodersen, Simone Crippa, Christopher Rumsey, Richard Wahls, J. Morrison, D. Mavriplis, and M. Murayama. Summary of data from the Fifth AIAA CFD Drag Prediction Workshop. *AIAA Paper 2013-0046.*, 2013.

[48] David W Levy, Thomas Zickuhr, John Vassberg, Shreekant Agrawal, Richard A Wahls, Shahyar Pirzadeh, and Michael J Hemsch. Data summary from the first aiaa computational fluid dynamics drag prediction workshop. *Journal of Aircraft*, 40(5):875–882, 2003.

[49] Douglas K Lilly. A proposed modification of the Germano subgrid-scale closure method. *Physics of Fluids A - Fluid Dynamics*, 4:633–635, 1992.

[50] Adrian Lozano-Duran, Sanjeeb T Bose, and Parviz Moin. Prediction of trailing edge separation on the NASA juncture flow using wall-modeled LES. In *AIAA Scitech 2020 Forum*, page 1776, 2020.

[51] Adrián Lozano-Durán, Sanjeeb T Bose, and Parviz Moin. Performance of wall-modeled les with boundary-layer-conforming grids for external aerodynamics. *AIAA Journal*, 60(2):747–766, 2022.

[52] Ivan Marusic and Jason P Monty. Attached eddy model of wall turbulence. *Annual Review of Fluid Mechanics*, 51:49–74, 2019.

[53] Timothy Mauery, Juan Alonso, Andrew Cary, Vincent Lee, Robert Malecki, Dimitri Mavriplis, Gorazd Medic, John Schaefer, and Jeffrey Slotnick. A guide for aircraft certification by analysis. NASA/CR–2021-0015404, 2021.

[54] John B McDevitt and Arthur F Okuno. Static and dynamic pressure measurements on a naca 0012 airfoil in the ames high reynolds number facility. Technical report, 1985.

[55] Rajat Mittal and Parviz Moin. Suitability of upwind-biased finite difference schemes for large-eddy simulation of turbulent flows. *AIAA journal*, 35(8):1415–1417, 1997.

[56] Parviz Moin, Kyle Squires, W Cabot, and Sangsan Lee. A dynamic subgrid-scale model for compressible turbulence and scalar transport. *Physics of Fluids A: Fluid Dynamics*, 3(11):2746–2757, 1991.

[57] George Ilhwan Park. Wall-modeled large-eddy simulation of a high Reynolds number separating and reattaching flow. *AIAA Journal*, pages 3709–3721, 2017.

[58] Ugo Piomelli and Elias Balaras. Wall-layer models for large-eddy simulations. *Annual Review of Fluid Mechanics*, 34(1):349–374, 2002.

[59] Thierry J Poinsot and Sanjiva K Lele. Boundary conditions for direct simulations of compressible viscous flows. *Journal of Computational Physics*, 101:104–129, 1992.

[60] Melissa Rivers and Ashley Dittberner. Experimental investigations of the NASA common research model. In *28th AIAA Applied Aerodynamics Conference*, page 4218, 2010.

[61] Melissa B Rivers, Ralf Rudnik, and Jürgen Quest. Comparison of the NASA common research model european transonic wind tunnel test data to NASA test data. In *53rd AIAA aerospace sciences meeting*, page 1093, 2015.

[62] Christopher L Rumsey and Jeffrey P Slotnick. Overview and summary of the second aiaa high-lift prediction workshop. *Journal of Aircraft*, 52(4):1006–1025, 2015.

[63] Christopher L Rumsey, Jeffrey P Slotnick, and Anthony J Sclafani. Overview and summary of the Third AIAA High Lift Prediction Workshop. *Journal of Aircraft*, 56(2):621–644, 2019.

[64] Christopher L Rumsey, Jeffrey P Slotnick, and Carolyn Woeber. Hlpw-4/gmgw-3: Overview and workshop summary. In *AIAA AVIATION 2022 Forum*, page 3295, 2022.

[65] Christopher L Rumsey, JP Slotnick, M Long, RA Stuever, and TR Wayman. Summary of the first AIAA CFD high-lift prediction workshop. *Journal of Aircraft*, 48(6):2068–2079, 2011.

[66] Juan A Sillero, Javier Jiménez, and Robert D Moser. One-point statistics for turbulent wall-bounded flows at Reynolds numbers up to $\delta^+ \approx 2000$. *Physics of Fluids*, 25(10):105102, 2013.

[67] Jeffrey Slotnick, Abdollah Khodadoust, Juan Alonso, David Darmofal, William Gropp, Elizabeth Lurie, and Dimitri Mavriplis. CFD Vision 2030 study: a path to revolutionary computational aerosciences. NASA/CR–2014-218178, 2014.

[68] Philippe R Spalart. Comments on the feasibility of LES for wings, and on a hybrid RANS/LES approach. In *Proceedings of first AFOSR international conference on DNS/LES*. Greyden Press, 1997.

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

[70] Edward N Tinoco. Analysis of the transonic NASA Common Research Model at high angle of attack. In *AIAA Aviation 2020 Forum*, page 2745, 2020.

[71] Edward N Tinoco, Olaf P Brodersen, Stefan Keye, Kelly R Laflin, Edward Feltrop, John C Vassberg, Mori Mani, Ben Rider, Richard A Wahls, Joseph H Morrison, et al. Summary

data from the sixth AIAA CFD drag prediction workshop: Crm cases. *Journal of Aircraft*, 55(4):1352–1379, 2018.

[72] Ali Uzun and Mujeeb R Malik. Large-eddy simulation of flow over a wall-mounted hump with separation and reattachment. *AIAA Journal*, 56(2):715–730, 2018.

[73] Edward R Van Driest. On turbulent flow near a wall. *Journal of the Aeronautical Sciences*, 23(11):1007–1011, 1956.

[74] John Vassberg, Edward Tinoco, Mori Mani, Ben Rider, Tom Zickuhr, David Levy, Olaf Brodersen, Bernard Eisfeld, Simone Crippa, Richard Wahls, et al. Summary of the fourth AIAA CFD drag prediction workshop. In *28th AIAA Applied Aerodynamics Conference*, page 4547, 2010.

[75] John C Vassberg, Edward N Tinoco, Mori Mani, Olaf P Brodersen, Bernhard Eisfeld, Richard A Wahls, Joseph H Morrison, Thomas Zickuhr, Kelly R Laflin, and Dimitri J Mavriplis. Abridged summary of the third aiaa computational fluid dynamics drag prediction workshop. *Journal of aircraft*, 45(3):781–798, 2008.

[76] AW Vreman. An eddy-viscosity subgrid-scale model for turbulent shear flow: Algebraic theory and applications. *Phys. fluids*, 16:3670–3681, 2004.

[77] Meng Wang and Parviz Moin. Dynamic wall modeling for large-eddy simulation of complex turbulent flows. *Physics of Fluids*, 14:2043–2051, 2002.

[78] M. P. Whitmore, K. P. Griffin, S. T. Bose, and P. Moin. Large-eddy simulation of a Gaussian bump with slip-wall boundary conditions. *Annual Research Briefs*, Center for Turbulence Research, Stanford University. pages 45–58, 2021.

[79] Owen J Williams, Madeline Samuell, Matthew L Robbins, Hariprasad Annamalai, and Antonino Ferrante. Characterization of separated flowfield over gaussian speed-bump CFD validation geometry. In *AIAA Scitech 2021 Forum*, page 1671, 2021.

[80] Carolyn Woeber, James S Masters, and David R McDaniel. Summary of exascale and remeshing efforts for the second geometry and mesh generation workshop. In *AIAA Aviation 2019 Forum*, page 3458, 2019.

[81] Xiang IA Yang and Kevin P Griffin. Grid-point and time-step requirements for direct numerical simulation and large-eddy simulation. *Physics of Fluids*, 33(1):015108, 2021.

[82] Yuzuru Yokokawa, Mitsuhiro Murayama, Takeshi Ito, and Kazuomi Yamamoto. Experiment and CFD of a high-lift configuration civil transport aircraft model. In *25th AIAA Aerodynamic Measurement Technology and Ground Testing Conference*, page 3452, 2006.

[83] Yuzuru Yokokawa, Mitsuhiro Murayama, Masahiro Kanazaki, Katsuichi Murota, Takeshi Ito, and Kazuomi Yamamoto. Investigation and improvement of high-lift aerodynamic performances in lowspeed wind tunnel testing. *AIAA Paper 2008-350.*, 2008.

[84] Yuzuru Yokokawa, Mitsuhiro Murayama, Hiroshi Uchida, Kentaro Tanaka, Takeshi Ito, Kazuomi Yamamoto, and Kazuomi Yamamoto. 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*, page 684, 2010.