

Enhanced Simulation Techniques in Predicting Sonic Boom Loudness Using CFD

S. Neuhoff*, C. Ashby*, J. Housman*, J. Duensing*

Corresponding author: scott.neuhoff@nasa.gov

*NASA Ames Research Center

Abstract

This paper outlines advancements in predicting sonic boom loudness within the Launch, Ascent, and Vehicle Aerodynamics (LAVA) computational framework. Traditionally, a two step process consisting of a steady state computational fluid dynamics problem for near-field analysis and a far-field propagation solver for calculation of loudness metrics has been used. Improvements to this process made in this work include utilizing a high-order space marching method for mid-field computations, developing a novel output-based mesh adaptation method targeting error in near-field pressure signatures, and developing a robust scripting system using curvilinear grids to increase robustness and simplify the process of running large databases of simulation cases. These advancements are detailed and applied to the simulation of the X-59, presenting comparative cost and timing analyses between the prior two step workflow and the current three step procedure. We achieve increased accuracy and robustness for loudness predictions with at least a 50% computational cost reduction.

1 Introduction

Loudness predictions for supersonic flight vehicles using CFD-based methods are traditionally made using a two step procedure first introduced in [1]. The first step uses computational fluid dynamics (CFD) to solve the steady Euler or Reynolds-Averaged Navier Stokes (RANS) equations in a domain around the aircraft out to a distance of 3 or more vehicle body lengths (the near-field). The second step propagates the signature from the near-field down to the ground (the far-field) and computes various loudness metrics, using a dedicated propagation solver. Decomposing the problem in this way has been used with success on other supersonic geometries and, for example, on the X-59 and its previous design iterations [2]. In Section 2, this standard two step loudness prediction procedure using the Launch, Ascent, and Vehicle Aerodynamics (LAVA) CFD framework is presented, in particular by using its structured curvilinear grid paradigm. There, a loudness prediction for a demonstrative flight condition of the X-59 will be made, and the grid, computational domain and methodology, and computational cost will be presented. The LAVA framework has been thoroughly validated against experimental wind tunnel results from the NASA Ames 9x7 ft. Supersonic Wind Tunnel [3] and the NASA Glenn 8x6 ft. Supersonic Wind Tunnel [4], as well as in supersonic jet noise predictions [5, 6].

The X-59 is an experimental aircraft being designed and built by industry partners in collaboration with NASA, shown in Figure 1. It has been designed to minimize the sonic boom generated when flying over land at supersonic speeds. Throughout its design process, and going forward into its first flight tests, simulation teams at NASA and in industry have built simulation workflows to analyze and predict the X-59's sonic boom loudness for a given flight configuration. It was designed around a target loudness of 75 dB for a particular baseline flight configuration, but due to the complex flow physics involved, certain changes in the flight configuration can result in louder or softer sonic booms on the ground. It is the role of the simulation teams to rapidly respond to changing requirements to make robust and accurate loudness predictions as the first flight of the X-59 approaches. There is also an ongoing collaboration between NASA teams to develop a flight planning software that requires large databases of training data in the form of near-field pressure signals for a wide variety of potential flight configurations. This software is responsible for generating flight conditions for a given desired loudness, so it requires high-fidelity, representative near-field solutions for reliability.



Figure 1: Artist rendering of the X-59 flying over land.

Several advancements on the traditional two step loudness prediction procedure have been developed to meet the increased accuracy and robustness requirements, such as introducing a third step by way of a high-order Space Marching (SM) method for mid-field computations, and developing a scripting system using curvilinear grids to increase robustness and simplify the process of running large databases of simulation cases. These will be introduced in Section 3 and applied to the X-59 in Section 3.3. Comparisons of the near-field and far-field results and costs between the original two step workflow and the new three step procedure will be made, and more robust and accurate predictions with a reduced computational cost are obtained using the new three step procedure.

High fidelity loudness predictions also require a well-designed CFD mesh to capture the shock-dominated flowfield. This typically requires careful placement of grid refinement regions in locations of known shock-inducing geometric features, with the capability of transforming depending on the flight condition. To address this, a novel output-based mesh adaptation method targeting error in near-field pressure signatures has been developed and it will be explored and applied to X-59 loudness prediction in Section 4.

2 Two Step Solution Procedure

In order to predict the loudness of any particular X-59 flight configuration, the problem is decomposed here into a two step procedure. First the flowfield around the vehicle is computed using the LAVA [7] Curvilinear CFD solver with a steady RANS formulation on a structured overset mesh that extends radially out from the vehicle to a predetermined distance. The pressure disturbance $dP/P_\infty = (P - P_\infty)/P_\infty$ is interpolated onto computational line probes at this distance at regularly spaced azimuths around the vehicle, forming an enclosing extraction cylinder. These pressure disturbance signals are then propagated to the ground using a specialized far-field propagation solver sBOOM [8], which solves the augmented Burgers' equation along acoustic ray paths.

sBOOM assumes that azimuthal correlations in the near-field pressure signals are negligible. This assumption only holds when the extracted pressure signals are sufficiently far enough away from the vehicle such that azimuthal variations have attenuated. The X-59 vehicle uses a standard length $L = 90$ feet to represent the length of the vehicle, and the radial distance from the vehicle centerline is expressed with R . Previous sonic boom prediction efforts such as the third Sonic Boom Prediction Workshop (SBPW3) [3] have used $R/L = 3$ as a standard location for the transition from near-field CFD to far-field propagation for earlier design iterations of the X-59; it has been verified to be at a sufficient distance to ensure azimuthal independence of each off-track angle.

Flight conditions specifying a flight Mach number of 1.4, an altitude of 53,200 feet, and an angle of attack of 2.04 degrees will be used as a representative flight condition for the X-59 for the remainder of this section.

2.1 LAVA Curvilinear Grid Generation

X-59 simulations in the LAVA Curvilinear module use structured overset meshes. These meshes are constructed to conform to the X-59 CAD geometry used in this study, shown in Figure 2. Though mostly span-wise symmetrical, the vehicle includes some vents and cameras that are not mirrored across the centerline. Some initial geometry cleanup in ANSA [9] was required in order to close gaps and create a watertight surface. Overlapping structured surface meshes were then created in Pointwise [10], shown in Figure 3. Meshing guidelines followed the best practices learned from creating similar meshes on earlier design iterations of the X-59, like the C608 design simulated in SBPW3 [3]. 0.1% local chord spacing was specified for leading edges, at least 5 points across blunt trailing edges was ensured, and the global maximum stretching ratio governing surface and volume growth was set to 1.25 for the coarsest grid level. A minimum 5 points of overlap between adjacent surface domains was also ensured for adequate communication between fringe points of overlapping grids. Chimera Grid Tools [11] is then used to grow the volume grids hyperbolically from the body-fitted surface grids. When growing from viscous surfaces, the initial cell height is designed to maintain $y^+ \approx 1.0$. The full X-59 grid system consists of 774 grids. The vehicle geometry is rotated to the required angle of attack in place so that the off-body grids can remain aligned with the Mach angle for any given flight configuration, with the incoming flow directionally fixed at $\alpha = 0^\circ$. Grids were also carefully segmented into modular groups for later manipulation, which will be discussed in Section 3.2.1.




Figure 2: Several views of the X-59 geometry.




Figure 3: Surface grids from Pointwise, exported and viewed in CGT.

become relevant in Section 4.

All surface meshes were constructed at their coarsest level before uniform mesh refinement was applied. Following the results observed in [3], a refinement level of 1.4 was determined to be sufficient for accurate loudness predictions. The final grid on the full CFD domain extending out to $R/L = 3.1$ has 545 million nodes and takes about 35 minutes to generate on a modern high-end workstation.

Off-body grids are generated using a quasi-1D method, where nonuniform refinements are placed only in the axial direction. The grids are uniform in the circumferential direction, and expand outward using the global stretching ratio ($SR = 1.25$) radially. An axial point distribution is defined on the centerline, where additional points can be placed to generate refinement regions for known shock-inducing parts of the geometry. Such refinement locations include the forward probes and cameras, forward and aft vents, antennas, and a large region for the aft deck and plume. This centerline is then extruded along the Mach angle out to the desired off-body distance, with a small offset of 0.2° applied to compensate for skewness in the Mach angle introduced by numerical artifacts. The radial extent for off-body grids is chosen as $R/L = 3.1$ to ensure the near-field pressure probe extraction location is sufficiently enclosed by the domain. This surface is then revolved to form an enclosing cylindrical domain. Due to memory constraints, the off-body cylinder is split into 8 point-matched abutting blocks, which will be

2.2 Problem Setup

The LAVA Curvilinear solver uses a second-order accurate finite-difference discretization to the steady-state compressible RANS equations in strong conservation law form. The convective fluxes are evaluated using a modified Roe scheme with third-order upwind-biased flux reconstruction [12, 13] and utilize the Koren flux limiter. The viscous fluxes are discretized using a mid-point and node-centered difference scheme. For consistency with NASA best practices and other collaborating CFD teams, the turbulence model selected for RANS closure is the one-equation Spalart-Allmaras (SA) turbulence model [14], where the convective terms are discretized using first-order upwinding. The discretized set of non-linear equations is marched in pseudo-time to steady-state using automatic CFL ramping to accelerate convergence. At each pseudo-time step, the linear system is solved using the preconditioned generalized minimal residual (GMRES) algorithm [15]. The preconditioner is obtained from an incomplete lower-upper factorization with fill level 0 of a first-order approximation of the residual Jacobian.

The flow in the engine from the intake to the exit nozzle is not modeled as a flow-through engine. Instead, a boundary face is placed inside the engine where the compressor would be, and a special Mach boundary condition designed during SBPW3 [3] is used. For this boundary condition, only the incoming Mach number is specified. The solver then attempts to dynamically find a single pressure value across the face such that the target Mach number is achieved in the average sense. The outlet of the engine is another boundary face in the plenum of the exit converging-diverging nozzle, where stagnation pressure and temperature are specified. This same cut-off strategy and boundary condition implementation is also used at the inlets and outlets of the environmental control system (ECS).

The boundary conditions at the engine and ECS inlets can sometimes develop into supersonic inlet conditions if started from purely ambient conditions. To prevent this, the flow-field is always initialized such that any solution node belonging to a grid inside the inlet regions is set to the post-shock results from basic normal shock relations at the flight condition Mach number. Additionally, solution instabilities can occasionally develop near the ECSs due to their geometric complexity and the potential difficulty of converging the Mach boundary condition; to alleviate these, the solver initially runs for 2000 nonlinear iterations using a more dissipative first order flux reconstruction, then switching to third order for the remainder of the iterations. The four ECS vents, v-tail vent, bottom drain, and sub-nozzle engine bypass have stagnation pressure and temperature supplied.

2.3 Near-field Pressures and Loudness Prediction

Figure 5 shows a visualization of the flowfield on a centerline slice in the aft portion of the flow, wherein the loudness prediction exhibits great sensitivity. Figure 6a shows the convergence behavior comparing the normalized L_2 norm of the flow equation residuals as a function of nonlinear iterations, highlighting the CFL ramping capability of LAVA Curvilinear. CFL ramping is prescribed to start from unity to a maximum of 150, which helps to accelerate convergence. A spike in the residuals can be seen at 2000 nonlinear iterations corresponding to when the simulation switches from 1st to 3rd order convective flux evaluations. The solution took 11.5 hours on 60 Electra Skylake nodes (2400 Intel Xeon Gold 6148 processors).

Pressure signatures in the form of dP/P_∞ are extracted from the volume solution at $R/L = 3$. The azimuthal angle ϕ is designated “on-track” when $\phi = 0^\circ$ underneath the vehicle. In this work, we extract the entire cylinder at an azimuthal resolution of 2.5 degrees, as shown in Figure 4. On-track near-field pressure signatures are given in Figure 6b.

Propagation was performed with sBOOM. The near-field signatures are pre-processed to have leading zeroes ahead of the nose shock and to linearly ramp the pressure from the downstream end of the probe back to ambient conditions to improve propagation consistency. Each on- and off-track angle is propagated using a sampling frequency of 350kHz in standard atmosphere without wind. Figure 7 shows the resulting loudness predictions for the boom carpet together with the on-track ground signature. Loudness predictions are most often provided using the perceived loudness (PL) level of Stevens [16] in decibels. The on-track PL on the ground is predicted at 76.8 dB, slightly louder than the target loudness of 75 dB, with off-track angles generally decaying in loudness.

The ground signature shown in Figure 7b highlights several ways in which the design of the X-59 mitigates loudness on the ground by shaping the boom signature; for instance, the long nose and canards help to break up the pressure rise, while the stabilator deflection interacting with the rounded bumps on the bottom of the vehicle gentle the slope of the pressure recovery. In fact, the under-pressure region of the ground signal is a crucial component in estimating the loudness of the X-59; significant loudness buildup using the B-weighted sound exposure level (BSEL) loudness metric can be observed




Figure 4: Line probe cylinders shown at $R/L = 1, 2, 3$. Starboard probes omitted for visual clarity.






Figure 5: Centerline slice of the volume solution, showing density normalized by free stream. Note where waves emitting from the geometry impinge on the nozzle exhaust plume and their resulting pressure disturbance underneath it.

near $t = 110\text{ms}$. A breakup of the wave just before significantly lowers the perceived loudness over other conditions. Phenomena like this can be traced back to the near-field probe, in this case in Figure 6b, near $x/L = 4.1$, and then traced back up the Mach angle to the near-field solution, here originating in the exhaust plume. This region is very complex involving not only the engine plume, but also shocks coming off the control surfaces atop the vertical tail, and also involving the pressure recovery portion after the stabilators, from $x/L = 4$ to $x/L = 4.1$. Refer to Figure 5 to see the region.




(a) Near-field CFD convergence.




(b) On-track pressure disturbance dP/P at $R/L = 3$.

Figure 6: Near-field CFD solution for given flight conditions.



(a) Perceived loudness (PL) in decibels over the boom carpet.



(b) Ground signature (in black) and loudness buildup (in blue).

Figure 7: Noise characteristics on the ground as propagated by sBOOM.

3 Advancements in Loudness Prediction Methods Part I: Space Marching and Database Capabilities

Recently, several improvements have been made to the two step noise prediction procedure within the LAVA computational framework. A high-order space marching method has been developed, yielding significant efficiency and accuracy improvements [17], which are covered in Section 3.1. Also, a robust scripting system covered in Section 3.2 has been developed that leverages the use of overset curvilinear grids to modularly re-appropriate or generate new grids based on any given flight condition, reducing user error and streamlining the production of flight simulation databases. Mesh redistribution is detailed separately in Section 4.

3.1 Space Marching

Space marching is a highly efficient middle step in the loudness prediction process, shown in Figure 8, that takes advantage of the supersonic nature of the flow by solving the steady state Euler equations

of gas dynamics on a Mach-cone aligned mesh. The space marching domain covers the region of space termed the “mid-field” between the interface location near the vehicle geometry and the interface location suitable for sBOOM ($R/L = 3$ in this work).




Figure 8: Introducing mid-field space marching to loudness predictions.

The near-field solution is computed as before in Section 2, but the CFD domain of the near-field problem can be truncated to just the domain of dependence. A simple way to estimate the domain of dependence is to construct a forward opening Mach cone placed just upstream of the nose of the vehicle and a Mach cone opening in the opposite direction placed near the aft of the vehicle. The placement of these cones along the streamwise axis can be adjusted until the entire vehicle geometry is encompassed by the two cones, taking care to include the full wing span and vertical tail. The radius where the cones intersect provides a good estimate of the required radial domain extent. For the X-59, the cones can be seen in Figure 9, and the intersection be found as $R/L = 0.6$.

In order to verify that truncating the domain in this way does not affect the near-field pressure disturbances, Figure 10 shows the pressure disturbance at $R/L = 0.25$ interpolated from the default near-field CFD domain of Section 2, which extends out to $R/L = 3.1$ to include the pressure probes, and on the truncated domain for use with space marching extending to $R/L = 0.6$. Visually, there is no difference in the solution. Figure 11 interpolates a line probe underneath the vehicle at the same radial extent $R/L = 0.25$, and again one can see no visible difference in the pressure disturbance for most of the signal, indicating that domain truncation indeed does not affect the near-field solution. A small difference can be seen at the downstream end of the signal due to ignoring the influence of the powered engine plume in the Mach cone domain of dependence analysis, however this region has a negligible influence on the loudness buildup (see Figure 7b). Additionally, truncating the CFD domain in this way reduces the size of the more costly RANS CFD problem from solving on a grid with 545 million nodes, 412 million of which are in the off-body grid, to the solution on a grid with 273 million nodes, 131 million of which are off-body, representing a 50% reduction in solution points.

After finding the minimum CFD domain extent, an interpolating elliptical surface is then setup around the vehicle, according to Equation 1.

$$\sqrt{a_0(y - y_0)^2 + b_0(z - z_0)^2} = 1 \quad (1)$$

This surface trims the interior portion of the space marching grid that is covered by the near-field CFD domain, effectively removing those grid points from the space marching solution procedure. Currently, the elliptical interface is determined manually; automatic methods of construction are currently being explored. The ellipse must be carefully constructed such that it is within the truncated CFD domain, and captures the vehicle geometry without cutting too close to it. For all X-59 cases in this work, $a_0 = 0.0325$, $b_0 = 0.15$, $y_0 = z_0 = 0$. See Figure 12 for a visualization of the elliptical cutting surface.

The space marching solver uses a specially designed structured curvilinear grid with (ξ, η, ζ) corresponding to (j, k, l) indices to discretize the mid-field. The grid aligns its ζ -direction with the freestream flow and uses a uniform spacing, reducing potential dispersion error. The ξ -direction is aligned with the Mach angle $\mu = \sin^{-1}(1/M_\infty) + \mu_\epsilon$ where $\mu_\epsilon = 0.2^\circ$, and is governed by a simple stretching rule $SR = 1.05$ with a prescribed maximum aspect ratio $AR = 20$. Finally, the η -direction covers the circumference with a uniform spacing.




Figure 9: Domain of dependence.




Figure 10: Near-field CFD solution at $R/L = 0.25$. No visible differences in solution observed from domain reduction.

Pressure probes are typically most sensitive to streamwise spacing. To evaluate this, a grid sensitivity study was performed on the same flight conditions as in Section 2. Here, the space marching domain in the streamwise direction is fixed at $2L$, starting just ahead of the vehicle nose. The finest mesh level has $n_\zeta = 1001$ where n_ζ is the number of streamwise points. Coarser grid levels halve the spacing. Circumferential resolution is held constant at 1 point per degree for simplicity ($n_\eta = 361$). See Figure 13 for a visualization of the on-track line probes interpolated at $R/L = 3$. At even the finest space marching mesh level, the cost of the computation is negligible compared to the cost of the near-field CFD solution, so the finest grid level with $n_\zeta = 1001$ will be used.




Figure 11: Line probes extracted at $R/L = 0.25$. In black is the default CFD domain extending out to $R/L = 3.1$. In blue is from the truncated CFD domain for use with space marching, which only extends to $R/L = 0.6$.




Figure 12: Construction of CFD-SM interface location.

n_ζ	125	251	501	1001
total dof	7.81M	28.4M	106M	411M
wall time	1m	2m	5m	18m

Table 1: Space Marching grid resolution study using a single Electra skylake node.

Because of the hyperbolic nature of supersonic flows, space marching solves the Euler equations in a time-like marching method in the ζ -direction, using backward differencing (BDF2). For the two space-like directions, the derivatives are discretized using a high-order Hybrid Weighted Compact Nonlinear Scheme (HWCNS) [18, 19], with fluxes given by a WENO [20] scheme. Several versions of these schemes are available in the space marching solver; in this work, we use HWCNS4-WENO3.




Figure 13: Space marching streamwise grid refinement study at $R/L = 3$

3.2 Database Capabilities

In addition to providing flight planning teams with fast and accurate loudness predictions, the LAVA team has also helped develop and execute several iterations on databases of X-59 flight conditions in order to provide the Acoustics Verification, Testing, and Planning (AVTPE) team with a large amount of data with which to train an as-yet unpublished flight planning machine learning model called CLEOPATRA. Several new capabilities were introduced to the X-59 simulation process in order to execute these databases efficiently, and in such a way that was less prone to human error. These capabilities extend their benefits to any current and future loudness prediction simulations, whether individual cases or large sets of cases, by way of a Python scripting system developed specially for this work. These database scripts control the full simulation process from grid generation to postprocessing, and enable users to perform X-59 loudness predictions with consistency and reliability.

3.2.1 Build-from-nominal Grid Process

One advantage of using structured overset meshes for sonic boom predictions is their modularity. When constructing a series or database of simulations, often the changes required from one flight configuration to the other are quite small in proportion to the full grid system. For instance, a stabilator sweep only requires meshes on the stabilator itself to rotate around the deflection axis. Regrowing the volume grids from the surface can be foregone if there is good coverage of the space; the existing stabilator mesh from the previous deflection point in the sweep can simply be rotated and overlaid into the rest of the grid system. In this way, only the grids affected by a change in the simulation inputs need be altered, often by simple transformations, and the rest of the volume grids can be copied over from the previous

grid to the new grid system. Constructing X-59 meshes in this way reduces the total time to generate a grid on a modern high-end workstation from 35 minutes down to 10 minutes, with no difference in the resulting grid system. Most of the 10 minutes spent in grid generation is in the overlapping and hole-cutting procedure to generate *iblack* information and interpolation stencils. To facilitate this process, overset meshes were carefully constructed in Pointwise and CGT such that control surfaces, like the ailerons, flaps, stabilators, rudder, and t-tail, highlighted in Figure 14, can be dynamically deflected. All of the highlighted components can be rotated around predetermined axes to some reasonable maximum deflection; or in the case of the nozzle, opened and contracted to some reasonable exit plane area.




Figure 14: Control surfaces able to be deflected in database scripts.

3.2.2 Dynamic nozzle and control surface deflections

Recent versions of the X-59 simulation database featured variable inputs for the throat and exit areas of the vehicle's engine, which in reality uses hydraulics to control the nozzle areas to achieve different thrust conditions. To account for this, the interior nozzle grids, which are point matched abutting grids built on the original geometry, are dynamically generated via CGT grid scripts. Exit plane widening and narrowing is usually small compared to the exit plane area, so the streamwise coordinate of the exit plane is kept constant, and the lip is kept at a constant width. This defines the top coordinate of the nozzle exterior. A root point is fixed to the geometry which then defines the boundary points of the exterior surface. The exterior is then formed by scaling the original nozzle curve by its distance from the fixed root point. This method maintains the existing surface curvature of the original, undeflected nozzle, and is a small improvement over using straight lines for the exterior surface. Most surface grids around the nozzle can then be projected onto the deformed surface.

3.3 Three-Step Solution Procedure

In this section, results from simulating the X-59 at the same representative flight conditions using the database scripting system for grid generation and preprocessing, and using space marching for mid-field computations will be presented. Comparisons with the results of Section 2 will also be highlighted.

3.3.1 Problem Setup

The grid for this case is generated from a previous run condition (in this case, using a geometry from a different angle of attack and different control surface deflections) following the build-from-nominal grid process. The multi-block off-body grids are extended out to $R/L = 0.6$, following the domain of dependence analysis from Section 3.1 and observed results from [17]. The space marching grid has 411 million nodes in total and it can be seen in Figure 15, where the center of the domain has been trimmed by the elliptical cutting surface. The solution is computed as before for 2000 nonlinear iterations using more dissipative first order convective fluxes before restarting with a third order scheme and running for a fixed 10000 total nonlinear iterations. The near-field solution is then interpolated onto the space marching grid and propagated to mid-field. Line probes are extracted at the $R/L = 3$ location from the space marching volume solution and given to sBOOM for propagation to the ground. sBOOM uses a frequency resolution of 350k in standard atmosphere and no wind.




Figure 15: Centerline slice of space marching domain, shown on a coarse resolution for visual clarity.

3.3.2 Results

Line probes extracted 3 body lengths away are shown on the space marching domain in Figure 16. They are indistinguishable for the most part, with the biggest discrepancies occurring in the downstream portion. The inset image focuses on this area. As observed in Section 2, this can be traced back to the plume region, where shocks from the t-tail impact the flow in the plume and continue along the Mach angle. A small difference in the freestream recovery region can also be observed, most likely due to more aggressively truncating the CFD domain such that turbulent and viscous features far downstream of the vehicle are being communicated to space marching and kept in the mid-field propagation. Ultimately, this has a negligible effect ($< 0.1dB$) on the predicted loudness. The near-field probes for positive off-track angles are collected into contour plots in Figure 17 for comparison. Visually, there is very little difference in the near-field solutions between the two methods.

Figure 18a shows predicted loudness over the boom carpet in dBs. Across the carpet, the pressure disturbance obtained using space marching consistently predicts higher loudness than the probes obtained using CFD alone. This is expected as space marching is much less dissipative than the RANS CFD in the mid-field. Space marching also introduces less dispersion error into the flow by having a refined uniform mesh in the streamwise direction without steps in resolution coming from refinement regions, like what is present in the CFD solution.

In Figure 18b, we observe that space marching maintains a slightly sharper initial rise to the pressure disturbance peak, and a difference in the under pressure peak characteristic which causes a slight step up in the loudness buildup, visualized with the BSEL loudness metric. The underpressure peak in the space marching solution is slightly higher in amplitude, and also the pressure recovery slope is a little sharper, contributing to a higher loudness buildup. This area is the most common location for the X-59 in loudness prediction differences across analysis teams due to how many physical effects are involved, including the plume structure (which itself includes engine conditions and nozzle geometry sensitivity), v-tail and t-tail wave propagation from the top of the geometry to the plume, and three dimensional effects from the aft deck expansion interaction with the plume. Using space marching allows for CFD domain truncation and cost reduction, enabling more refinement to be given to the plume and t-tail




Figure 16: Near-field CFD solution using space marching for mid-field propagation.




Figure 17: Pressure disturbance contours for positive off-track angles.


regions. As before, the ground signature and loudness buildup past $t = 120\text{ms}$ show that loudness predictions have little sensitivity to the signal closeout.

Loudness metrics	ASEL	BSEL	CSEL	PL
CFD	63.172	76.057	90.467	76.781
Space Marching	63.443	76.509	90.901	77.167


Table 2: On track loudness metrics compared between the two methods.

In Table 2 we compare the various loudness metrics (A-, B-, and C-weighted sound exposure levels and PL) for the two procedures directly under the X-59 at $\phi = 0^\circ$. Across the board, space marching predicts a slightly higher loudness, with PL being 0.38 dB higher.

Table 3 compares computational cost in CPU-hours for the traditional two step procedure of Section 2 and the current best practice three step procedure using space marching. We find a 53% cost reduction for a loudness prediction that is more accurate than before. CFD domain truncation comprises the



(a) Perceived loudness (PL) in decibels over the boom carpet.



(b) Ground signature (solid lines) and BSEL loudness buildup (dashed lines)

Figure 18: Far-field propagation of pressure signatures from space marching domain using sBOOM.

CPU-hours	2 step	3 step
Grid Generation	0.50	0.16
Near-Field Solution	690	400
Space Marching	—	0.28
sBOOM	0.1	0.1
Total	690.6	400.54

Table 3: Computational cost comparison between two step (CFD, sBOOM) and three step (CFD, SM, sBOOM) procedures.

majority of cost savings between the two methods, while the more efficient grid structure and solution marching procedure from the near-field to the mid-field provide accuracy enhancement and enable space marching to be run using a single Skylake node.

4 Advancements in Loudness Prediction Methods Part II: Anisotropic Mach Cone Aligned Mesh Redistribution

Structured curvilinear Mach cone aligned grids are used within both the CFD and space marching domains. These grids align the radial computational space direction with the Mach cone angle to minimize grid induced errors when capturing the shock-dominated off-body flow field. Their extents are defined by the domain of dependence which may be manually or automatically determined through intersecting Mach cones (see Figure 8). Previous work highlighted an automated procedure for constructing these grids from a small set of user parameters [21, 22]. For database generation, only the Mach cone angle need be redefined based on the baseline flow conditions. Generation of this mesh is negligible, much less than a minute, for most production-level grid systems. The true difficulty comes in automatically defining refinement regions which provide sufficient resolution of the off-body flow features required for accurate ground-level noise predictions. For such automation, we rely on a novel anisotropic Mach cone aligned mesh adaptation technique.

Adaptation is achieved through an r-refinement algorithm, which fixes the number of grid points and redistributes them to coarsen or refine certain regions of the grid. This adaptation approach can be viewed as a method for generating an optimal off-body grid for a given number of degrees of freedom, and for the current application a fixed mesh orientation. By fixing the degrees of freedom we are able to generate a more accurate solution with a fixed simulation cost. Consider the integration of the pressure

disturbance waveform at any radial location in the continuum,

$$J(p) := \int \frac{\Delta p}{p_\infty} d\ell \quad (2)$$

In the asymptotic range, the discrete integration of an extracted waveform satisfies

$$J_h = J_{\text{exact}} + C_p h^p + \mathcal{O}(h^{p+1}) \quad (3)$$

where J_h is the approximation of J_{exact} from a quadrature rule, p is the order of the CFD discretization, C_p is a grid-independent constant, and h is the characteristic grid spacing in the streamwise direction. Therefore, by locally refining the off-body grid in regions where the discretization error is large, we reduce the total error in our waveform prediction.

Governing the adaptation is a one-dimensional error equidistribution principle of the form

$$w_i \Delta s_i = C \quad (4)$$

for $i = 1, \dots, N$. Here, N is the number of grid points along the one-dimensional curvilinear coordinate line, w_i is a positive weighting function, Δs_i is the length of the i^{th} interval in the grid, and C is an arbitrary constant. The weighting function is defined as

$$w_i = 1 + A f_i \quad (5)$$

where f_i is the error indicator which drives the redistribution and A is a user-specified clustering constant. Methods for automatically defining A are discussed by Ashby et al. [21, 22]. Since information travels along characteristic lines through the off-body which is aligned with the radial direction of the Mach cone aligned grids, we restrict redistribution in either the streamwise or circumferential directions using Equation (4). Both feature-based and adjoint-based indicators may be used to drive the adaptation.

Feature-based indicators are simple to construct from a computed flow solution, however, they may target regions in the off-body domain which may not directly contribute to the error in the waveform. For this reason, adjoint-based indicators may be more suitable to only refine regions contributing directly to the error. We have observed little benefit from solving large adjoint systems particularly for complex geometries where exact adjoint methods are twice as expensive as a flow solve. Comparison between a first-order adjoint method relying on finite differencing to construct the numerical Jacobian to feature-based indicators in LAVA were published by Ashby et al. [21, 22].

4.1 Feature-based Indicators

Feature-based indicators are an error measure for a given flow quantity. Here we present a simple and effective feature-based error indicator based on the work of Buning and Pulliam [23]. The one-dimensional undivided second-difference of a flow quantity q is computed at point i

$$\frac{\partial^2 q}{\partial \xi^2} = q_{i+1} - 2q_i + q_{i-1} \quad (6)$$

This is a computational coordinate biased proxy for q at grid point i and is proportional to the remaining error term of the linear interpolation of q from neighboring points on a two times coarser grid level. Common examples of q for supersonic flow include Mach number and pressure disturbance. For a given computational coordinate direction, non-blanked points are computed

$$e(\xi_j, \eta_k, \zeta_l) = \max_{i=j,k,l} \frac{|q_{i+1} - 2q_i + q_{i-1}|}{q_{\text{ref}}} \quad (7)$$

where q_{ref} is a reference quantity used to normalize the indicator.

Feature-based indicators are advantageous for adaptation since they only require a flow solution. For CFD simulations, the discretization error is convected and diffused throughout the domain. Since feature-based indicators target error and not the sources of error, adaptation methods will resolve regions of the domain that do not impact the discretization error. For this reason, adjoint-based error indicators are a popular alternative.

4.2 Adjoint-based Indicators

In CFD simulations, output functionals such as lift and drag are often the quantities of interest. Adjoint-based indicators are formulated from estimates of output functional discretization error estimates. The derivation of these estimates follow from discrete adjoint-weighted residuals for finite-volume and finite difference discretizations for near-field CFD simulations. For brevity, we present a condensed formulation and refer readers to the work of Venditti et al. [24, 25].

Consider an open bounded domain $\Omega \subset \mathbb{R}^d$, $1 \leq d \leq 3$. We denote the output quantity of interest as $J(\mathbf{Q})$, where $\mathbf{Q}$ is the analytic solution to a set of model equations,

$$\mathbf{R}(\mathbf{Q}) = \mathbf{0} \text{ in } \Omega \quad (8)$$

For this work we apply the RANS equations in strong conservative form where Ω is the near-field CFD domain. Suppose Ω has been discretized to form a working mesh Ω_H with characteristic cell-size H and N vertices. The flow equations are then discretized and form a nonlinear discrete system of partial differential equations,

$$\mathbf{R}_H(\mathbf{Q}_H) = \mathbf{0} \quad (9)$$

where $\mathbf{Q}_H = [\mathbf{Q}_1, \dots, \mathbf{Q}_N]^T$, $\mathbf{Q}_i \in \mathbb{R}^M$, and $\mathbf{R}_H \in \mathbb{R}^{M \cdot N}$. If J was an integral operator, then J_H represents the quadrature rule used to compute J on Ω_H . An error estimate for the computation of the output functional discretization error is sought,

$$\varepsilon = |J(\mathbf{Q}) - J_H(\mathbf{Q}_H)| \quad (10)$$

Let Ω_h be the embedded uniformly refined mesh with characteristic cell-size satisfying $h = H/r_f$ where $r_f \in \mathbb{R}_+$ is the refinement factor. In the asymptotic range, Richardson extrapolation gives relation

$$J(\mathbf{Q}) - J_H(\mathbf{Q}_H) = \frac{r_f^p}{r_f^p - 1} (J_h(\mathbf{Q}_h) - J_H(\mathbf{Q}_H)) + \mathcal{O}(h^{p+1}) \quad (11)$$

Therefore, to estimate the error in the coarse grid functional estimate, we need only to form an error estimate between embedded and working space functionals.

The discrete adjoint (or dual) problem is defined as the linear system of equations,

$$\left[\frac{\partial \mathbf{R}_h}{\partial \mathbf{Q}_h} \Big|_{\mathbf{Q}_h^H} \right]^T \psi_h = \left[\frac{\partial J_h}{\partial \mathbf{Q}_h} \Big|_{\mathbf{Q}_h^H} \right]^T \quad (12)$$

where $\mathbf{Q}_h^H$ is a prolongation of the working mesh flow solution into the uniformly embedded space Ω_h . Taylor expansions of the embedded space flow residual and output functional about the prolonged flow solution give an error estimate,

$$J_h(\mathbf{Q}_h) - J_h(\mathbf{Q}_h^H) \approx -\psi_h^T \mathbf{R}_h(\mathbf{Q}_h^H) \quad (13)$$

Adding and subtracting prolonged coarse grid adjoint solutions gives

$$J_h(\mathbf{Q}_h) - J_h(\mathbf{Q}_h^H) \approx -\left(\psi_h^H\right)^T \mathbf{R}_h(\mathbf{Q}_h^H) - \left(\psi_h - \psi_h^H\right)^T \mathbf{R}_h(\mathbf{Q}_h^H) \quad (14)$$

To avoid solving the large sparse adjoint linear system on the embedded mesh, we apply low- and high-order prolongation operators to the working mesh primal and adjoint solutions, $\mathbf{Q}_h^{LO} = \mathbf{P}_{LO} \mathbf{Q}_H$, $\psi_h^{LO} = \mathbf{P}_{LO} \psi_H$, and $\psi_h^{HI} = \mathbf{P}_{HI} \psi_H$. Replacing ψ_h , ψ_h^H , and $\mathbf{Q}_h^H$ in (14) gives

$$J_h(\mathbf{Q}_h) - J_h(\mathbf{Q}_h^{LO}) \approx -\left(\psi_h^{LO}\right)^T \mathbf{R}_h(\mathbf{Q}_h^{LO}) - (\psi_h^{HI} - \psi_h^{LO})^T \mathbf{R}_h(\mathbf{Q}_h^{LO}) \quad (15)$$

Substituting this approximation into (11),

$$J(\mathbf{Q}) - J_H(\mathbf{Q}_H) \approx \frac{r_f^p}{r_f^p - 1} (J_h(\mathbf{Q}_h^{\text{LO}}) - J_H(\mathbf{Q}_H)) \quad (16)$$

$$- \frac{r_f^p}{r_f^p - 1} \left[\left(\psi_h^{\text{LO}} \right)^T \mathbf{R}_h(\mathbf{Q}_h^{\text{LO}}) + (\psi_h^{\text{HI}} - \psi_h^{\text{LO}})^T \mathbf{R}_h(\mathbf{Q}_h^{\text{LO}}) \right] \quad (17)$$

The output-based error estimate on the embedded off-body grid is defined to be the error in the computable correction,

$$e_h = (\psi_h^{\text{HI}} - \psi_h^{\text{LO}})^T \mathbf{R}_h(\mathbf{Q}_h^{\text{LO}}) \quad (18)$$

A full-weighting restriction operator is applied to e_h to produce the adaptation indicator on the working off-body mesh.

Adjoint-based indicators are advantageous for improving mesh resolution to accurately capture the output quantity of interest. A downside of this approach is the associated computational cost of obtaining an adjoint solution. Although prolongation operators are utilized to mitigate these costs, adjoint solutions may cost twice as much as the flow solution. The resulting mesh may result in a significant reduction in grid nodes, potentially offsetting this cost for database generation [21, 22].

4.3 Direction-based Indicators

To ensure direction-based redistribution produces high-quality volume grids, we apply a novel procedure to the aforementioned error indicators. Let $e_0 = e_0(\xi, \eta, \zeta)$ be an initial error indicator computed on the Mach cone aligned off-body grid where ξ , η , and ζ are the streamwise, circumferential, and radial computational grid axes respectively. Since supersonic flow information travels in a common time-like direction, we form a new indicator from e_0 such that the adaptation will generate locally refined spacing along the time-like direction while maintaining required mesh alignment and smoothness. Without loss of generality, we proceed by outlining the procedure for forming a streamwise adaptation indicator. An initial streamwise adaptation indicator g is defined

$$g(\xi) = \max_{\eta} |e_0(\xi, \eta, \zeta^*)| \quad (19)$$

where ζ^* is the cylindrical surface with minimum radial distance from the aircraft that is not blanked by the near-body grid system. This ensures the capturing of the errors in the circumferential direction at each streamwise location on the fixed radial surface. An elliptic smoothing algorithm is applied to the indicator,

$$g_i^s = \frac{g_{i-1} + 4g_i + g_{i+1}}{6} \quad (20)$$

where $g_i = g(\xi_i)$ [26]. Applying this algorithm for a low number of iterations reduces numerical noise and sets values in the blanked points with reasonable approximations for adaptation. Let $\bar{g}$ be the average of the smooth streamwise indicator. Since $g^s \geq 0$ and we require $g^s \in [0, 1]$ in order to specify the clustering constant A in (5), we define the mapping

$$f(\xi, \eta; \zeta^*) = \tanh(\beta g^s(\xi)) \in [0, 1] \quad (21)$$

where $\beta = \tanh^{-1}(0.8)/(C \bar{g})$ and $C \geq 1$ is a threshold constant ensuring the targeting of statistically significant regions of the error indicator. Once the surface indicator $f(\xi, \eta; \zeta^*)$ is defined, we interpolate its values onto the rest of the grid using a normalized distance metric

$$\tau = \frac{x - r\sqrt{M_\infty^2 - 1}}{L} \quad (22)$$

where r is the radial distance from the aircraft centerline and L is the aircraft body length.

4.4 Redistribution Verification

The goal of mesh redistribution is to reduce the grid-induced error in the CFD off-body domain without increasing the degrees of freedom. If we are targeting important regions of the flow, applying mesh redistribution to the coarse grid and refining the resultant grid system should result in improved solutions through all mesh levels. To verify our methodology, we consider a simple extruded diamond airfoil CFD simulation with a truncated off-body domain. A two-dimensional diamond airfoil with a thickness of 0.07 is simulated in a supersonic inviscid flow with a Mach number of 2 [27]. The goal of this study is to investigate the impact of adaptation exclusively on extracted off-body CFD signatures. Our adaptation methodology relies on the assumption that off-body grid errors in the CFD simulation contribute significantly to the error in the near-field space marching solution. These errors accumulate and propagate to the ground-level noise metrics. Therefore, we use this test case to demonstrate that we can reduce off-body CFD simulation errors.




Figure 19: Initial (left) and adapted (right) diamond airfoil grids comparisons.




Figure 20: Diamond airfoil line probe (red) at $R/L = 0.35$

The off-body grid is generated using 301 streamwise points, 101 radial points, and constant radial spacing 0.1 meters from the geometry. A family of grids is created using a refinement factor of 1.4, resulting in five grid levels. A line probe of $\Delta p/p_\infty$ is extracted at $R/L = 0.35$. We use the fifth grid




Figure 21: Diamond airfoil line probe convergence at $R/L = 0.35$.




Figure 22: Waveform error convergence for diamond airfoil.

level as a reference solution to measure convergence. Adaptation is performed in the streamwise direction using a pressure indicator on the coarsest mesh level. Comparisons between the initial coarse and adapted grids are shown in Figure 19. A family of adapted grids is similarly generated by refining the adapted coarse grid with a refinement factor of 1.4. To compare results, we measure the L_2 norm of the line probe error compared to the extracted reference line probe. Figure 22 shows the error convergence for both the uniform streamwise and adapted grids. Observe that adaptation enables accuracy matching that of the next finest grid level. The convergence behavior between the uniform and adapted grids is similar; however, we see slower convergence from the adapted grid for the first three levels and faster convergence between the third and fourth grid levels. Finally, a qualitative comparison between the extracted line probe convergence is provided in Figures 21a and 21b.

4.5 Database Adaptation Feasibility

To demonstrate the feasibility of including adaptation within the database generation workflow, we apply the described adaptation methodology to the X-59 at the flight conditions of Sections 2 and 3.3. An automatic off-body mesh with uniform streamwise spacing was generated for the CFD domain. As before, the CFD solver is initially run for 2000 first-order iterations. Instead of serving as our warm-start, we

use this solution to build the adaptation indicator. Figure 23 shows a Mach number direction-based indicator on the multi-block off-body mesh topology. Adaptation is performed using this indicator and the first-order solution is interpolated onto the adapted grid. We then warm-start the CFD solver for the adapted grid using the interpolated first-order solution.




Figure 23: Feature-based multi-block adaptation indicator.

A comparison between a uniform, manually refined, and adapted off-body grid streamwise spacing is shown in Figure 24. The streamwise local refinement regions of the manual and adapted grids are comparable. Meaningful differences are observed at the leading shock, where equidistribution coarsens the uniform grid to provide more resolution downstream. Such point distributions are non-intuitive and suggest the need for adaptation, especially when running database simulations with possibly large variations in flight conditions.




Figure 24: Streamwise grid spacing comparisons.

Figure 25 compares the two step, three step, and adapted three step near-field signatures. Minor differences between the three step and adapted three step methods exist, resulting from the manually refined and adapted off-body grids. Grid sensitivity is strongest in the downstream region traced back to the t-tail shock interaction with the plume. The corresponding on-track ground-level signatures are presented in Figure 26. Despite the differences in near-field signatures in this region, the adapted grids





Figure 25: Near-field signatures for the two step and three step prediction procedures.

were able to provide loudness predictions with similar accuracy to those of the manually refined grids, as seen in Figure 26(a).

With strong agreement between manually refined grids and adapted grids, one may safely conclude for this case that adaptation does not sacrifice accuracy for automation. These results suggest a gain in robustness to varying flight conditions. A practical concern with adaptation is the cost of including it within the database workflow. Table 4 contains a breakdown of the time required to compute an adapted CFD solution. All steps of the process were ran using 2400 Intel Skylake Xeon Gold 6148 processors. The combination of MPI/openMP hybrid parallelism and the direction-based adaptation algorithm design render the cost of including adaptation in the database workflow negligible.



(a) Perceived loudness over the boom carpet.

(b) On-track ground-level overpressure signatures.

Figure 26: PL boom carpet and ground-level overpressure signatures.

Task	Wall-Clock Time
Initial Flow Solve	1 hr 51 min 58 s
Adaptation	2 min 37 s
Restarted Flow Solve	8 hr

Table 4: Adapted CFD timing breakdown for the X-59 on 2400 Intel Skylake Xeon Gold 6148 processors.

5 Conclusion and Future Work

Advancements on the traditional two-step loudness prediction solution procedure have been described and applied to the X-59. We have shown that truncating the CFD domain and using Space Marching for mid-field computation yields an accuracy enhancement at roughly half the cost. Database capabilities using LAVA Curvilinear were developed that allow series of cases at a wide range of flight conditions to be generated and preprocessed with more robustness and in a way that is less prone to human error. In the context of the database scripting system, leveraging the modularity of structured curvilinear overset grids by using a novel build-from-nominal grid generation process was introduced. This process has a grid built from nominal database flight conditions, modifying only the necessary parts of the geometry without complete grid regeneration to account for new specified flight conditions. This results in a 71% reduction in grid generation time.

A novel direction-based anisotropic Mach cone aligned mesh redistribution procedure was verified and applied to the X-59. The inclusion of adaptation into the three-step low boom simulation procedures requires only 2.5 minutes of wall-clock time. With low computational cost, the robustness of automatically adapting off-body CFD grids without sacrificing accuracy is an exciting advancement.

In the future much of the database capabilities designed for the X-59 have great potential to be used together with the NASA developed Computational Aerosciences Productivity and Execution (CAPE) framework [28]. The mesh redistribution procedure will also be explored using different indicators and convergence quantities, with the ultimate goal of including within our best practice workflow.

6 Acknowledgments

This work was partially supported by the Commercial Supersonic Technologies (CST) and Transformational Tools and Technologies (TTT) Projects within NASA's Aeronautics Research Mission Directorate (ARMD). Computational 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. The authors acknowledge the exceptional efforts of the LAVA team, particularly James Koch, James Jensen, Brandon Lowe, and Emre Sozer.

References

- [1] Juliet Page and Kenneth Plotkin. An efficient method for incorporating computational fluid dynamics into sonic boom prediction. In *9th Applied Aerodynamics Conference*, page 3275, 1991.
- [2] Michael A Park and Melissa B Carter. Nearfield summary and analysis of the third aiaa sonic boom prediction workshop c608 low boom demonstrator. In *AIAA Scitech 2021 Forum*, page 0345, 2021.
- [3] Jared C Duensing, James C Jensen, Jeffrey A Housman, Michael G Piotrowski, Daniel Maldonado, Emre Sozer, Cetin C Kiris, and Gaetan K Kenway. Structured and unstructured simulations for the third aiaa sonic boom prediction workshop. *Journal of Aircraft*, 59(3):624–646, 2022.
- [4] Melissa B Carter, Alaa A Elmiligui, Michael A Park, Patrick R Shea, Courtney S Winski, Donald A Durston, James C Jensen, Jacob M Wagner, Scott M Neuhoﬀ, and John D Wolter. Experimental and computational study of the x-59 wind tunnel model at glenn research center 8-by 6-foot supersonic wind tunnel. *NASA/TM-2022-0011496*, 2022.
- [5] Gerrit-Daniel Stich, Olivia G Martin, Chase Ashby, Jeffrey A Housman, and Jared Duensing. Noise prediction of multi-stream internally mixed jets with external and internal plugs using large-eddy simulation. In *30th AIAA/CEAS Aeroacoustics Conference (2024)*, page 3377, 2024.
- [6] 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.
- [7] Cetin C Kiris, Jeffrey A Housman, Michael F Barad, Christoph Brehm, Emre Sozer, and Shayan Moini-Yekta. Computational framework for launch, ascent, and vehicle aerodynamics (lava). *Aerospace Science and Technology*, 55:189–219, 2016.
- [8] Sriram K Rallabhandi. Advanced sonic boom prediction using the augmented burgers equation. *Journal of Aircraft*, 48(4):1245–1253, 2011.
- [9] BETA CAE Systems. Ansa pre-processor, 2024.
- [10] Cadence. Fidelity pointwise for cfd meshing, 2024.
- [11] William Chan. Developments in strategies and software tools for overset structured grid generation and connectivity. In *20th AIAA Computational Fluid Dynamics Conference*, page 3051, 2011.
- [12] Jeffrey A. Housman, Cetin C. Kiris, and Mohamed M. Hafez. Time-Derivative Preconditioning Methods for Multicomponent Flows—Part I: Riemann Problems. *Journal of Applied Mechanics*, 76(2):021210, 02 2009.
- [13] Jeffrey A. Housman, Cetin C. Kiris, and Mohamed M. Hafez. Time-Derivative Preconditioning Methods for Multicomponent Flows—Part II: Two-Dimensional Applications. *Journal of Applied Mechanics*, 76(3):031013, 03 2009.
- [14] P. Spalart and S. Allmaras. A one-equation turbulence model for aerodynamic flows. *30th Aerospace Sciences Meeting and Exhibit*, 1992.
- [15] Youcef Saad and Martin H. Schultz. Gmres: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM Journal on Scientific and Statistical Computing*, 7(3):856–869, 1986.
- [16] S. S. Stevens. Perceived Level of Noise by Mark VII and Decibels (E). *The Journal of the Acoustical Society of America*, 51(2B):575–601, 02 1972.
- [17] Jeffrey Housman, James Jensen, Gaetan Kenway, and Cetin Kiris. Algorithmic improvements to a high-order space marching method for sonic boom propagation. In *Advanced Modeling & Simulation Seminar Series*, 2022.
- [18] Xiaogang Deng. New high-order hybrid cell-edge and cell-node weighted compact nonlinear schemes. In *20th AIAA computational fluid dynamics conference*, page 3857, 2011.
- [19] Taku Nonomura and Kozo Fujii. Robust explicit formulation of weighted compact nonlinear scheme. *Computers & Fluids*, 85:8–18, 2013.
- [20] Chi-Wang Shu. High order weighted essentially nonoscillatory schemes for convection dominated problems. *SIAM review*, 51(1):82–126, 2009.
- [21] Chase Ashby, Jeffrey A Housman, and Jared C Duensing. Anisotropic mach cone aligned mesh adaptation for low boom simulations. In *AIAA AVIATION 2023 Forum*, page 4333, 2023.
- [22] Chase P. Ashby, Brandon M. Lowe, Jeffrey A. Housman, and Jared C. Duensing. Anisotropic mach cone aligned mesh adaptation for low boom simulations. *AIAA Journal*, 62(6):2076–2094, 2024.
- [23] Pieter Buning and Thomas Pulliam. Cartesian off-body grid adaption for viscous time-accurate flow simulations. In *20th AIAA Computational Fluid Dynamics Conference*, 2011.
- [24] David A Venditti and David L Darmofal. Adjoint error estimation and grid adaptation for functional outputs: Application to quasi-one-dimensional flow. *Journal of Computational Physics*, 164(1):204–

- 227, 2000.
- [25] David A. Venditti and David L. Darmofal. Grid adaptation for functional outputs: Application to two-dimensional inviscid flows. *Journal of Computational Physics*, 176(1):40–69, 2002.
 - [26] Aniruddha Choudhary, William C. Tyson, and Christopher J. Roy. Implementation and Assessment of a Residual-Based r-Adaptation Technique on Structured Meshes. *Journal of Verification, Validation and Uncertainty Quantification*, 3(4), 05 2019. 041005.
 - [27] Michael A. Park, Nicolas Barral, Daniel Ibanez, Dmitry S. Kamenetskiy, Joshua A. Krakos, Todd R. Michal, and Adrien Loseille. Unstructured grid adaptation and solver technology for turbulent flows. In *2018 AIAA Aerospace Sciences Meeting*, 2018.
 - [28] Derek Dalle. CAPE: Computational aerosciences productivity and execution, 2024.