

Predicting Orion Launch Abort Acoustics

Francois Cadieux*, Michael F. Barad*, James C. Jensen*, Jordan B. Angel*, Cetin Kiris*
Corresponding author: francois.cadieux@nasa.gov

* NASA Ames Research Center, Moffett Field, CA, USA

Abstract:

Computational fluid dynamics (CFD) scale-resolving simulations were performed to help assess the Orion launch abort system's vibro-acoustic environment in collaboration with the Orion Loads and Dynamics team at the Johnson Space Center (JSC). The primary objectives are to establish best practices and to validate CFD predictions for surface fluctuating pressure spectra with ground, wind tunnel, and flight test data. Excellent agreement is obtained between CFD predictions by the Cartesian adaptive mesh refinement (AMR) module of the Launch, Ascent, and Vehicle Aerodynamics (LAVA) solver and measurements from a ground test representative of the start of a pad abort scenario, and with measurements from a wind tunnel test case corresponding to a transonic, high total angle of attack ascent abort scenario. A simplified (static and at rest) simulation of a pad abort flight test shows good agreement with the flight data, whereas a more costly simulation where the vehicle is accelerating relative to the ground from rest to Mach 0.4 shows consistent over-prediction of the spectra. Simulation results for a supersonic ascent abort flight test are in reasonable agreement with the measurements, but show larger excursions in regions where non-linear interactions between the plume turbulence, shocks, and turbulent boundary layer are strongest – likely due to CFD modeling assumptions, the effect of which are discussed. Overall, the CFD results demonstrate that scale-resolving simulation methodologies like the one presented can help characterize full-scale launch abort vehicle near-field noise and can complement wind tunnel, ground and flight tests by providing trends and reducing uncertainty for scenarios that cannot or are not planned to be tested.

Keywords: Computational Fluid Dynamics, Aero-acoustics.

1 Introduction

The National Aeronautics and Space Administration's (NASA) Artemis missions plan to bring humans to the Moon – and eventually Mars – and return them to Earth safely onboard the Orion spacecraft. Orion will be propelled into space by the powerful Space Launch System, shown in Figure 1 with the Orion crew module and its launch abort system – the focus of this work – outlined in red. The launch abort system (LAS), detailed in Figure 2 is attached to the top of the Orion crew module during launch and is jettisoned once the spacecraft safely reaches the edge of space. The LAS consists of a tower where three rocket motors are located, and a fairing assembly which protects the crew module. Together, the crew module and the LAS form the launch abort vehicle (LAV). The fairing assembly is further split into two sections: the fillet which connects to the tower above, and the ogive which serves as a shield to the crew module. Near the top of the tower, the attitude control motors (ACM) are solid rocket motors with a total of eight nozzles used to steer the LAV. The jettison motor, which serves to separate the LAS from the crew module once the LAV is far enough away from any potential danger, is located closer to the center of the tower. At the base of the tower, above the fillet, sits the abort motor with its four nozzles. The abort motor is a solid rocket motor designed to provide enough thrust to rapidly take the LAV and its crew away from danger in the event of a launcher malfunction. When triggered, the abort motor produces four hot supersonic exhaust plumes that flow along the sides of the LAV to propel it away from a potentially exploding rocket.




Figure 1: Diagram describing the components of the Space Launch System (SLS) Block 1 Crew, including the Orion Multi-Purpose Crew Vehicle and its launch abort system, which together form the launch abort vehicle.




Figure 2: Infographic describing the Orion launch abort system (LAS).

These plumes are highly turbulent and produce significant unsteady pressure waves (near-field acoustics)

that impinge on the surface of the LAS and can cause large vibro-acoustic responses in many of the systems and subsystems. Design and certification to the amplitudes, frequencies, and modalities of these vibrations is critical to ensure the structural integrity of the LAS, and in turn, the safety of the astronauts. In order to derive engineering design requirements for the LAS, an aero-acoustics environment – also called forcing functions for vibro-acoustics calculations – covering the worst possible abort scenarios must be developed.

A campaign of turbulence-resolving simulations were completed in collaboration with the NASA Orion Loads and Dynamics at the Johnson Space Center (JSC) to help answer important engineering questions about the vibro-acoustic environment of the Orion LAS. Chief among those is whether computational fluid dynamics (CFD) can be trusted to give accurate near-field noise predictions on the surface of the LAV, and on the crew module (CM) heat shield. As such, focus is placed on the process of validating CFD predictions by reproducing wind tunnel, ground, and flight tests in as much fidelity as possible. This simulation campaign spanned over 4 years. We refined both numerical method and modeling best practices as we learned from our successes and failures. This article is thus also a means to disseminate these lessons learned to the wider CFD community.

Unfortunately, this launch abort vehicle is not in the public domain and falls under the export restrictions related to the International Traffic in Arms Regulations (ITAR). So we cannot divulge any details of the vehicle’s design, its propellant, potential trajectories, or the acoustic levels expected on its surface. Colormap legends and y-axis of most plots have been thus been removed. Nevertheless, we make every effort to provide apples-to-apples comparisons between CFD predictions and wind tunnel, ground, and flight test measurements without providing numbers. For example, a consistent target accuracy band determined with the NASA Orion Loads and Dynamics acoustics experts will be shown around experimental measurements on all plots. The article is organized as follows. The methodology is presented in section 2. Results of our validation efforts for ground, wind tunnel, and flight tests are detailed in sections 4, 5, 6, and 7 respectively. We conclude with overall lessons learned and an outlook for the future of computational aero-acoustic simulations for near-field acoustics involving hot exhaust jets in 8.

2 Methodology

We use the Cartesian adaptive mesh refinement (AMR) module of the Launch, Ascent, and Vehicle Aerodynamics (LAVA) CFD software [1] to perform simulations of the Orion launch abort system (LAS). The choice of Cartesian mesh paradigm is motivated by the initial goals of the project:

1. To simulate detailed and complex as-built test stands and flight vehicles, potentially in relative motion to the ground,
2. To track ignition over-pressure (IOP) waves as it propagates and interacts with vehicle, measurement devices, and support structures,
3. To accurately model high speed, hot turbulent plume near-field acoustics,
4. To provide predictions for pressure spectra at various locations in a short-enough turnaround time to impact engineering decisions.

The Cartesian AMR with immersed boundary representation can seamlessly handle extremely detailed CAD geometry with no manual effort other than to generate a water-tight triangulated representation of the CAD geometry. High-order shock-capturing methods on Cartesian grids are robust, mature, and perfectly suited to capture moving shocks like IOP waves when coupled with solution-based AMR. Cartesian grids are automatically generated to ensure user-specified mesh spacing $\Delta x \approx 0.02D$, where D is the abort nozzle exit diameter, near the surface of the vehicle (or test stand) and in user-specified regions (via simple boxes or flared cylinders) encompassing the extent of the plumes to accurately model the turbulent acoustics from where they are generated to the vehicle’s surface. It can be computationally challenging to resolve thin attached boundary layers with a Cartesian isotropic AMR approach. Thankfully, the surface pressure spectra – the quantity of interest and the figure of merit – is preeminently driven by the turbulent plumes (free shear layers) for which low aspect ratio cells are ideally suited according to recent jet noise prediction research [2, 3]. The block-structured nature of the Cartesian AMR approach makes it particularly computationally efficient, and its automatic grid generation capability accelerates the overall turnaround time.

2.1 Numerics

We solve the multi-species compressible Navier-Stokes equations in conservative finite difference form (fluxes are reconstructed at the cell faces). The convective fluxes are computed using the fifth-order weighted essentially non-oscillatory (WENO) scheme with improved weights (Z-WENO5) which was shown to be robust and well-suited to shock-turbulence interaction [4]. It uses flux-vector splitting with Rusanov upwind physical fluxes at cell-centers, transforms them to characteristic space based on the Roe-averaged face state, applies the Z-WENO5 reconstruction and transforms the resulting fluxes at the face back to physical space before taking the divergence of fluxes. To ensure the positivity of pressure and density, we replace the Z-WENO5 flux by the local first order Lax-Friedrich’s flux if and only if the Z-WENO5 flux is expected lead to zero or negative pressure or density [5].

All gases and gas mixtures are considered non-reactive and follow the ideal gas equation of state $p = \rho RT$, where p is pressure, ρ is density, R is the gas constant, and T is temperature. Combustion, heat radiation, or particle physics (e.g. aluminum pellets or soot) are ignored. The viscous flux is computed using the typical second-order central conservative finite difference formulation. Subgrid viscous and heat fluxes are not explicitly modeled because WENO schemes have been shown to already provide more numerical dissipation than modern subgrid scale models would [6].

Time integration is performed explicitly with the classical fourth-order Runge-Kutta method for most simulations – AA-2 simulations use the third-order strong stability-preserving Runge-Kutta method [7]. The time step size is chosen or adapted such that the acoustic CFL number is less than 1 at all times, and is typically between 0.5 and 0.95.

2.2 Boundary Conditions

LAVA Cartesian uses an immersed boundary method based on the ghost-cell method [8] extended to compressible flows and modified to handle thin geometry using the ghost-in-fluid approach [9]. Only one layer of ghost cells is filled inside or near the watertight computer-aided design (CAD) model’s discrete surface representation (triangulation). Prior to AA-2 simulations, the WENO scheme transitions to the first-order upwind Rusanov flux near geometry to avoid using incorrect data that is potentially on the other side of thin geometry due to the 7-point stencil of Z-WENO5. This means the effective order of accuracy of the method is dropped to between first and second order for the first two cells outside of the model’s surface. More recent simulations like those for AA-2 use an improved approach where the missing physical fluxes at cell centers beyond the available ghost cells are extrapolated from the interior at up to second order based on the incident Mach number of the flow, allowing to use the WENO scheme all the way to the surface, nominally yielding second order accuracy for the convective flux even for cells next to the surface.

The boundary conditions applied at the surface of the vehicle (or test stand) correspond to an adiabatic slip wall. This is justified because we cannot afford to resolve the very thin turbulent boundary layer developing on the vehicle with Cartesian isotropic cells given the vehicle’s large flight Reynolds number, and the even higher Reynolds number in the abort motor chamber and where the plumes interact with the vehicle surface. AA-2 simulations use no-slip adiabatic walls on the heat shield where the flow is massively separated.

The design of the abort motor’s internal mould line is shown in Figure 3. Typically we would obtain time-accurate and radially-varying profiles of the primary variables just upstream of the nozzle exit from a precursor 2D axisymmetric Reynolds averaged Navier-Stokes (RANS) simulation to apply as supersonic inlet boundary condition. However, the abort motor exhaust flow must turn a sharp corner ahead of the nozzle throat – the so-called “arm pit” – which results in non-axisymmetric (circumferentially-varying) flow at the nozzle exit. The abort motor is thus modeled by an inflow face upstream of the splitter plate inside the abort motor chamber where we apply total pressure and total temperature boundary conditions (Mach number is extrapolated from the interior). The total conditions can be varying in time, as given by abort motor chamber measurements from tests. The abort motor exhaust gas is modeled as a single homogeneous mixture of the combustion by-products (Al_3O_2 powder, steam, HCl , CO , CO_2 , etc) with no species mass fraction diffusion. Its thermodynamic and transport properties were determined by the proprietary NASA-Lewis thermochemistry program applied to the particular nozzle design with equilibrium chemistry for the specific solid rocket fuel and oxidizer used. Helium gas is used for the abort motor plumes in the 80-AS wind tunnel test.




Figure 3: Diagram of abort motor internal mould line and CFD inlet boundary condition face.

The attitude control motor (ACM) is not modeled in the simulations presented. The ACM was not present in the qualification motor test (QM-1), nor the 80-AS wind tunnel test. It was deemed to be a second order effect following successful comparisons of spectra between the 80-AS wind tunnel test measurements and pad abort flight test (PA-1) data [10]. The good agreement between PA-1 and 80-AS spectra also justifies the decision to forgo combustion chemistry modeling for CFD because 80-AS used hot Helium, an inert ideal gas, to model the abort motor exhaust.

The far-field boundary conditions are automatically applied in first order form following the initial conditions. For example, if the flow is supersonic with a specified angle of attack α and side slip angle β , the inlet faces are automatically set to supersonic inlet, and the outlet faces as supersonic outlet, and any remaining faces are assumed to have no gradients at the far-field face (extrapolation from the interior). For all simulations, we ensure the far field boundary conditions are at least 5 vehicle lengths away and we use a minimum of 7 and up to 9 levels of AMR with 2:1 refinement ratios to ensure that any flow features that reach a far-field boundary face have been significantly damped by the coarse mesh out there (usually $\Delta x \approx O(2)D$). As such we do not observe any significant vortical or pressure wave reflections from our otherwise simple far-field boundary conditions.

2.3 Acoustic Processing

The certification requirement for vibro-acoustic analysis is to obtain an average spectra of surface pressure fluctuations for each of the zones defined in Figure 4. To this effect, surface pressure is collected at all the nodes on the watertight triangulated surface representation of the various CAD models. The surface information is sampled at various sampling frequencies depending on the simulation, but always at a rate equal to or greater than a Strouhal number ($St = fD/U_e$, where U_e is the area-weighted nozzle exit velocity) of 1.6 to ensure that the spectra over the range of interest of 0.003 to 0.8 St are not affected by the sampling. Numerical sensors at particular points corresponding to measurement device locations in wind tunnel, ground or flight tests are generally sampled at a much higher rate, closer to every time step taken in the simulation, which is typically on the order of $St = 100$. For many cases, the numerical sensor values are extracted from the nearest surface node information because the sensor locations were not known a priori. This process results in a spatial error in location of at most 0.04 D due to the resolution of the surface mesh. This gives a pressure time series for each point with the same sampling frequency as the surface. Once we have a time series for a given location in space, the acoustic post-processing methodology is fairly straightforward. We first exclude the ignition over-pressure wave from the signal. To further reduce the influence of frequency modes below 10 Hz, we remove from the signal its seventh-order polynomial fit (regression). Then the Welch method is used to compute the power spectrum density S_{xx} with the Hann window. We typically use windows of roughly 60 convective time units ($t^* = L/U_e$, where L is the length from the nozzles to end of zone J) with 50%




Figure 4: Diagram defining the acoustic zones of the Orion launch abort system (LAS) used for vibro-acoustic analysis.

overlap for both CFD predictions and wind tunnel, ground and flight test measurements. This unfortunately corresponds to a single window for many time series derived from CFD simulations because of their short duration. To meet the requirement, we perform spatial averaging of S_{xx} (area-weighted or simple average) over sensors within each acoustic zones, as defined in Figure 4, to improve confidence in CFD predictions at low frequencies despite short simulation durations. From $S_{xx}(f)$ we then perform a third octave band-filter to obtain a single average spectrum per acoustic zone $S_{xx}(f_{\frac{1}{3}})$. Spectra are non-dimensionalized by the area-weighted nozzle exit dynamic pressure q_e : $S_{xx}^*(f_{\frac{1}{3}}) = \frac{U_e}{D} \frac{1}{q_e^2} S_{xx}(f_{\frac{1}{3}})$. Great care is taken to make sure the acoustic post-processing is consistent between test measurements and CFD to ensure that comparisons are as informative and as “apples-to-apples” as possible.

3 Results Overviews

Table 1: Parameters of CFD simulations performed, where M_∞ is the freestream Mach number, $p_{t,\infty}/p_{t,\infty,QM-1}$ is the ratio of total freestream pressure with respect to ground test QM-1, α is the angle of attack, β is the side slip angle, and Δt_{ac}^* is the interval of data collected used for acoustics post-processing in non-dimensional time units $t^* = \frac{U_e}{L} t$. Computational cost is shown in terms of number of cells (N_{cells}), number of cores (N_{cores}) and time to simulation completion (Walltime).

Model	M_∞	$\frac{p_{t,\infty}}{p_{t,\infty,QM-1}}$	(α, β)	$\Delta t_{ac,cfld}^*$	$\Delta t_{ac,exp}^*$	$N_{cells} \times 10^6$	N_{cores}	Walltime (days)
QM-1	0.00	1.00	0.00	85	480	596	3200	18
80-AS	0.70	0.93	Large	1523	16000	388	2240	21
PA-1	0.00	1.04	0.00	95	480	632	8000	21
PA-1	0→0.4	1.15	0→Large	240	500	635	16000	45
AA-2	1.17	0.71	Small	104	480	907	12800	7
AA-2	1.62	0.89	Small	100	480	907	12800	6

The parameters describing all the cases simulated with CFD are detailed in Table 1. Please note that the CAD models for Exploration Mission 1 (EM-1) and Ascent Abort 2 (AA-2) are both representative of the launch abort vehicle (LAV) that will be used in future Artemis missions, and have only minor differences in their outer mold line (OML). In contrast, the qualification motor 1 (QM-1) model is a ground test stand, and 80-AS corresponds to a simplified 6% scale wind tunnel model. The Pad Abort 1 (PA-1) flight test

model corresponds to the launch abort system attached to the crew module with a simple connector instead of the fairing assembly (fillet and ogive) that was used in that flight test.

4 Abort Motor Qualification Ground Test



Figure 5: Picture of the Abort Motor Qualification ground test (QM-1) showing the abort motor plumes, the test stand, the PIAT tower (used to measure the IOP wave strength) to the left of the test stand, the two near-field pressure array (NFPA) measurement towers to the right of the test stand, the heat shield measurement apparatus just above the test stand and the crane from which it is suspended further above.

The Qualification Motor ground test (QM-1) occurred on Jun 15, 2017. During this test, the latest design of abort motor was fired with the nozzles pointing away from the ground, producing thrust into the ground that the test stand was designed to resist. The solid rocket motor burn lasted roughly 2.5 seconds. A picture of the test showing the various apparatus used to measure pressure over a wide frequency range (equivalent to a Strouhal number of up to 11.4) is found in Figure 5.

The grid generated to simulate the QM-1 test is detailed in Figure 6. It contains a total of 596×10^6 cells, 60% of which are on the finest refinement level which encompasses the plumes and the heat shield measurement apparatus. Given the target was to validate acoustic spectra taken on this heat shield measurement device, some simplifications were made: the ignition over-pressure wave was not tracked with solution-based AMR and fixed refinement regions were used throughout, and the time-varying total pressure and temperature conditions on the abort motor inlet face were assumed taken to be constant and equal to the mean conditions over the time interval of interest (to promote faster arrival to statistical stationarity).

No turbulence is resolved within the internal mould line of the abort motor: the flow transitions to turbulence naturally due to the mixture of shear and temperature gradient instabilities of the hot high speed exhaust with respect to the quiescent ambient temperature air around it. This is clearly visible in Figure 7 where clear turbulence structures develop immediately downstream of the nozzle exit and evolve even beyond the height of the crane used to suspend the heat shield measurement apparatus before being dissipated away by mesh coarsening.

The geometry of the test stand loosely corresponds to the launch abort vehicle aft of the jettison motor, but turned upside down. The test stand itself does not follow the outer mould line of the vehicle due to the requirement of bolting it down, and to save costs. Third octave spectra comparisons between CFD predictions and QM-1 ground test measurements is shown in Figure 8. The sensors located closest to the plume in our CFD simulation are unfortunately not on the finest mesh level, and as such display an early roll-off compared to the test measurements. This is an oversight in the mesh generation process which focused predominantly on the heat shield predictions. Despite this, CFD results are mostly within the accuracy




Figure 6: Slice through the computational grid showing outlines of Cartesian boxes each containing 32^3 cells. The cut is colored by instantaneous Mach number, where blue is low and red is high. Plume refinement regions and their connection to the heat shield measuring apparatus are outlined in green. The computational representation of the test apparatus is shown in gray.




Figure 7: Snapshot of QM-1 simulation showing iso-surfaces of Q-criterion colored by pressure, where blue is low and red is high. The iso-surfaces showcase the detailed turbulent structures captured by the simulation, while the variations from blue-to-white-to-red indicate pressure fluctuations or waves in space that can travel and impact the vehicle's surface. Credit: Timothy Sandstrom

target across all sensors and display spectral shapes very well correlated with the measurements. It is clear that the CFD spectra are not quite converged given the significant oscillatory behavior especially visible at low frequencies for zones H to I2 given the limited number of sensors over which to perform spatial averaging. The duration of this simulation is at best one quarter of the ground test measurements' duration, which explains some of these discrepancies. The CFD captures both the frequency at which the peak occurs, and the peak value within the accuracy target. Results for sensors further away from the plume axis (radially) are not shown because no attempt was made to keep a fine enough mesh to capture acoustics there, and so the spectral content rolls off at much lower frequencies. More importantly, the heat shield sensors show excellent agreement with the ground test measurements. This is evident in the area-weighted third octave

spectra where CFD results track very closely the experimentally measured values from St of 0.06 to 0.8, and only deviate in the very low Strouhal number range due to short duration effects. Obtaining this level of agreement was a significant milestone for this project and for the agency in terms of the development and deployment of computational aero-acoustics capabilities.




Figure 8: Comparison of non-dimensional third octave filtered spectra between CFD and QM-1 measurements plotted on a logarithmic scale for NFPA sensors from nearest to abort motor nozzle (zone H), to farthest (zone I2), as well as heat shield sensors (zone J). Zones H to I-2 only have two sensors over which to perform spatial average, whereas zone J has enough sensors placed on a grid to compute area-weighted average spectra.

5 80-AS Wind Tunnel Test

The purpose of the 80-AS wind tunnel test was to assess the aero-acoustic environment of the launch abort system for aborts occurring over a wide range of angles of attack, side-slip, Mach numbers, and ambient pressure (to simulate different altitudes) [11]. Simulations of the 80-AS wind tunnel model with hot Helium plumes were conducted in order to anchor the CFD predictions for the Artemis-1 vehicle specifically for transonic abort scenarios at high total angle of attack (sum of angle of attack and side slip angle). The 80-AS wind tunnel model was outfitted with 237 pressure transducers placed strategically over the launch abort system and crew module heat shield surface [11]. The experimental measurements were acquired at a sampling Strouhal number of 22.8. The CFD numerical sensors match the locations of the wind tunnel model and are output at every time step, which changes as the simulation advances due to the CFL-based time stepping, but yields a mean sampling Strouhal number of approximately 1580. The maximum acoustic wave frequency the CFD can capture due to the spatial discretization (mesh size and convective scheme) is closer to $St=0.8$, and drops to $St=0.4$ in coarser regions like the heat shield – see Figure 9. Given the unknown effect of the high angle of attack and side slip, we leveraged solution-based AMR to track pressure gradients and vorticity in an effort to automatically capture the plumes' deviation from their nozzle axis. The resulting grid total number of cells oscillates around 350×10^6 cells. Unfortunately, our choice of AMR parameters did not guarantee fine mesh over much of the heat shield. The wind tunnel measurements benefit from the




Figure 9: Slice through snapshot of the Cartesian AMR CFD domain for the 80-AS wind tunnel model simulation at transonic conditions and high total angle of attack. Box outlines are shown in black, vehicle is shown in gray. Each box contains 32^3 cells. The slice is colored by Mach number where blue is low and red is high.

very long duration corresponding to 16,000 convective time units based on the length from the nozzle exit to the end of the vehicle (t^*). The geometric scale of the wind tunnel model is 6% of its flight size, but the nozzle exit velocity is roughly matched (albeit with heated Helium instead of solid rocket motor combustion by-products). Keeping an equivalent resolution to the QM-1 test per nozzle diameter means that Δx was also scaled by 0.06, drastically reducing the maximum stable explicit time step the CFD could take and limiting our simulation duration severely ($1523 t^*$) compared to the wind tunnel ($16000 t^*$).

Due to this large duration difference between CFD and wind tunnel, and the lack of many cycles at low frequencies for CFD, it is more instructive to look at spatially averaged spectra instead of individual sensors. This is because the spatial average is similar to an average across multiple realizations and thus reduces the spread of the CFD predictions at low frequency. So to quickly evaluate the CFD predictions, we do a simple average of the sensors that lie within the extent of each acoustic zone and present the comparison in Figure 10. The spectra derived from the CFD simulation for zones G to I2 show agreement with the wind tunnel measurements to within the target accuracy over a wide range of frequencies. The agreement is still excellent for zone J, but the CFD displays premature roll-off due to insufficient resolution in the wake of the vehicle near the heat shield. The only salient disagreement occurs in zone F where the CFD under-predicts the mid to high frequency range. We attribute this to insufficient resolution near the nozzle exit to properly capture the breakdown into turbulence of the supersonic jet. Despite this limitation, this demonstrates that the methodology employed here is sound and can yield predictions consistent both qualitatively – in terms of capturing the trends going axially down the zone of the vehicle – and quantitatively with careful wind

tunnel measurements.




Figure 10: Comparison of simple average of non-dimensional third octave filtered spectra (plotted on a logarithmic scale) of all sensors in each zone between 80-AS wind tunnel measurements and CFD predictions.

6 Pad Abort Flight Test

The Pad Abort (PA-1) flight test occurred on May 6, 2010 at White Sands Missile Range. During this test, a mockup of the launch abort vehicle accelerated from rest on the ground to Mach 0.7 at roughly 2000 m altitude – see Figure 11. Simulations reproducing different segments of the pad abort 1 (PA-1) flight test were conducted in order to validate the CFD-derived acoustic predictions with flight data. The goal was also to investigate the effect of acceleration on the vehicle surface fluctuating pressure. To begin this effort, we simulated the PA-1 vehicle in free-air (no ground effect) at zero velocity ($M_\infty = 0$, $\alpha = 0^\circ$, $\beta = 0^\circ$) to approximate the first 0.4 seconds of flight where the vehicle is moving very slowly away from the ground. Although the vehicle is still close to the ground during this segment of the PA-1 flight test, we chose to remove the influence of the ground because we are primarily interested in the near-field acoustics, which are driven by the plumes and are not directly affected by acoustic reflections coming from the ground (there were no sensors on the heat shield for this test, where this effect might have been important). This simulation used a number of improvements to our best practices:

- subcycled time integration for higher parallel efficiency,
- user-generated mesh with no solution-based AMR to guarantee fine mesh over entire plumes and vehicle heat shield, and
- constant abort motor chamber pressure and temperature taken from PA-1 measurements at $t = +0.2$ s, similar to what was done for QM-1.

These improvements increased the total mesh size from 350×10^6 cells to 632×10^6 cells, without negatively affecting the turnaround time. A cut through the mesh is shown in Figure 12. Notice that all the plumes



Figure 11: Picture from the PA-1 flight test, May 6 2010, White Sands Missile Range. Photo credit: NASA.

are captured on the finest level with a conical refinement region extending from just upstream of the abort motor nozzles all the way to 0.5 meter downstream of the end of the heat shield. This will guarantee that CFD can capture frequencies up to $St=0.8$ accurately over the whole vehicle downstream of the nozzles. Notice also that the vehicle nozzle axis are not aligned with the Cartesian directions, unlike many of the previous simulations. This was done in an effort to keep the vehicle in the structural dynamics coordinate system and avoid having to rotate and translate sensor locations.

Following these improvements, we simulated the accelerating and banking vehicle following the best estimated trajectory reconstructed from the PA-1 flight test (prescribed motion) for roughly 1.25 seconds after ignition. The accelerating simulation also used the time-varying measured abort motor chamber pressure, as opposed to a fixed value. The vehicle moves with respect to the Cartesian grid and the ground, so there are significant additional computational expenses incurred: moving the vehicle at each Runge-Kutta substep at which the time is advanced, recomputing the geometric quantities at each of the cells where immersed boundary conditions need to be applied, and re-gridding to ensure that the fine mesh follows the vehicle and the plumes as they move through the domain. A slice through the CFD domain shown in Figure 13 exemplifies the setup, where the bottom of the domain is the ground. Figure 14 illustrates the dynamic nature of this accelerating PA-1 simulation and the complexity of the fluid dynamics captured with a passive particle animation derived from the nearly 200 terabytes of volumetric data produced. This full PA-1 flight test simulation thus required additional HPC resources and a longer turnaround time – see Table 1.

The flight test vehicle was outfitted with 60 dynamic pressure sensors on the connector and the crew module replica that were numerically reproduced in the CFD simulations. For both of these simulations, the sensor data is extracted from the vehicle surface data, which is output at different frequencies, both well-above the Strouhal number range of interest. The flight test vehicle on-board data recorder has a sampling Strouhal number of 1.5 but uses an 8-pole filter to remove all content beyond $St=0.46$. As such, third octave band-filtered data is shown only up to $St=0.46$ to present a fair comparison between CFD and test measurements.

We use the acoustic zone definitions to bin the PA-1 sensors into each zone, and provide a simple average




Figure 12: Slice through the fixed Cartesian CFD domain with box outlines shown in black, vehicle is shown in gray. Each box contains 32^3 cells. The slice is colored by Mach number taken from a snapshot of the PA-1 $M_\infty = 0$ simulation, where blue is low and red is high.

spectrum of all sensors in a given zone. This also allows convenient rapid overview of the accuracy of CFD, and the effect of acceleration in Figure 15. Post-processed results show close agreement between the flight test measurements and the static PA-1 $M_\infty = 0$ CFD for zone H: the CFD predictions fall within the accuracy target across the entire range of Strouhal numbers of interest. In zone I1 and I2, the static CFD results consistently under-predict the mid-range of Strouhal number from 0.02 to 0.08 while correctly capturing the low and high Strouhal range. The accelerating PA-1 simulation consistently over-predicts the measurements across all zones, except I2 where it falls just within the accuracy target over much of the range. Neither meets simulation strictly meets our validation criteria, but they bound the data from above and below which is encouraging.

It is clear that moving the vehicle with respect of the Cartesian background grid resulted in an increase in pressure fluctuations in the CFD. This is well visualized in Figure 16 where the overall sound pressure level has clearly increased across the entire surface of the vehicle and on the cut plane through the nozzles. If we transform the Navier-Stokes equations into the accelerating frame of the body, we end up with time-dependent (and spatially constant for pure translational motion) face velocities (corresponding to the vehicle's velocity with respect to the fluid at any given time) that are to be super-imposed on top of our reconstructed face fluxes [12]. Changes in vehicle translational velocity do not affect spatial gradients in the accelerating body frame. Only rotational motion could create additional gradients in velocity. The great majority of the vehicle's acceleration is translational in nature and thus should not increase the amplitude of the fluctuations in the plume, let alone in their core as is visible in Figure 16. We tracked down this unexpected rise in fluctuation amplitudes to the immersed boundary method. Ghost cell methods (GCM) in general do not guarantee conservation of mass at cells that are cut by the vehicle geometry, but produce small mass-




Figure 13: Slice through a snapshot of the Cartesian CFD domain with box outlines shown in black, vehicle is shown in green. Each box contains 16^3 cells. The refinement zones are user-specified like in Figure 12 but move with the vehicle. The slice is colored by gauge pressure taken from a snapshot of the PA-1 accelerating simulation where blue is low and red is high.

conservation errors for flows over smooth bodies and are much computationally cheaper than cut-cell methods that do guarantee conservation of mass. However, GCM are particularly sensitive to under-resolved curvature and sharp corners coupled with motion relative to the Cartesian background mesh. In such situations, the GCM can cause spurious acoustic disturbances of order 1% of the local pressure. This is a known problem for which fixes have been proposed [13]. Unfortunately, such improvements to mass conservation at cells neighboring the geometry are typically incremental and come at high computation cost because they require re-computing cut-cell face areas (borrowed from cut-cell methods). The relatively sharp corners inside the abort motor chamber both at the contraction before the “armpit” and at the “armpit” itself produce acoustic perturbations that are quite large given the pressure inside the inner mould line of the abort motor is extremely high. These spurious perturbations produce sharp wave fronts that propagate to the nozzle exit and energize the free shear layers of the plumes, explaining the overall increase in pressure fluctuations. In retrospect, a better approach would have been to perform this PA-1 simulation in the accelerating reference frame and disregard the effects of the presence of the ground given the lack of sensors on the heat shield for this test (which is likely the only region that would be directly affected by the presence of the ground).

Acceleration of a supersonic jet’s co-flow itself is not a driving parameter in noise levels. From a jet noise perspective, the expectation is that the noise levels should scale with the convective Mach number $M_c = |U_e - U_c| / (\frac{1}{2}(a_e + a_c))$, where a_e and a_c are the sound speeds of the exhaust at the nozzle exit and of the ambient fluid in the vicinity of the nozzle exit respectively and remain largely independent of Reynolds number for $Re_j > O(10^5)$ [2]. Thus, as the freestream velocity increases (in the subsonic range), one expects less shear, and consequently a decrease in noise levels. However, this decrease in noise levels is counter-acted by the increase in freestream dynamic pressure (momentum) which pushes the angled plumes closer to the vehicle surface, moving the noise source closer to the sensors under the plume and causing them to record relatively constant levels as the vehicle reaches higher velocity during its trajectory. Some of the sensors located near rapid changes in vehicle cross-sectional area actually record higher OASPL as the vehicle moves into the transonic range, likely due to the start of plume-shock interaction, regardless of their azimuthal position (under versus in-between plumes).




Figure 14: Animation of the accelerating PA-1 simulation showing passive particles seeded at the nozzle and convected with the volumetric data generated. Particles are colored by velocity magnitude where white is high and dark orange is low velocity. The vehicle is shown in white, and the ground is shown as a light gray grid to clarify the perspective and relative motion of the vehicle. Credit: Timothy Sandstrom

This supports the argument that the difference in time scale of the plumes (fast) to the vehicle's change in velocity and attitude (slow) are large enough to consider performing individual (non-accelerating) CFD simulations to cover different points in a given abort trajectory instead of performing a single accelerating CFD simulation across the full abort trajectory. As such, using trajectory points to obtain statistically stationary solutions at the trajectory extremes (i.e. abort initiation, and abort max velocity during abort motor burn) should be sufficient to bound the flight data, even with the knowledge that some low-to-mid frequency phenomenon (due to vehicle acceleration and change in orientation) may not be accurately captured (likely under-predicted).




Figure 15: Left: LAV acoustic zones overlaid with transparency on top of the PA-1 vehicle. Right: Comparison of simple average of non-dimensional third octave filtered spectra of all sensors in each zone between PA-1 flight test measurements and CFD predictions from static $M_\infty = 0$ and accelerating simulations.




Figure 16: Post-processed CFD results for the static and accelerating PA-1 case showing overall sound pressure level (OASPL) on vehicle surface and cut plane through nozzles; blue is low and red is high.

7 Supersonic Ascent Abort Flight Test




Figure 17: Picture from the AA-2 flight test showing the Orion LAV (top) and the Minotaur launcher below it, July 2 2019, Cape Canaveral, FL. Photo credit: NASA.

The ascent abort (AA-2) flight test took place on July 2, 2019 at Cape Canaveral. During this test, the launch abort vehicle sat atop a Northrop Grumman provided Minotaur launcher that accelerated it to supersonic speed, at which point the abort motor was triggered, quickly taking the launch abort vehicle away from the launcher – see Figure 17 which shows the launch abort vehicle accelerating away from the launcher below it. The AA-2 flight test provided the most detailed and accurate vibro-acoustic data to date. Acoustics recorded onboard exceeded previous predictions based on scaling of previous flight data (PA-1) and wind tunnel test data (80-AS). It thus serves as a particularly critical validation case for CFD. To bound the freestream conditions of the AA-2 flight vehicle during abort, two non-accelerating CFD cases were simulated. $M_\infty = 1.17$ corresponds roughly to the first 0.5 seconds after abort initiation. $M_\infty = 1.62$ is meant to represent the last 0.5 seconds before the abort motor chamber pressure drops off. Most of AA-2 flight test saw very little angle of attack (α) or side slip (β), so those values, although different between $M_\infty = 1.17$ and $M_\infty = 1.62$, are dropped from the CFD case names for brevity. Additional assumptions are made to simplify simulations' setup for the validation effort:

1. ACM are not firing,
2. $M_\infty = 1.62$ simulation uses the same time interval of abort motor chamber pressure as $M_\infty = 1.17$: which slightly over-predicts chamber pressure during that part of AA-2 flight, but is useful to isolate the effects of changes in Mach number and dynamic pressure; and
3. Minotaur launcher is not included despite remaining fairly close to the vehicle for the first 0.75 seconds after abort initiation.

The CFD mesh definition for both cases is identical for both cases and is shown in Figure 18 along with a snapshot of the Mach number toward the end of the $M_\infty = 1.17$ simulation.




Figure 18: Slice through CFD domain for AA-2 simulations with Cartesian box outlines shown in black, vehicle is shown in gray. Each box contains 32^3 cells. The slice is colored by Mach number taken from a snapshot of the $M_\infty = 1.17$ simulation. Refinement zones are defined by flared cylinders and regular cylinders, so this cut is representative of the resolution at any other azimuthal angle.

Large differences in the AA-2 flight test data for sensors in-between plumes as opposed to under a plume prompted the splitting of acoustics zones into in-between and under plume segments of 22.5° each, as well as splitting the heat shield zone axially for sensors near the shoulder versus those near the center of the heat shield. Figure 19 contains all the information needed to quickly evaluate the CFD predictions in terms of how well they agree with flight data and whether they predict similar trends over these zones. It is immediately apparent that the neither CFD simulation meets our accuracy criteria across all zones, particularly when it comes to the spectra from sub-zones in-between plumes. Many qualitative trends are nevertheless captured well:

- relatively flat spectra for zones F and G,
- very little difference for sub-zones under and between plumes for zones F and G,
- spectral shape (slope and roll-off) for zones I to J.

We surmise it may be due to the effect of the ACM actuation frequency or stem from shedding off of an upstream protuberance. The CFD predictions fail to capture the substantial differences recorded by the

flight sensors in-between plumes compared to the sensors underneath the plumes in zones I1 and I2. They similarly fail to capture the large low Strouhal number content in zone G. And in zone H, the $M_\infty = 1.62$ severely over-predicts the spectral content in-between plumes across a wide Strouhal number range. Both CFD simulations predict a drop in spectral content at mid-to-high Strouhal numbers in going from zone I1 to zone I2, but the flight data shows those remaining almost exactly constant. Lastly, the CFD under-predicts the levels in zone J where it has previously showed great agreement (QM-1 and 80-AS). The discrepancies in zone J can at least partially be explained by the presence of the Minotaur launcher nearby during the first 0.75 seconds. The presence of solid objects almost always has a large impact on near-field acoustics because they change the mean pressure environment and can create significant pressure gradients in supersonic flows.

In zone F, the flight test recorded a rather pronounced but broad tonal peak centered around $St=0.25$ for the in-between plumes sensors. This type of tonal peak was not present in PA-1 or 80-AS data, so we surmise it is related to upstream influence, either from the ACM, or from vortex shedding behind a protuberance that was not present in the previous tests. The fact that the CFD does not show such a peak thus seems to point to the effect of the ACM as the cause of this peak because geometric effects from protuberances should be relatively well-captured in the CFD.

To shed light on why we are seeing these discrepancies, Figure 20 compares the OASPL of each CFD simulation on the surface of the vehicle and on the cut plane through the nozzles. Notice the bow shock upstream of the abort motor nozzles due to the plumes changing the effective cross-sectional area of the vehicle, and the bow shock just upstream of the fillet-ogive junction due to the abrupt axial change in cross-sectional area of the vehicle. The latter bow shock is rather weak, so the OASPL in its region has a lower peak and is spread across a wider region between zones G and H as it gets pushed around by the plumes more than in the $M_\infty = 1.62$ case where it is stronger and more localized to zone H. Notice also how the diameter of the plume cores denoted by the dark blue region has actually shrunk from $M_\infty = 1.17$ to $M_\infty = 1.62$. The increase in freestream dynamic pressure has brought the abort motor bow shock closer to the nozzles, increasing the local mean pressure environment near the nozzle exit, decreasing the level of under-expansion of the plumes relative to their surrounding environment. The closer presence of the shock also means the jet convective Mach number actually increases instead of decreasing when comparing $M_\infty = 1.17$ to $M_\infty = 1.62$. It also pushes the plumes closer to the vehicle, increasing the scrubbing effect. This explains why $M_\infty = 1.62$ predicts higher OASPL over most of the vehicle, as evidenced also by the spectra in Figure 19. But these trends are not seen in the flight data – the recorded OASPL stayed remarkably constant throughout AA-2's abort motor burn.

We deduce from this that ignoring the influence of the ACM was an over-simplification that caused the CFD to miss some important physics, particularly in the supersonic regime. The ACM by design are always on after the abort is triggered and produce 8 very small high speed jets oriented normal to the vehicle's longitudinal axis around the nose of the vehicle. The momentum of these jets change the effective cross-sectional area of the vehicle and cause a stronger bow shock to occur ahead of the vehicle's nose. This in turn significantly changes the abort motor plumes' mean pressure environment. For example, we would not expect to see a strong bow shock ahead of the abort motor nozzles anymore, as is especially salient at $M_\infty = 1.62$ in Figure 20. Similarly, this will affect the strength and location of the plume-shock interaction driven by the fillet-ogive junction. Ultimately, the decision not to include them was driven by computational cost and turnaround time: to model them accurately, we would have needed significantly finer resolution (and thus smaller time step) near the ACM which could have easily tripled our overall turnaround time.

The lack of ACM modeling however cannot explain why the CFD does not predict the drop in fluctuations in-between plumes on the ogive. Instead, this is more likely to be driven by the assumption of zero wall-shear stress: we expect that the fillet-ogive shock would interact with the evolving turbulent boundary layer on the vehicle and likely cause a recirculation zone that would protect the surface pressure sensors from the large pressure fluctuations caused by unsteady shock movement. This could be addressed in the future with a wall-modeled large-eddy simulation (WMLES) approach, but was deemed too costly due to the higher grid density requirements of WMLES for this problem.




Figure 19: Comparison of simple average of third octave spectra in each acoustic zone between AA-2 flight data and CFD predictions for conditions matching AA-2 abort initiation ($M_\infty = 1.17$) and conditions at end of “early-burn” ($M_\infty = 1.62$).




Figure 20: OASPL from AA-2 CFD results on vehicle surface and cut plane through nozzles where blue is low and red is high.

8 Conclusions

Scale-resolving CFD simulations were carried out to develop a capability to predict Orion launch abort vehicle near-field acoustics and validate it across ground, wind tunnel, and flight tests. The effort leveraged the LAVA Cartesian AMR module in an effort to adapt to various CAD models rapidly with its automatic volume grid generation and provide a fast enough turnaround time to impact engineering decisions to be made about the vehicle's requirements in terms of its expected aero-acoustic environment. Comparison of predicted to measured average surface pressure spectra demonstrated the accuracy of this approach for the qualification motor (QM-1) ground test and for transonic high total angle of attack conditions for the 80-AS wind tunnel test. The predictions for the pad abort flight test (PA-1) assuming a static vehicle and no ground effect turned out to be closer to the flight data (over most of the vehicle) than our best attempt at replicating the flight test with the vehicle accelerating from rest relative to the ground due to a pitfall in our immersed boundary methodology. Although neither simulation matched our accuracy target, the predictions from both bounded the PA-1 flight test data from above and below. Simulations of the supersonic ascent abort flight test also did not meet our accuracy target for surface pressure spectra, but captured a number of trends well. Investigation into these discrepancies revealed that some assumptions made early on, namely that the attitude control motors (ACM) and wall shear stress on the vehicles surface could both be ignored, were probably adequate in the subsonic regime, but caused the simulations to miss important physics in the transonic and supersonic regime. We hope that the turbulence-resolving simulation results presented here will convince readers of the value of CFD predictions in determining aero-acoustic environments in settings where flight measurements are scarce – despite some discrepancies, the CFD predictions were typically in the right ball-park and often conservative.

8.1 Lessons Learned

- Solution-based AMR is great for capturing moving fronts, but we found it to be un-reliable to obtain converged surface pressure spectra because any coarsening between sound source and receiver will cause nearly half the spectral content to disappear, and cause rather unpredictable behavior if refinement and coarsening occurs periodically.
- Some upstream effects that do not appear to be primary noise sources, like the ACM, turn out to be important in the transonic and supersonic regime due to their indirect effect on the mean pressure environment.
- Writing down every assumption made and re-evaluating their individual and collective potential impact with fluid dynamics and acoustics experts should become part of the process before starting any new aero-acoustics simulations at this level of computational cost.
- Shock boundary layer interaction cannot be ignored at transonic and supersonic speeds when the goal is to predict surface pressure spectra.
- Much resources were spent to simulate the vehicle accelerating with respect to the ground for PA-1 when the data collected did not show evidence of ground effects: better results could have been obtained at lower cost by ignoring the ground and transforming the equations to the vehicle's frame of reference.
- It is often better to include more geometry elements than less: the AA-2 Minotaur launcher should have been included at a representative distance from the Orion launch abort vehicle in our AA-2 $M_\infty = 1.17$ (non-accelerating) simulation to better match the acoustics on the heat shield recorded during the start of the abort.
- If CFD is not accurate enough to match flight data, it is still be valuable to use it to extract trends by looking at delta's between simulations to gain more insight where test data is not available.

8.2 Outlook

Throughout this work, we have identified a number of areas that would improve the state-of-the-art in scale-resolving simulations to better predict near-field acoustics relevant to the Orion launch abort system. First and foremost, any methodology that would allow us to capture upstream effects that change the mean pressure field like the ACM, without significantly impacting the turnaround time of a single simulation could increase the accuracy of the predictions. Similarly, although significant progress has been made to develop best practices for WMLES of (unpowered) aircraft[14] and of supersonic jets [3] separately, the community could benefit from the development of best practices for simulations that meld the two (supersonic jets, and some type of aircraft), especially for situations where there is significant interaction between the jet and the vehicle. There is also a need for computationally efficient immersed boundary methods that handle non-smooth geometry without creating spurious pressure waves when the geometry is accelerating with respect to the background Cartesian mesh.

9 Acknowledgements

The authors are deeply grateful for the collaboration with the Orion Loads and Dynamics team: Quyen Jones, Vincent Fogt, Jay Panda, Kenneth Fiorelli, and Christian Behrend. Their knowledge and expertise were extremely valuable throughout the project, from establishing best practices to gaining a deeper understanding of the particular problems of interest and sources of uncertainty in the design and certification of the Orion launch abort system. Special thanks to Jay for providing the kulite data for all experimental measurements and example scripts on how to read it and post-process it.

The authors would also like to thank Jeffrey Housman who provided critical insights into rocket plume modeling and acoustics. Many thanks also to Timothy Sandstrom who maintains our high-performance geometry kernels and made several eye-catching visualizations for this project. This work was made possible thanks to funding from the Orion project. Computing resources supporting this work were provided by the NASA High-End Computing (HEC) Program through the NASA Advanced Supercomputing (NAS) Division at Ames Research Center.

References

- [1] C. Kiris, J. Housman, M. Barad, C. Brehm, E. Sozer, and S. Moini-Yekta. Computational Framework for Launch, Ascent, and Vehicle Aerodynamics (LAVA). *Aerospace Science and Technology*, 55:189–219, August 2016.
- [2] Guillaume A Brès and Sanjiva K Lele. Modelling of jet noise: a perspective from large-eddy simulations. *Philosophical Transactions of the Royal Society A*, 377(2159):20190081, 2019.
- [3] Gerrit-Daniel Stich, Aditya S Ghate, Jeffrey A Housman, and Cetin C Kiris. Wall modeled large eddy simulations for NASA’s jet noise consensus database of single-stream, round, convergent jets. In *AIAA SCITECH 2022 Forum*, page 0684, 2022.
- [4] C. Brehm, M. Barad, J. Housman, and C. Kiris. A Comparison of Higher-Order Finite-Difference Shock Capturing Schemes. *Computers & Fluids*, 122:184–208, November 2015.
- [5] Xiangyu Y Hu, Nikolaus A Adams, and Chi-Wang Shu. Positivity-preserving method for high-order conservative schemes solving compressible euler equations. *Journal of Computational Physics*, 242:169–180, 2013.
- [6] G. Castiglioni, J.A. Domaradzki, V. Pasquariello, S. Hickel, and M. Grilli. Numerical simulations of separated flows at moderate reynolds numbers appropriate for turbine blades and unmanned aero vehicles. *International Journal of Heat and Fluid Flow*, 49:91–99, 2014.
- [7] Sigal Gottlieb. On high order strong stability preserving runge-kutta and multi step time discretizations. *Journal of Scientific Computing*, 25(1):105–128, 2005.
- [8] Rajat Mittal, Haibo Dong, Meliha Bozkurtas, FM Najjar, Abel Vargas, and Alfred von Loebbecke. A versatile sharp interface immersed boundary method for incompressible flows with complex boundaries. *Journal of Computational Physics*, 227(10):4825–4852, 2008.

- [9] Kazuhiro Nakahashi. Immersed boundary method for compressible euler equations in the building-cube method. *AIAA Paper*, 3386:2011, 2011.
- [10] J. Panda. Aero-acoustic Environment on MPCV Launch Abort Vehicle during Abort Motor Burn. Technical Memorandum NASA/TM-2012-216059, National Aeronautics and Space Administration, 2012.
- [11] Jayanta Panda, George James, Nathan Burnside, Robert Fong, Vincent Fogt, and James Ross. Use of heated helium to simulate surface pressure fluctuations on the launch abort vehicle during abort motor firing. In *17th AIAA/CEAS Aeroacoustics Conference (32nd AIAA Aeroacoustics Conference)*, page 2901, 2011.
- [12] Robert Biedron and James Thomas. Recent enhancements to the fun3d flow solver for moving-mesh applications. In *47th AIAA Aerospace Sciences Meeting including The New Horizons Forum and Aerospace Exposition*, page 1360, 2009.
- [13] Jung Hee Seo and Rajat Mittal. A sharp-interface immersed boundary method with improved mass conservation and reduced spurious pressure oscillations. *Journal of Computational Physics*, 230(19):7347–7363, 2011.
- [14] Cetin C Kiris, Aditya S Ghate, Jared C Duensing, Oliver M Browne, Jeffrey A Housman, Gerrit-Daniel Stich, Gaetan Kenway, Luis M Dos Santos 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.