Unstructured Grid Development for the Space Launch System Liftoff and Transition Lineloads Computational Analysis

Nalin A. Ratnayake*, Steven E. Krist, Ph.D. †, Farhad Ghaffari ‡, and Varda Ahmed §
NASA Langley Research Center, Hampton, VA, 23681-2199

Production of aerodynamic lineloads databases for the Space Launch System (SLS) vehicle at Liftoff and Transition (LOT) conditions has required the development of a Computational Fluid Dynamics (CFD) workflow capable of producing high-quality solutions for this unique phase of flight. Aerodynamic considerations included a wide range of flow angles (from 0° up to 90° total angle of attack), resulting leeside separation, and interaction effects between the three bodies of the integrated SLS vehicle, as well as the nearby launch tower. Computational mesh development for similar problems at the NASA Langley Research Center (such as for the Constellation/Ares launch vehicles) has primarily relied on in-house tools such as VGRID/POSTGRID, with grids designed for NASA-developed and maintained flow solvers such as USM3D and FUN3D. The workflow for such problems has evolved over the development of the various SLS configurations to incorporate new tools such as the Heldenpatch/Heldenmesh grid generator (Helden Aerospace) and CREATE-AV Kestrel (US Department of Defense) flow solver. This paper describes efforts to benchmark a grid generation approach for LOT problems using Heldenpatch/Heldenmesh and Kestrel, verified against prior best practices from VGRID/POSTGRID. Parameters studied include surface grid density, first-layer viscous cell height, and volume grid growth rate parameters. Resulting solutions are compared based on total force and moment values, sectional lineloads, and surface pressures, all validated against existing wind tunnel aerodynamic databases where available for the SLS Block 1B Cargo configuration.

Nomenclature

C_P	=	Coefficient of pressure
$\frac{dCN}{\Delta x/L_{ref}}$	=	sectional normal force coefficient lineload
$\frac{dCY}{\Delta x/L_{ref}}$	=	sectional side force coefficient lineload
L_{ref}	=	reference length (nominal centerbody diameter)
M_∞	=	freestream Mach Number
R1	=	volume grid growth rate 1
R2	=	volume grid growth rate 2
α_P	=	total angle of attack (missile axis system), deg
δ_j	=	cell height of the jth grid layer above a surface
$\Delta X/L_{ref}$	=	normalized axial location
ϕ_P	=	aerodynamic roll angle (missile axis system),deg
AMR	=	Adaptive Mesh Refinement
CFD	=	Computational Fluid Dynamics
DAC	=	Design and Analysis Cycle
GN&C	=	Guidance, Navigation, and Control
LaRC	=	Langley Research Center
LAS	=	Launch Abort System
LOT	=	Liftoff and Transition (phase of flight)

*Research Aerospace Engineer, Configuration Aerodynamics Branch, AIAA Senior Member
†Research Aerospace Engineer, Configuration Aerodynamics Branch, AIAA Senior Member
‡Research Aerospace Engineer, Configuration Aerodynamics Branch, AIAA Associate Fellow
§Universities Space Research Association Intern, Configuration Aerodynamics Branch

MPCV	=	Multi-Purpose Crew Vehicle
MSFC	=	Marshall Space Flight Center
NASA	=	National Aeronautics and Space Administration
OML	=	Outer Mold Line
RANS	=	Reynolds-Averaged Navier-Stokes
SA − DDES	=	Spalart-Allmaras turbulence model with Delayed Detached Eddy Simulation
SBU	=	Standard Billing Unit
SLS	=	Space Launch System
SRB	=	Solid Rocket Booster
SWORDS	=	Soldier-Warfighter Operationally Responsive Deployer for Space
TetrUSS	=	Tetrahedral Unstructured Software System
14x22	=	NASA Langley Research Center 14- by 22-Foot Subsonic Tunnel

I. Introduction

The Space Launch System (SLS) consists of a common, multistage centerbody containing the liquid-fueled main rocket engines, and two five-segment SRBs attached to the sides. For the SLS-1000x (Block 1 Crew) and SLS-2800x (Block 1B Crew) series, the Orion Multi-Purpose Crew Vehicle (MPCV) sits atop the Upper Stage, and the Launch Abort System (LAS) is attached to the MPCV. For the SLS-1050x (Block 1 Cargo) and SLS-2700x (Block 1B cargo) variant, the upper stage MPCV/LAS is replaced with a payload shroud. Figure 1 shows these three configurations, while also depicting the larger Block 2 Crew and Cargo variants. For all configurations, the last two digits are reserved for indicating the specific version of the configuration as the vehicles pass through multiple Design and Analysis Cycles (DACs).

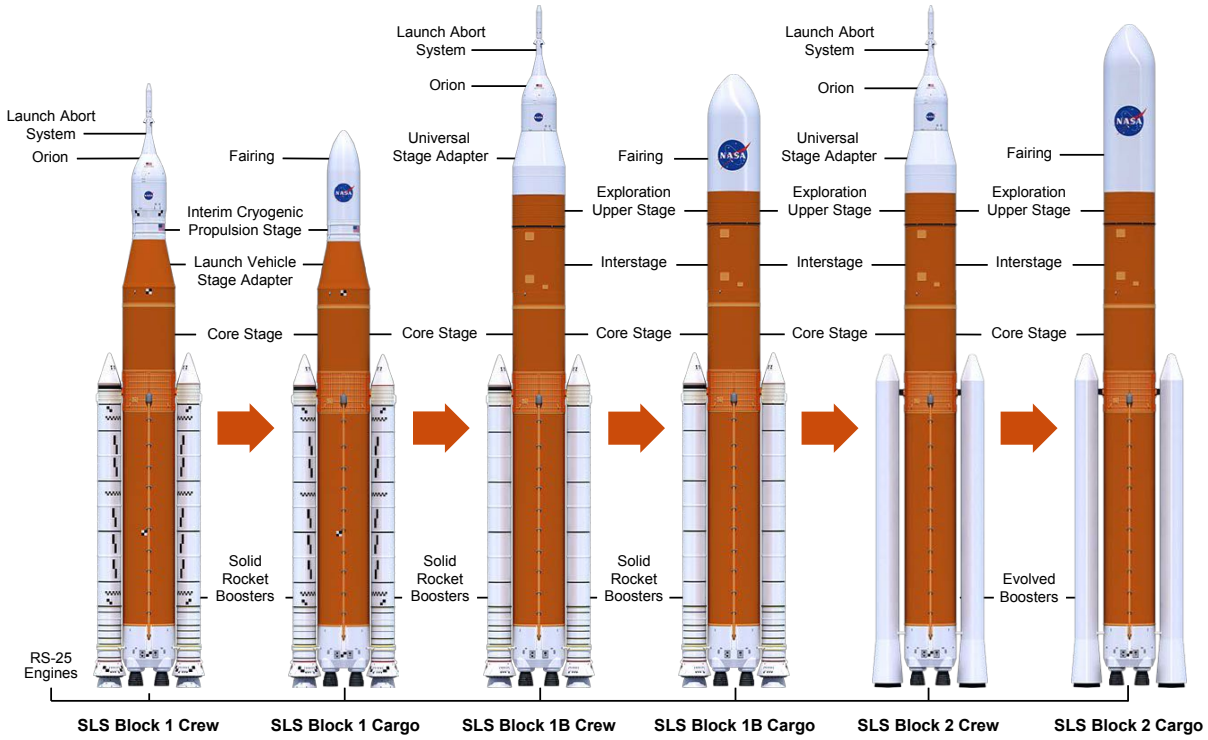

Fig. 1 SLS Vehicle Configurations.

The present study focuses on the Block 1B Cargo configuration (specifically, the 27005 version), due to its simpler Outer Mold Line (OML) – which enables faster solution turnaround – and the availability of vetted force and moment wind tunnel databases.

Since the primary deliverable for the problem of interest is aerodynamic lineloads, these will serve as the primary means of verifying solutions based on the computational meshes under study against solutions made at the same conditions with legacy methods, in particular those that informed SLS aerodynamic lineloads database SLS-27-D-LLA-001. Further, validation against wind tunnel data will be made using the integrated forces and moments for the SLS-27005 vehicle (Block 1B Cargo DAC2), obtained from force and moment aerodynamic database SLS-27-D-LFA-001 and based on Test 633 at the NASA Langley Research Center (LaRC) 14- by 22-Foot Subsonic Tunnel (14x22). Details concerning the wind tunnel testing and database creation are provided by Chan in Ref. [1].

The SLS aerodynamic lineloads databases provide the distributed (sectional) aerodynamic forces and moments acting along the length of the various SLS configurations from vehicle on-pad, through liftoff, then through transition to the nominal ascent trajectory. They are provided individually for the centerbody and each of the two Solid Rocket Boosters (SRBs). The main end users of the Liftoff and Transition (LOT) lineloads databases are the Loads team and the Guidance, Navigation, and Control (GN&C) team at the Marshall Space Flight Center (MSFC), who use the databases for sectional loading analysis and gimbal rate sensitivity analysis, respectively.

While the overall vehicle force and moment coefficients can be efficiently obtained through wind tunnel testing, computational analyses have become indispensable for obtaining the extensive amount of surface information (i.e., surface pressures and skin friction coefficients) required to generate accurate lineloads (distributed loads as a function of axial location).

A. Prior Work

The approach used to develop the database for the Ares-I launch vehicle using USM3D, as published by Abdol-Hamid and Ghaffari [2], provided a strong foundation for grid generation strategy and best practices on which to build the present SLS computational effort. Though the Ares-I was a different vehicle, many of the numerical approaches and modeling strategies that were used are essentially the same as those of the SLS.

Abdol-Hamid and Ghaffari performed grid convergence studies for the Ares-I vehicle using a factor of 2/3 on the initial viscous layer height for each successive level of grid refinement. The coarse, base, and fine grids were used to compute the flow at a total angle of attack (α_P) = 8°, Mach numbers (M_∞) = 0.9, 1.6, and 3.0, and an appropriate wind tunnel Reynolds number. These computations were performed at aerodynamic roll angles (ϕ_P) ranging from 0° to 330° in increments of 30°.

Full, rigorous grid convergence studies for each SLS configuration were not performed prior to the initial development of LOT lineloads databases for SLS Block 1 Crew and 1B Crew and Cargo configurations, due to the large size of the grids and associated demands on resources at the time. However, the best practices from Ares-I, as well as from analysis of the Soldier-Warfighter Operationally Responsive Deployer for Space (SWORDS) and Titan III vehicles [3] were carried forward into the SLS work. The present study is intended to build upon this necessary verification and validation work prior to the next design analysis cycles for all SLS configurations.

B. Description of the Present Study

All CFD cases in the present study were produced using the CREATE-AV Kestrel flow solver [4] version 7.1.2 [5], with an unstructured near-body grid overset with a Cartesian off-body grid undergoing Adaptive Mesh Refinement (AMR) based on vorticity magnitude. Second-order spatial and temporal discretization (third-order spatial in the off-body Cartesian grid) and the Spalart-Allmaras turbulence model with Delayed Detached Eddy Simulation (SA-DDES) were employed. Computations were performed on the Pleiades cluster of the NASA Advanced Supercomputing facility located at the NASA Ames Research Center. Typical CFD solutions of the LOT problems required grids on the order of one to three billion grid points. Typical computer resource requirements involve the use of 2000 IvyBridge processors (or their total equivalent on other architectures) running from 24 to 120 wall hours, depending on the specific case.

Kestrel uses a two-grid / two-solver approach, with a near-body unstructured grid that is trimmed to a specified distance off the surface of the vehicle. Kestrel then wraps a structured, Cartesian volume grid around the near-body grid, which then undergoes AMR. The near-body solution is computed using the kCFD solver and the off-body solution is computed using SAMAir, with the overlap region handled through an Overset [6] approach.

An example flowfield and mesh slice that shows the two-grid approach used by Kestrel is provided in Fig. 2. The inset at top left shows the transition between the near-body unstructured grid and the off-body Cartesian grid, while the main portion of the figure shows the AMR of the Cartesian grid in response to vorticity magnitude, as expressed by the color contours. The present study primarily focuses on developing a best-practice approach for generating an efficient near-body unstructured grid for the present SLS class of vehicles.

The cases were run at the wind-tunnel Reynold's number of 615,000 (based on nominal reference diameter), Mach number of 0.182, and a freestream static temperature of 522°R. More details about the evolution of the CFD approach and the rationale behind the choice of flow solver settings, turbulence model, and other aspects of the approach are provided by Ratnayake, et al. in Ref. [7].

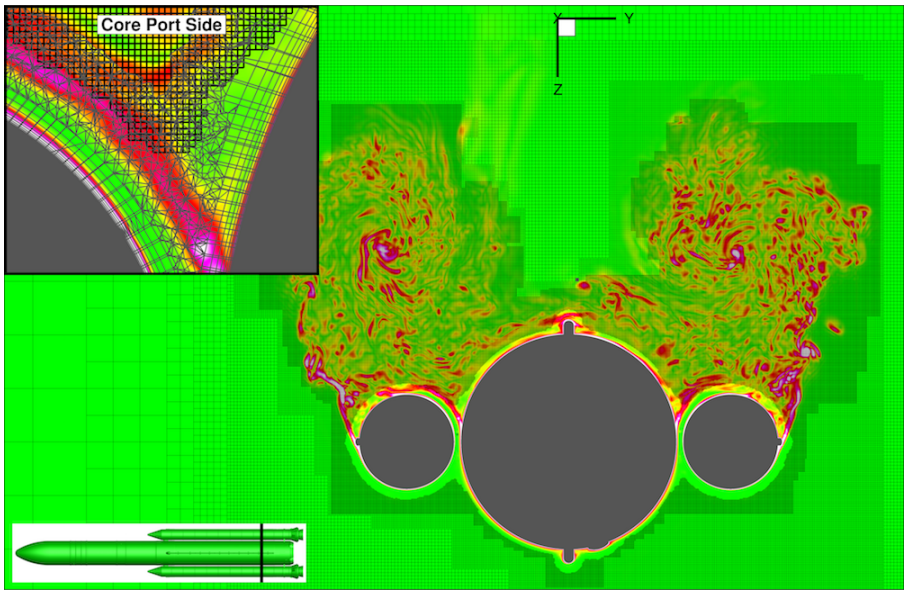

Fig. 2 Grid and flowfield slice showing Kestrel dual-mesh and adaptive mesh refinement based on vorticity for the SLS-27005 at α_P=30° and ϕ_P=0°. The point of view is aft looking forward, and the inset shows a zoomed view of the region between the port SRB and the centerbody.

GridTool and VGRID are part of the Tetrahedral Unstructured Software System (TetrUSS), which is a package of CFD tools including the flow solver USM3D, developed and maintained by the Configuration Aerodynamics Branch at the NASA Langley Research Center [8]. The combination of GridTool and VGRID has been used extensively at NASA LaRC, and best practices for these tools are well-established within the TetrUSS user community. Recently however, the addition of Heldenpatch and Heldenmesh (which make use of GridTool/VGRID restart files and sources) to the set of tools used by the team has greatly reduced the turnaround time for the generation of new unstructured surface/volume grids, largely by automating many geometric cleanup and sourcing tasks that were previously completed manually in GridTool.

Heldenmesh is hypothesized to provide a higher quality grid, on factors such as cell skewness and distribution of sources, resulting from much finer distribution and control of sources. VGRID requires that point, line, and volume sources be specified manually, either individually or in small groups (for example breaking up a curve into line source segments based on curvature angles). The legacy reference grid used in this study contained 5780 individual sources, defined manually in this manner using GridTool Cocoa. In Heldenmesh by contrast, the user specifies a more abstract sourcing strategy for patch families, including parameters such as minimum cell size, what level of surface angle change constitutes a curve or an edge, and maximum stretching ratio. The sources themselves are created algorithmically, which allowed for the baseline Heldenmesh grid used in this study to contain 38,746 individual sources. Though the user has less direct control over specific sources (unless they manually add additional sources via GridTool), the overall level of sourcing control over the resulting cell quality is much finer in Heldenmesh.

Both VGRID and Heldenmesh use Eq. 1 to determine the volume grid growth rate.

$$\delta_j = \delta_1(1 + R1(1 + R2)^{j-1})^{j-1} \tag{1}$$

In Eq. 1, δ_1 represents the height of the first cell in the normal direction from the surface, $R1$ and $R2$ are the two grid growth rates, and j is the layer number.

This study will determine the effects of four grid parameters on the computed sectional force lineloads:

1) Surface grid density: A baseline grid will be generated in Heldenmesh that is roughly comparable in overall cell count and surface mesh distribution to the legacy VGRID approach. Then, the source sizes will be adjusted

uniformly using the REFINE_FAC variable in the Heldenmesh input file for four additional refinement levels. For two of the coarser refinement levels, the cell height of the first viscous layer will also be expanded.

2) Volume grid layer growth rate 1: The baseline surface grid will then be used to study the effects of grid growth rate parameters of R1 and R2. R1 will be varied from 0.12 to 0.20 in increments of 0.02 with R2 held constant. The nominal value from previous best practices is 0.16.
3) Volume grid layer growth rate 2: Similar to the above, R2 will be varied from 0.01 to 0.04 with R1 held constant. The nominal value from previous best practices is 0.02.
4) Near-body grid trim distance from the vehicle surface: The trim distance will be varied from 7.2% of the centerbody reference diameter to 12.6% in increments of 1.8%. The nominal value from previous best practices is 10.8%.

The study focused on total angles of attack of $\alpha_P = 50°$, which is representative of a wide range of total angle of attack, where complex aerodynamic phenomena such as asymmetric vortices and large separation wakes form.

II. Establishing a Baseline Heldenmesh Grid

A. Baseline Surface Mesh

Figures 3 through 5 show surface mesh comparisons between the baseline (1.00-scale) Heldenmesh surface mesh and the legacy surface mesh from VGRID/POSTGRID. It is not possible to replicate the VGRID surface mesh exactly in Heldenmesh even with manually specified sources in GridTool, which would circumvent many of the advantages of using Heldenmesh in the first place. However, it can be seen that the two surface meshes can be made generally similar in cell size and distribution for the same features.

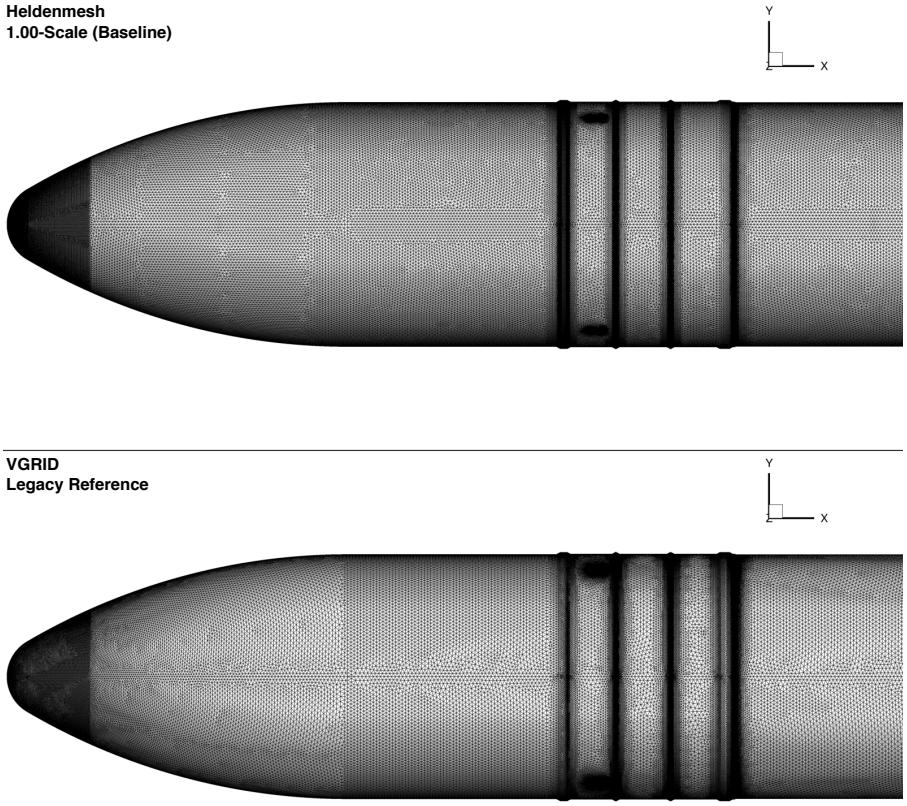

Fig. 3 Forward cargo fairing area, baseline Heldenmesh surface mesh vs VGRID.

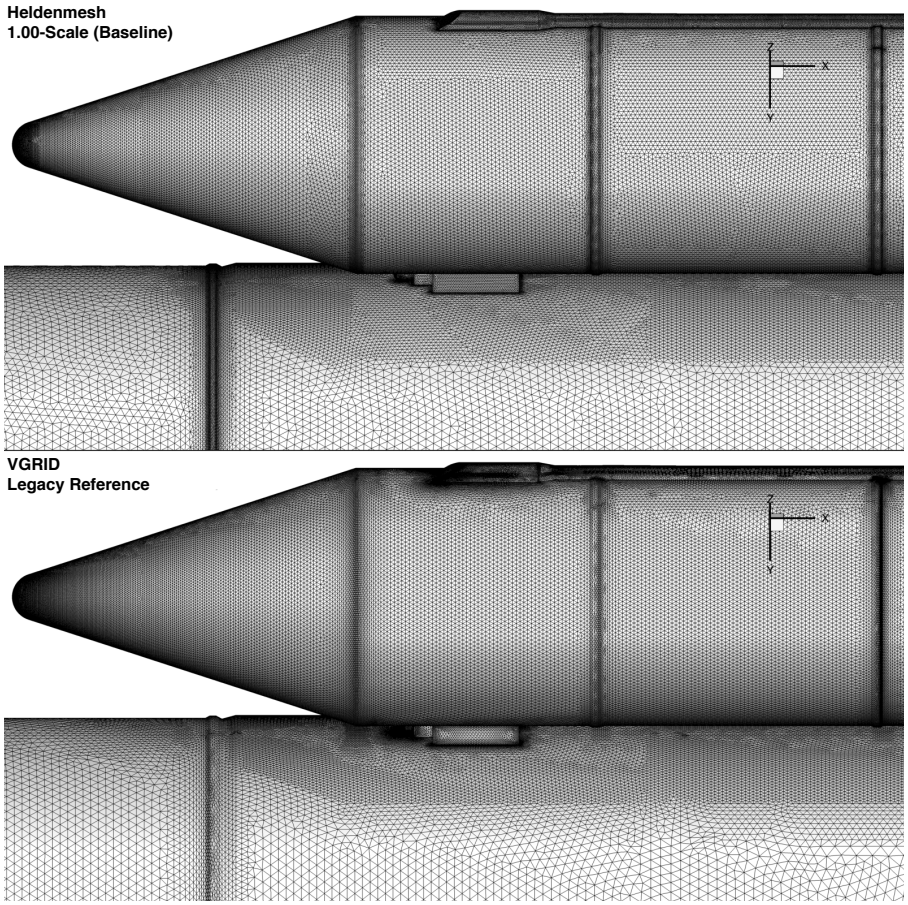

Fig. 4 Forward SRB attachment area, baseline Heldenmesh surface mesh vs VGRID.

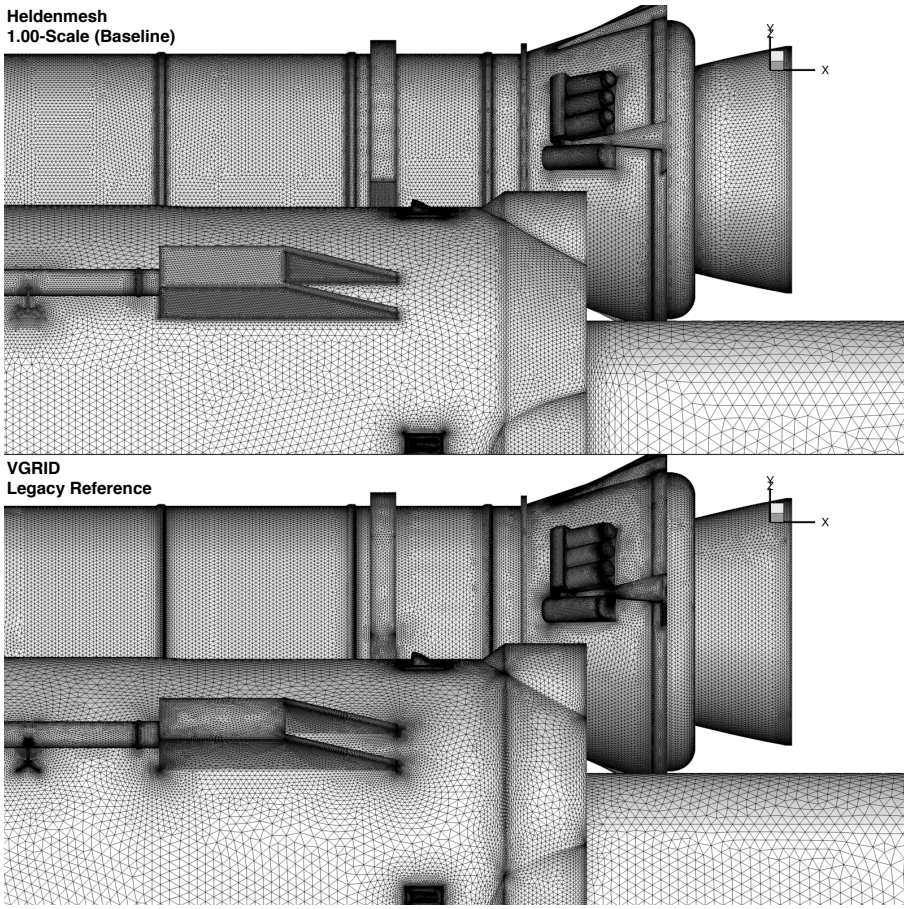

Fig. 5 Aft SRB attachment area, baseline Heldenmesh surface mesh vs VGRID.

The 1.00-scale volume grid was generated using the same boundary layer and growth settings in Heldenmesh as for the VGRID grid. Namely, an initial layer height of 0.01 inches, a value for R1 of 0.16, and a value for R2 of 0.02.

A comparison of the 1.00-scale Heldenmesh volume grid to the legacy VGRID volume grid is shown in Fig. 6. An area just aft of the fairing nose shoulder is depicted, with the two grids shown back-to-back for easier comparison between the volume slices. Broadly similar cell sizes can be seen between the surface meshes, which have 3.19 million surface triangles (VGRID) and 2.83 million surface triangles (Heldenmesh). There is good agreement between the volume layers all the way up to the outer boundary of the near-body grid, which is trimmed to 10.8% of the nominal reference diameter off the surface in the case of both grids. A noticeable difference is that Heldenmesh grew one additional viscous layer despite identical settings, indicating perhaps a slight difference with VGRID in how the stop condition for the advancing layer algorithm is implemented.

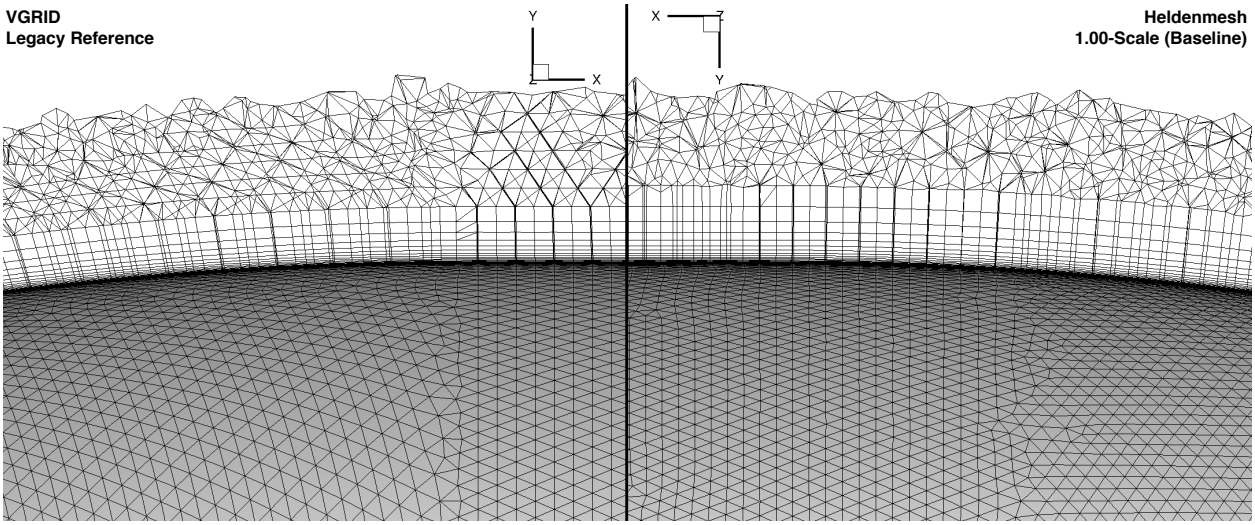

Fig. 6 Nose shoulder Z=0 slice, VGRID legacy vs baseline Heldenmesh volume grid (shown back-to-back).

Once trimmed to 10.8% and with viscous layer tetrahedrons merged to prisms in Carpenter, the baseline Heldenmesh near-body volume grid contains 123.4 million volume cells, as compared to the legacy VGRID 128.8 million volume cells. The additional layer of merged tetrahedral cells contributes to the slight discrepancy between the two cell counts, along with the aforementioned slight differences in the surface mesh that arise from using two different grid generation tools. When run in Kestrel at a condition of $\alpha_P = 50°$ and $\phi_P = 0°$ with identical off-body Cartesian grid refinement settings, the baseline Heldenmesh grid results in the size of the smallest off-body Cartesian cell being 1.2746 inches, versus 1.3461 inches from the legacy grid.

B. Quality Assurance of the Baseline Heldenmesh Grid

The baseline grid was run at $\alpha_P = 50°$ and several aerodynamic roll angles, and the surface pressure compared to static pressure measurements from the wind tunnel in order to validate that the cases were capturing relevant flow phenomena and adhering to trusted experimental data.

A representative solution history plot is provided in Fig. 7. The specific quantitative information pertaining to the geometry and performance of the SLS vehicle are designated Sensitive But Unclassified and subject to the International Traffic in Arms Regulation. Therefore, some quantitative information has been removed from this figure, and others that follow in the paper. The left two subplots show the solution history of the forces and moments in units of standard deviations over the time-averaging segment (timesteps 6,000-10,000). The right two subplots show the near-body (NB) and off-body (OB) mass residual histories, the NB and OB sweeps, and the number of off-body Cartesian grid points (reaching a maximum of roughly 1.25 billion in this case).

For comparison, the same history plot for the VGRID reference case is shown in Fig. 8.

Figures 9 and 10 show CFD-generated surface pressures at five axial stations (relative locations indicated on the figure) compared to static pressure measurements at $\alpha_P = 50°$ and $\phi_P = 0°$ from Test 633 at the LaRC 14x22 tunnel. Figures 11 and 12 show the same comparisons at $\alpha_P = 50°$ and $\phi_P = 60°$. On these plots, the thick circle represents the outer surface of the vehicle at each slice location, as well as the polar radius that represents zero pressure coefficient.

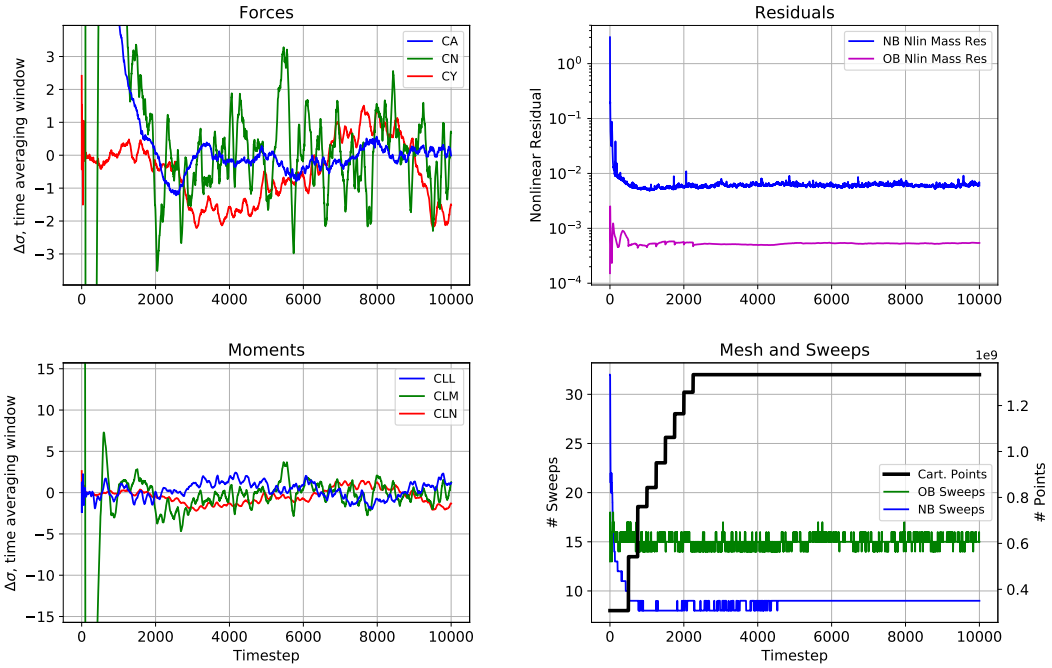

Fig. 7 Solution history for 1.00-scale Heldenmesh baseline grid at $\alpha_P = 50^\circ$ and $\phi_P = 0^\circ$.

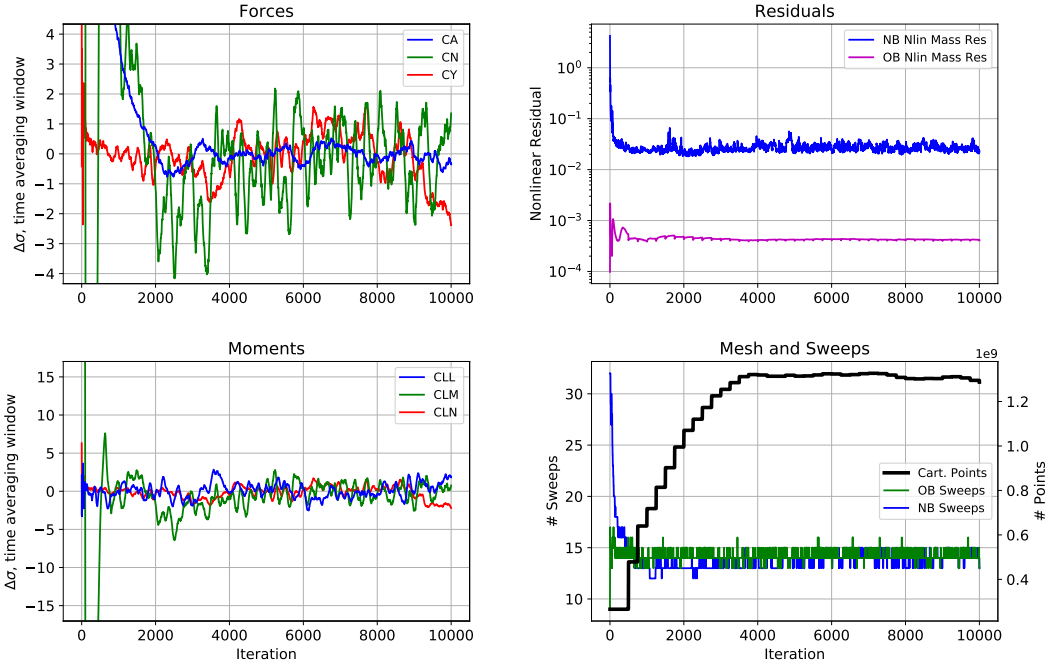

Fig. 8 Solution history for VGRID reference grid at $\alpha_P = 50^\circ$ and $\phi_P = 0^\circ$.

The arrow is for orientation and represents the relative direction of incoming flow. The wind tunnel data are represented by the open circles; no data are available for the SRBs from this test. The CFD solution is represented by the solid line bordering the contour region, in which cool colors represent negative relative pressure (suction) and warm colors represent positive relative pressure.

It can be seen that the Heldenmesh solutions compare reasonably well to the wind tunnel data at both flow conditions. With these highly-separated flows, it should be no surprise that the leeward side is not as well predicted as the windward side of the body.

III. Effect of Surface Mesh Refinement

After the baseline 1.00-scale surface mesh was created in Heldenmesh that approximated the legacy VGRID mesh, four additional Heldenmesh grids were created to result in source size factors of 0.90, 1.20, 1.40, and 1.60, with larger numbers indicating a coarser mesh. Due to roughly square-root dependence, a factor of 1.41 corresponds to approximately half as many elements.

A. Grid Refinement

Based on the same sourcing strategy (including relative sizes, strengths, stretching, and weights) as the 1.00-scale baseline, the sources were globally scaled using the REFINE_FAC variable in the Heldenmesh input file with values of 0.90, 1.20, 1.40, and 1.60 to create new surface meshes. Then each of these surface meshes was grown into a volume grid using the same volume growth settings as the baseline. Ideally, the study would have included scaling factors of 0.71 and 2.0; however, on the fine end, this resulted in volume grid cell counts that exceeded memory limits of the tools being used, and on the coarse end, it became difficult to close a grid sucessfully.

A representative visual example of the effect of scaling up the sources is shown in Fig. 13, which shows the aft region of the 1.40-scale surface mesh as compared to the baseline.

Near-body, trimmed volume cell counts for all scaling factors are shown in Table 1. Cell counts range from over 150 million in the 0.90-scale grid down to just under 50 million in the 1.60-scale grid.

B. Results

1. Comparison of Required Computational Resources

Ideally, all cases would have been completed on the same processor architecture with the same number of processors per job as the legacy VGRID reference case. However, available time and supercomputer queue constraints prevented such a straightforward comparison. The legacy VGRID reference case for $\alpha_P = 50°$ and $\phi_P = 0°$ completed 10,000 timesteps in 82.8 hours using 1,200 Haswell CPUs on 50 nodes (99,360 total CPU-hours, 9.94 CPU-hours per timestep). The baseline Heldenmesh 1.00-scale grid completed 10,000 timesteps in 48.4 hours using 2,000 Ivybridge CPUs on 100 nodes (96,800 total CPU-hours, 9.68 CPU-hours per timestep).

However, the two processor architectures are not equivalent in their efficiency in performing work. The NASA High-End Computing Program uses a Standard Billing Unit (SBU) factor to normalize each CPU architecture's equivalent work throughput to a baseline architecture (Broadwell) as described in Ref [9]. The conversion factor for Haswell is 0.80, and the factor for Ivybridge is 0.66. Application of this conversion factor to this comparison requires multiplying the wall time of each case by the number of nodes and then by the SBU conversion factor.

Table 1 shows a normalized resource requirement of 3314 SBUs for the VGRID reference case and 3194 SBUs for the baseline Heldenmesh 1.00-scale grid, an approximate reduction of 3.6%. On the coarsest end of the spectrum, the 1.60-scale Heldenmesh grid showed a normalized resource requirement of 1321 SBUs, a 60% reduction in required computational resources vs the VGRID reference case.

2. Comparison of Lineloads

Lineloads for the centerbody component were extracted by taking 1001 x-slices of the time-averaged pressure and skin friction surfaces and integrating around each ring-shaped slice to find local sectional forces in each direction. The lineloads for each case were compared to the reference lineload from the legacy VGRID volume grid, and the integrated value compared to the wind-tunnel-based SLS force and moment database.

The value of the sectional lineload at each slice is defined as the length-normalized contribution of that slice to the overall aerodynamic loading on the vehicle. The cumulative lineload at each x-slice is defined as the cumulative sum

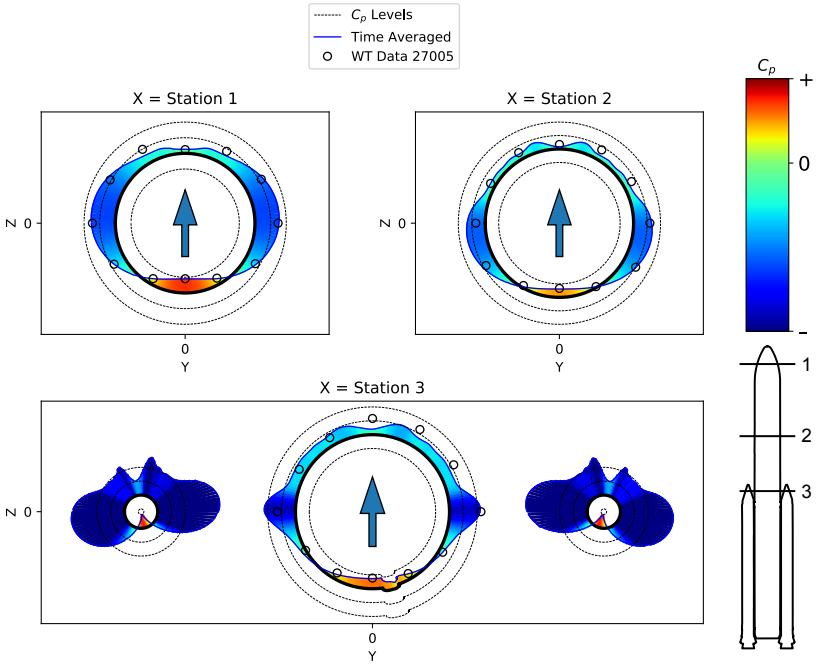

Fig. 9 Surface pressure comparison to wind tunnel static pressures at α_P = 50° and ϕ_P = 0° for Stations 1, 2, and 3.

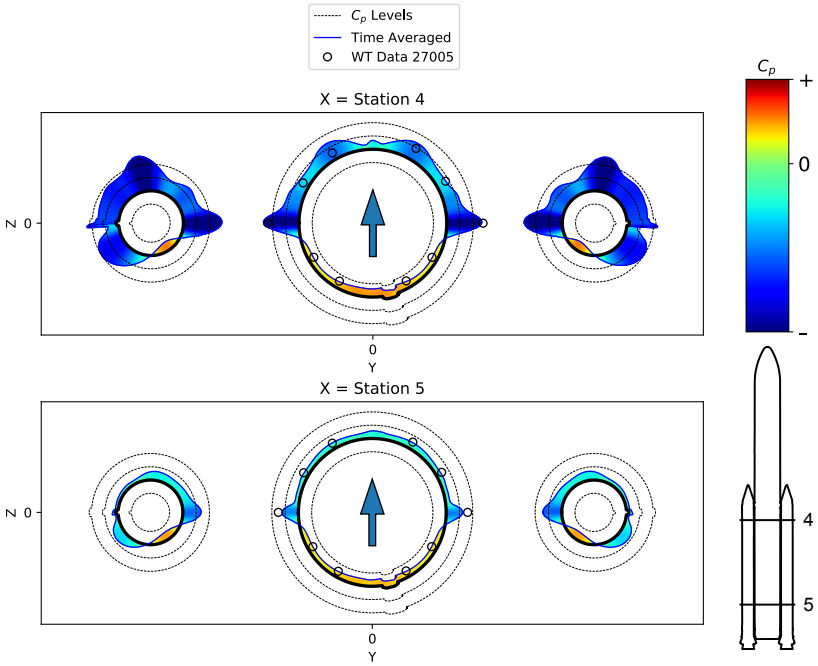

Fig. 10 Surface pressure comparison to wind tunnel static pressures at α_P = 50° and ϕ_P = 0° for Stations 4 and 5.

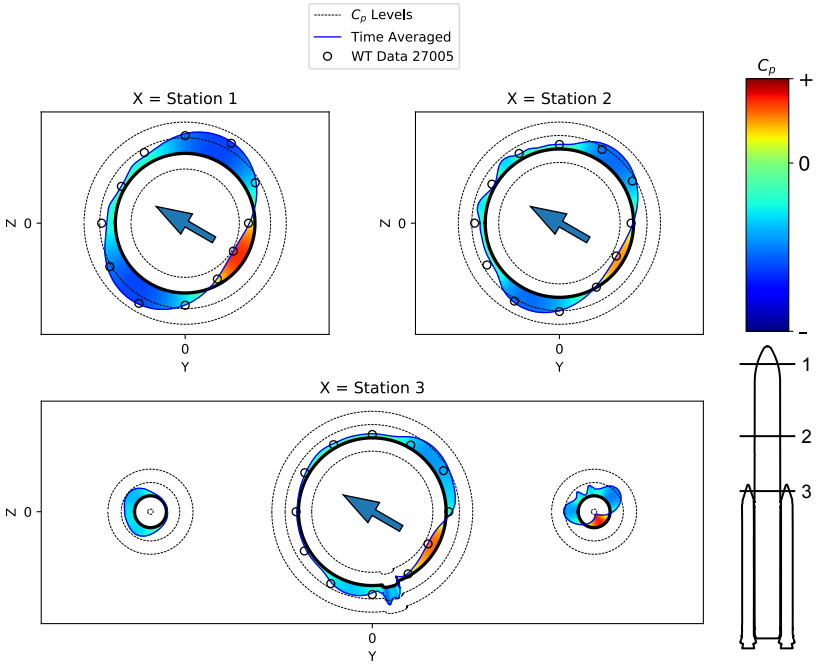

Fig. 11 Surface pressure comparison to wind tunnel static pressures at α_P = 50°and ϕ_P = 60°for Stations 1, 2, and 3.

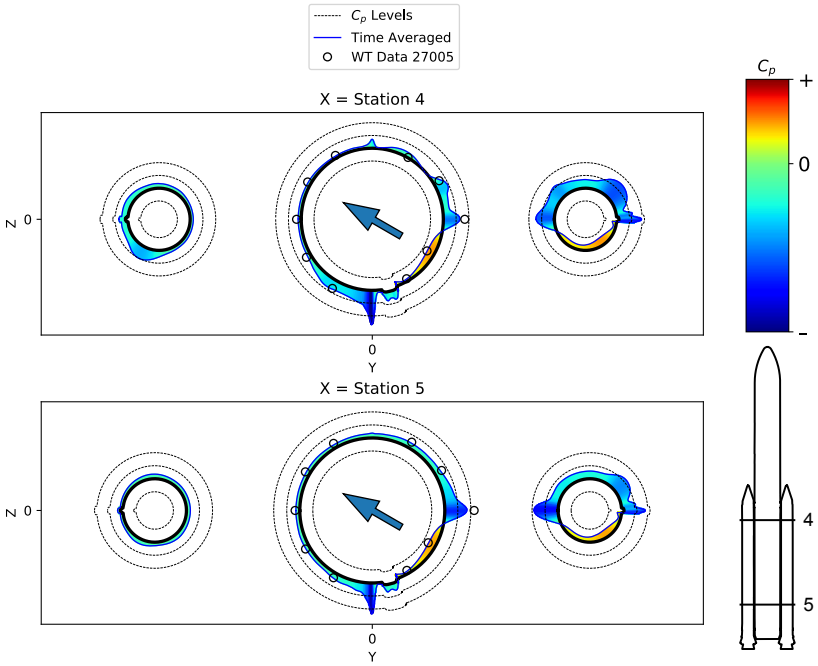

Fig. 12 Surface pressure comparison to wind tunnel static pressures at α_P = 50°and ϕ_P = 60°for Stations 4 and 5.

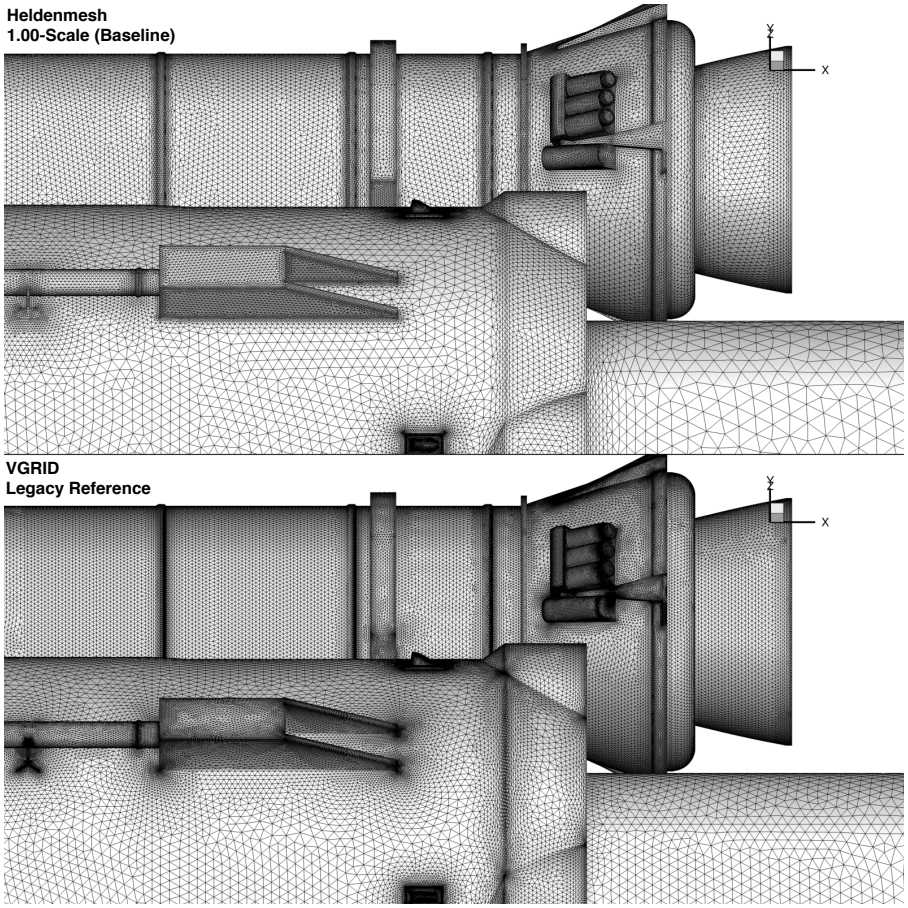

Fig. 13 Aft SRB attachment area, 1.40-scale Heldenmesh surface mesh vs VGRID.

Table 1 Resource requirements for levels of surface mesh source scaling in Heldenmesh versus the VGRID reference grid, to reach 10,000 timesteps.

Grid	NB Vol Cells (M)	Wall Time (hrs)	Nodes	Arch	Factor	SBUs	Δ
VGRID Reference	128.8	82.8	50	Has	0.80	3314	0.00%
Heldenmesh Baseline	123.4	48.4	100	Ivy	0.66	3194	-3.60%
0.90-scale Surface	154.8	51.5	100	Ivy	0.66	3399	2.58%
1.20-scale Surface	84.6	38.2	100	Ivy	0.66	2521	-23.91%
1.40-scale Surface	62.4	34.8	100	Ivy	0.66	2297	-30.69%
1.60-scale Surface	48.7	20.1	100	Ivy	0.66	1327	-59.96%

of the redimensionalized sectional lineload values at all slices from the fore tip of the body up to and including the slice itself. In the limit where the number of lineload x-slices approaches the surface mesh resolution, the cumulative lineload value at the last slice should equal the total integrated force on the vehicle.

In the lineloads comparisons that follow, it is important to note that there is no explicit requirement that the integrated value of the CFD lineloads actually match the total force and moment database value in SLS-27-D-LFA-001, as the two databases are developed from different data sources and delivered to different end users with differing purposes. Further, the lineload from the legacy VGRID grid should not be taken as a truth model, but merely as a reference lineload that represents the team's previous efforts. Further, the scales on all lineload plots are not necessarily the same. This was done because it was deemed more important to show differences between solutions, rather than compare all solutions to a target value.

Example sectional and cumulative normal force lineloads for various grid levels at the condition $\alpha_P = 50°$ and $\phi_P = 0°$ are shown in Figs. 14 and 15, respectively. The sectional side force lineloads for the same condition are shown in Fig. 16.

In each of these figures, the abscissae show axial position aft of the end of the centerbody, normalized by the nominal reference diameter, and the ordinates have been scrubbed of values except for zero. Below the lineloads plot is an approximate tracing of the outer contour of the centerbody, so that features on the lineload can be seen in the context of where they are geometrically on the vehicle component. The ordinate proportion has been exaggerated to better show geometric features.

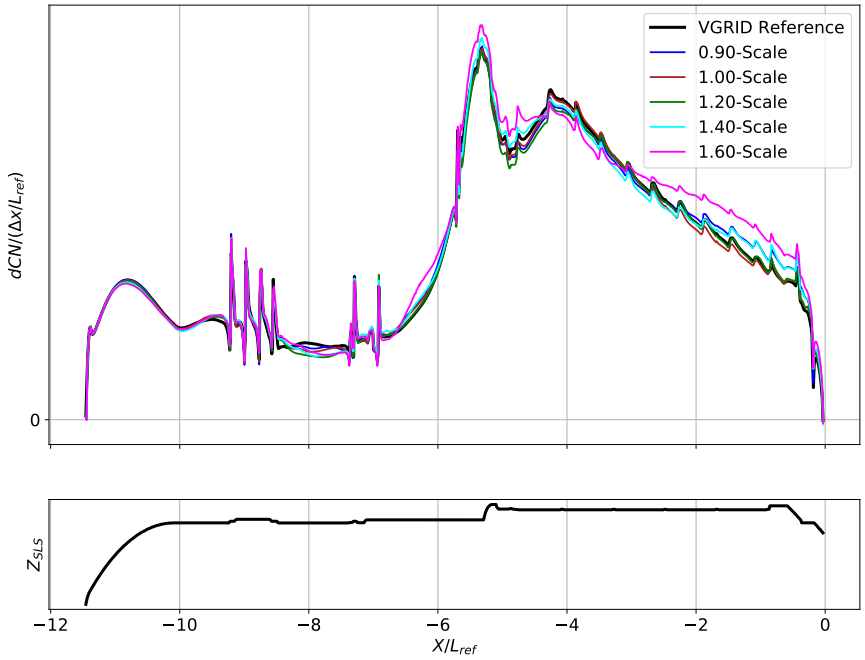

Fig. 14 Body-axis sectional normal force lineload for the centerbody, for various grid levels at $\alpha_P = 50°$ and $\phi_P = 0°$.

It is evident that the Heldenmesh cases deliver substantially similar normal force lineloads to the VGRID legacy case, diverging slightly toward the aft end of the rocket, with the coarsest grid solution being most out of family. As shown in Fig. 15, the total normal force is somewhat further from the value in the wind tunnel database SLS-27-D-LFA-001, which represents the value obtained from Test 633 at the 14x22.

The side force lineloads in Fig. 16 present an excellent example of the difficulties that arise from the vortex asymmetry phenomenon, in which for vehicles with slender forebodies, large, stable asymmetric vortices can be generated at midrange angles of attack [10]. Here the results show that changes to the grid spacing (with nothing else about the simulation changed) can cause the vortex asymmetry to flip to either side of the vehicle, or even not appear at all. This is consistent with numerous studies [11] [12] [13] that have shown very slight perturbations (in both wind tunnel experiment and computational simulation) can cause a flip in directionality of the attached vortex. Further details

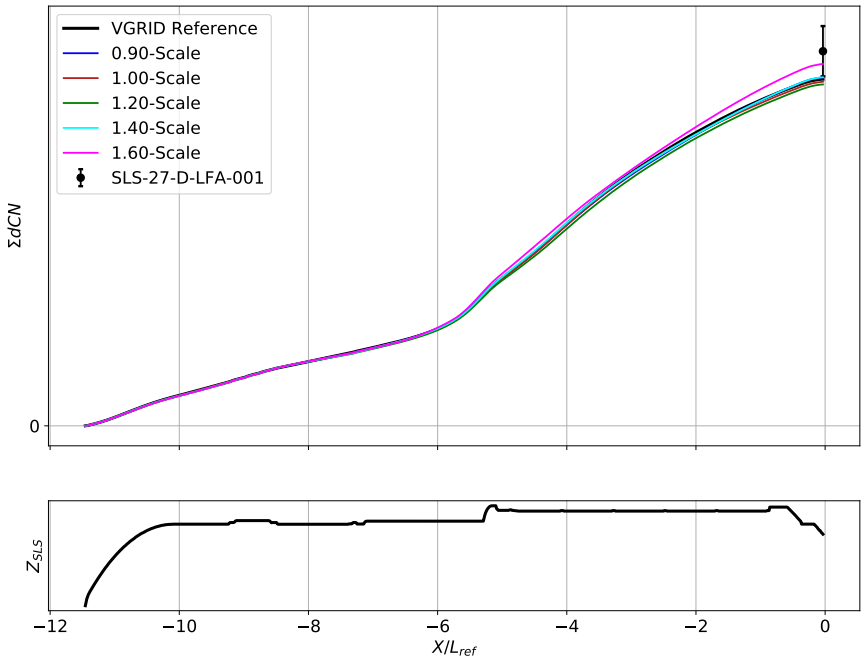

Fig. 15 Body-axis cumulative normal force lineload for the centerbody, for various grid levels at α_P = 50° and ϕ_P = 0°.

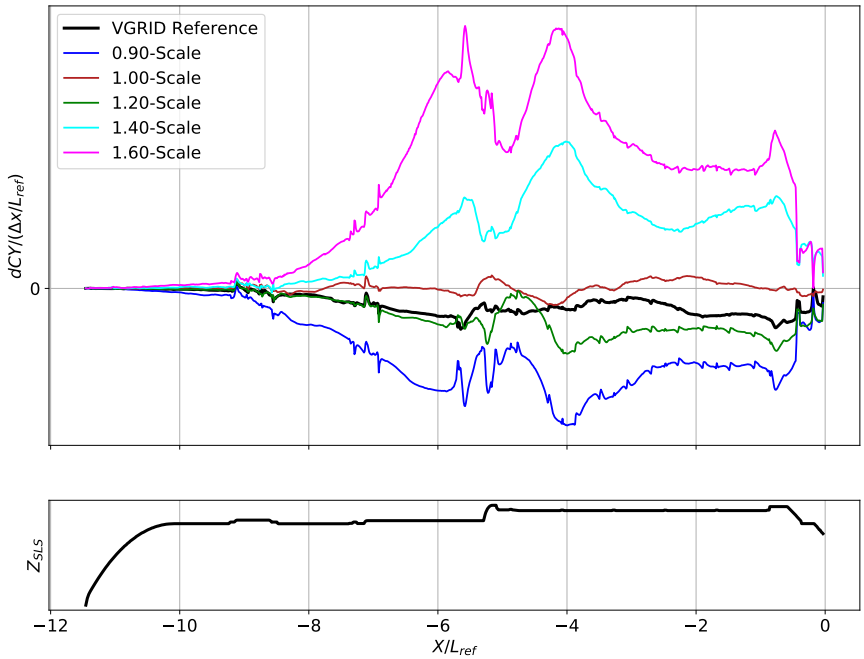

Fig. 16 Body-axis sectional side force lineload for the centerbody, for various grid spacing levels at α_P = 50° and ϕ_P = 0°.

about the phenomenon and the team's approach to characterizing and addressing the effects of vortex asymmetry in the context of the SLS Liftoff and Transition lineloads problem is described by Ratnayake, et al. in Ref. [7].

The results suggest that the baseline surface mesh resolution on the centerbody is generally more than adequate and could likely be coarsened in most areas of the centerbody. The exception would be toward the aft end of the centerbody, where the finest grid seems to be locally not quite grid converged between 0 and 2 diameters forward of the aft end of the rocket. This would suggest that the baseline surface mesh may not be sufficiently capturing some geometric features, such as protuberances, toward the aft end of the vehicle.

Example normal force lineloads from the starboard SRB are shown in Figs. 17 and 18. The plots show more dramatic variation due to surface mesh refinement than for the centerbody, though the overall conclusions remain largely the same.

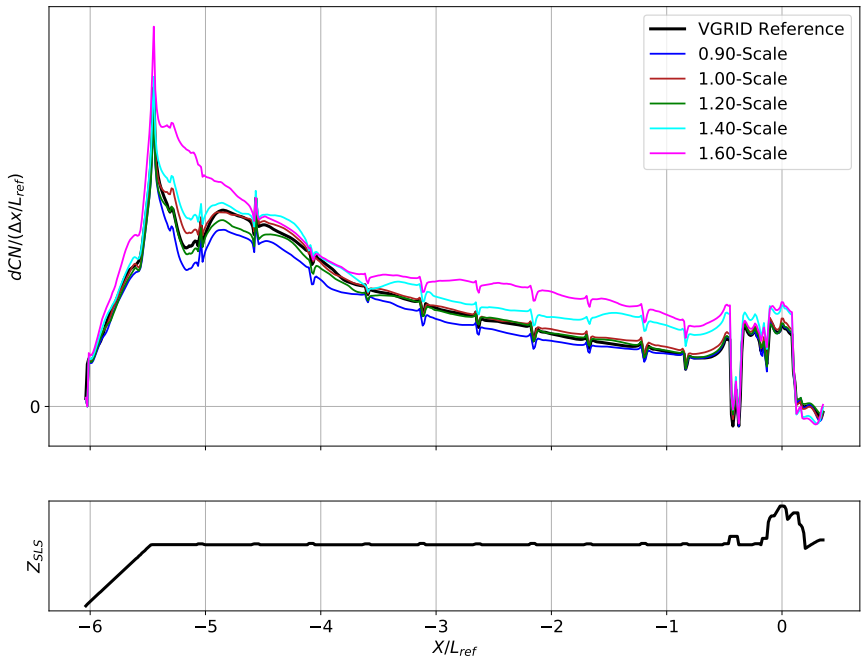

Fig. 17 Body-axis sectional normal force lineload for the starboard SRB, for various grid spacing levels at α_P = 50° and ϕ_P = 0°.

3. Comparison of Surface Pressures

Surface pressure contours on the upper and lower surfaces of the vehicle for the VGRID reference case are shown in Fig. 19 and for the baseline Heldenmesh case in Fig. 20. Though slight differences are noticeable, particularly in the region surrounding the forward attach brackets for the SRBs (just aft of the SRB noses), the results appear quite similar and appear to capture the same major flow structures on the vehicle surface.

Comparisons to centerbody surface pressure measurements taken in wind tunnel testing are shown in Figs. 21 and 22. In general, the coarser grids deviate further from the wind tunnel surface pressure measurements, even though they more closely approximate the total integrated normal force value shown in Fig 15.

The surface pressure comparisons reveal no glaring discrepancies between surface refinement values.

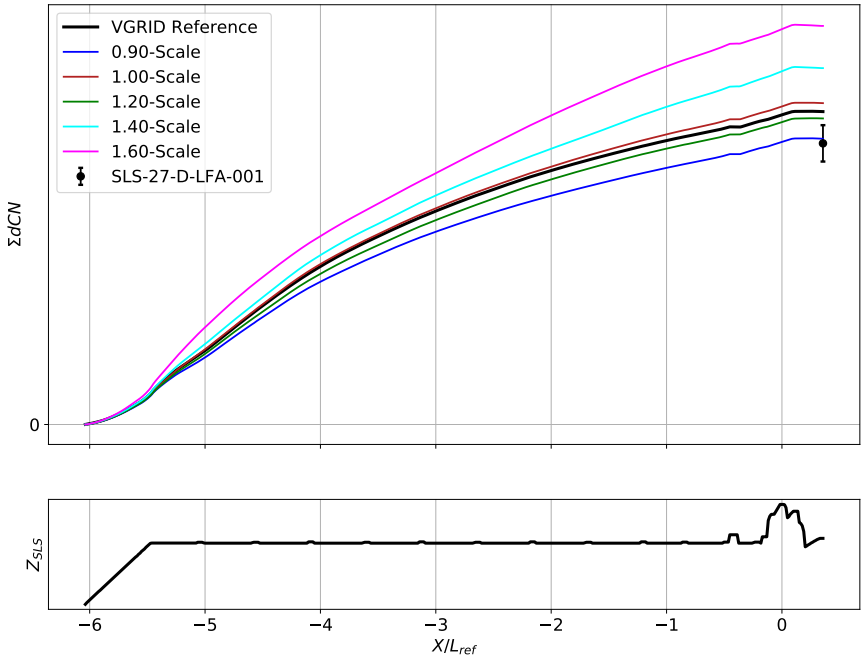

Fig. 18 Body-axis cumulative normal force lineload for the starboard SRB, for various grid spacing levels at $\alpha_P = 50^\circ$ and $\phi_P = 0^\circ$.

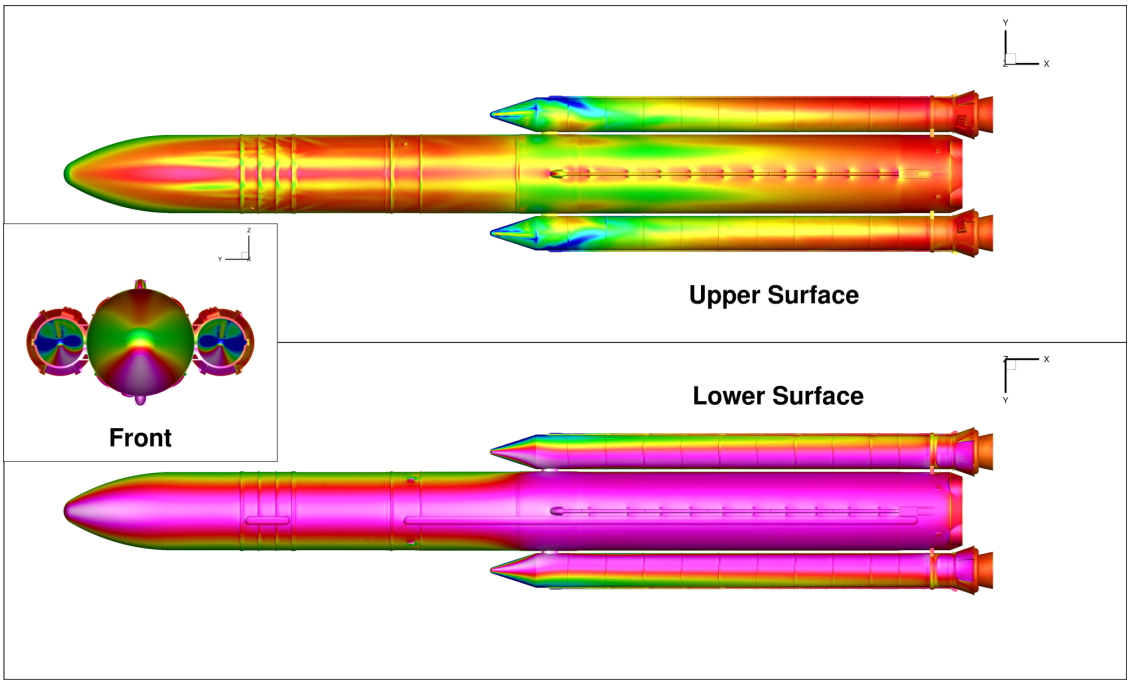

Fig. 19 Surface pressure contours for the VGRID reference grid at $\alpha_P = 50^\circ$ and $\phi_P = 0^\circ$.

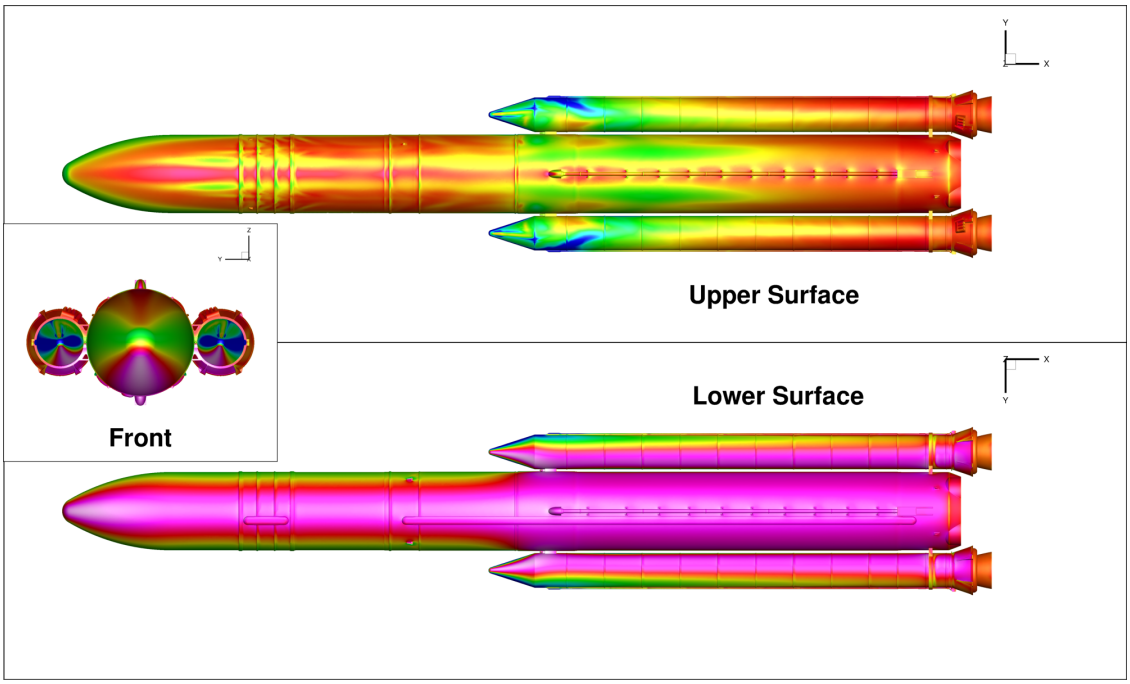

Fig. 20 Surface pressure contours for the Heldenmesh baseline grid at α_P = 50° and ϕ_P = 0°.

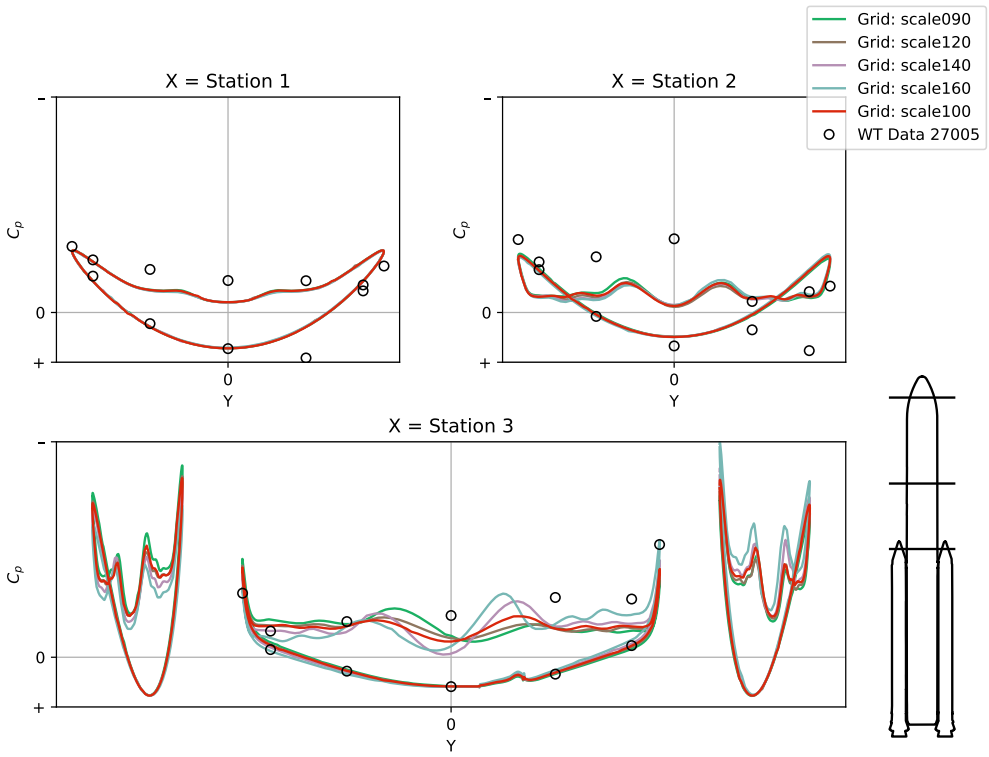

Fig. 21 Surface pressure comparisons between CFD and wind tunnel, for various grid spacing levels at α_P = 50° and ϕ_P = 0°.

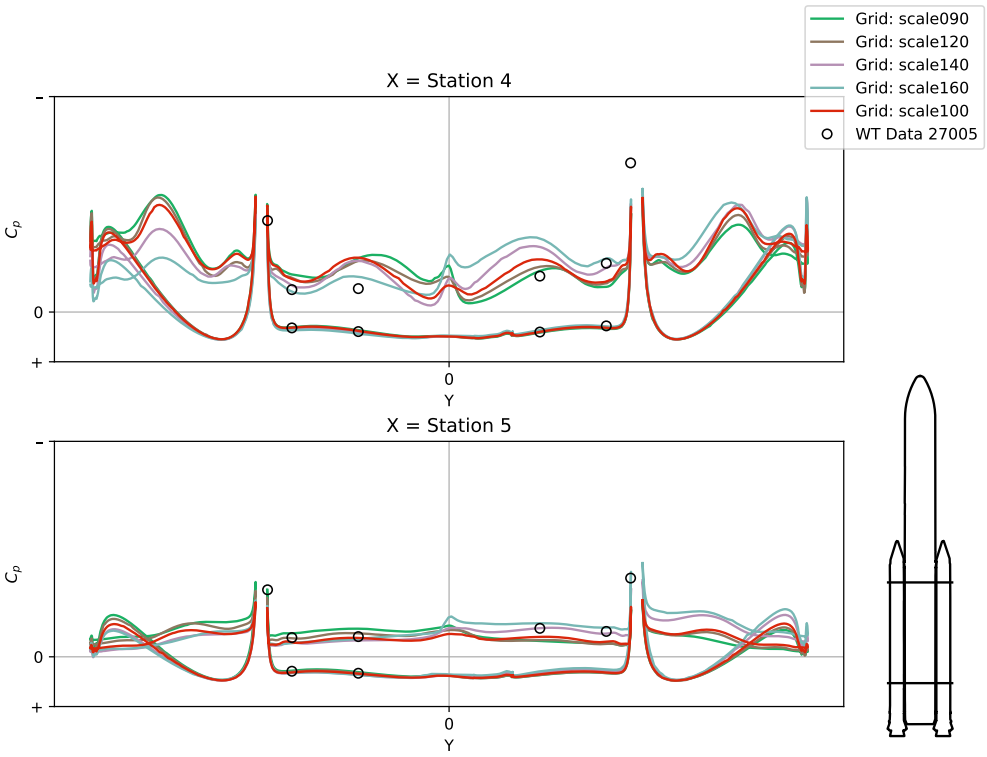

Fig. 22 Surface pressure comparisons between CFD and wind tunnel, for various grid spacing levels at α_P = 50° and ϕ_P = 0°.

IV. Effect of Volume Grid Growth Settings

Volume grids were created from the baseline 1.00-scale surface mesh that either varied R1 while keeping R2 constant, or the converse. A four-digit label was assigned to each grid, with the first two digits representing R1 in percent, and the second two digits representing R2 in percent. Thus, the baseline grid is labeled with rate 1602, representing R1 = 0.16 and R2 = 0.02. It should be noted for clarity that the 1.00-scale and the Rate 1602 volume grids are identical, and both referred to as the "baseline" grid; they are labeled differently depending on the comparison being made (source refinement or growth rate).

The near-body volume cell counts for the R1 and R2 grids are shown in Table 2. No significant variation or dependence was noted in the time required to produce a solution based on varying either R1 or R2. At first glance, this may be a surprising outcome, since the near-body volume cell counts are changing significantly between these grids. However, it may be that the differences between these grids are truncated by the near-body trim before they have a chance to become significant enough to vary the total volume cell count, which is dominated by the off-body Cartesian mesh, or to significantly change the character of the off-body Cartesian mesh itself.

Table 2 Resource requirements for various values of volume grid growth setting in Heldenmesh versus the VGRID reference grid, to reach 10,000 timesteps.

Grid	NB Vol Cells (M)	Wall Time (hrs)	Nodes	Arch	SBU Factor	SBUs	Δ
VGRID Reference	128.8	82.84	50	Has	0.8	3314	0.00%
Heldenmesh Baseline	123.4	48.4	100	Ivy	0.66	3194	-3.60%
R1=0.12 R1=0.02	134.2	51.3	100	Ivy	0.66	3386	2.18%
R1=0.14 R1-0.02	128.1	52.9	100	Ivy	0.66	3491	5.37%
R1=0.18 R1=0.02	119.6	50.3	100	Ivy	0.66	3320	0.19%
R1=0.20 R2=0.02	116.1	53.6	100	Ivy	0.66	3538	6.76%
R1=0.16 R2=0.01	127.7	49.5	100	Ivy	0.66	3267	-1.41%
R1=0.16 R2=0.03	118.3	53.1	100	Ivy	0.66	3505	5.76%
R1-0.16 R2=0.04	114.0	53	100	Ivy	0.66	3498	5.56%

Figure 23 shows very little difference between the sectional normal force lineloads for various values of volume grid growth rate 1, except in the same aft area from 0 to 2 diameters forward of the aft end of the centerbody. It should be noted that all of the volume grids for this comparison were grown from the same 1.00-scale baseline surface mesh. Combined with the fact that there is no obvious trend with Rate 1 (the 2002, 1802, and 1202 grids cluster together apart from the 1402, 1602, and VGRID reference grid in this region), this may indicate that both the surface mesh and volume grid in this local region may need to be finer to be fully grid converged. This conclusion should not be surprising, as this is a region of significantly separated wake flow at high angles of attack. In any case, Fig. 24 shows that these localized discrepancies between the solutions make very little difference to the overall cumulative lineload.

Figure 25 shows the expected variation in side force lineload on the centerbody due to the asymmetric vortex phenomenon, and the side force lineloads for the SRBs in Figs. 26 and 27 show that the phenomenon is in fact isolated to the centerbody for all grids.

The results in Figs. 28 and 29, which are for Rate 2, suggest identical conclusions as those arrived at for Rate 1.

V. Effect of Near-Body Grid Trim Distance

Resource requirements for the various grid trim values are provided in Table 3. Recall that both the Heldenmesh baseline and the VGRID reference grid have trim values of 10.8%.

Figures 31 and 32 show a similar story to the previous sections that were focused on volume grid growth rate, but shown here for varying near-body grid trim distance. There is no observable significant difference between the normal force sectional or cumulative lineloads, and the expected variation in the side force lineloads, due to varying the trim distance.

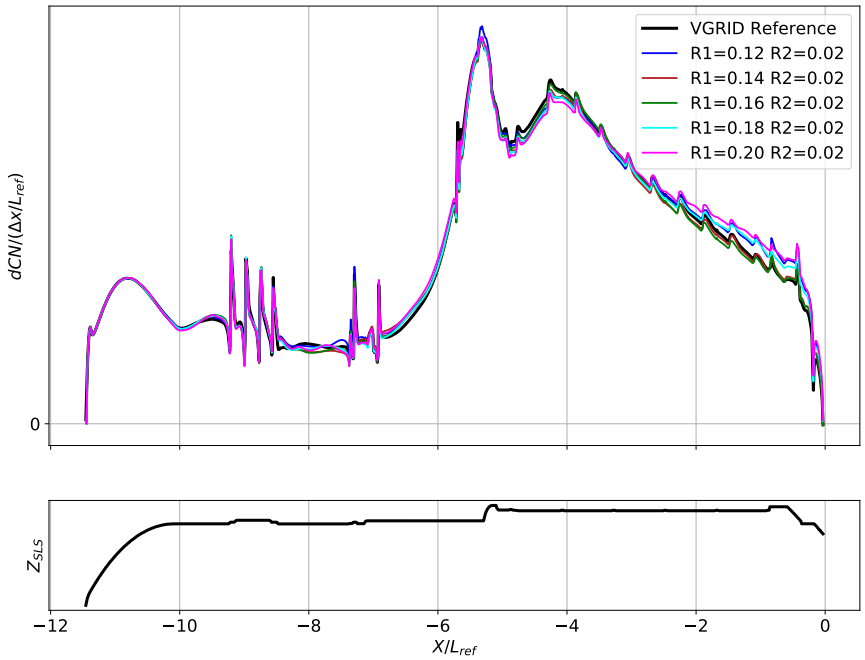

Fig. 23 Body-axis sectional normal force lineload for the centerbody, for various values of volume grid growth rate 1 at α_P = 50° and ϕ_P = 0°.

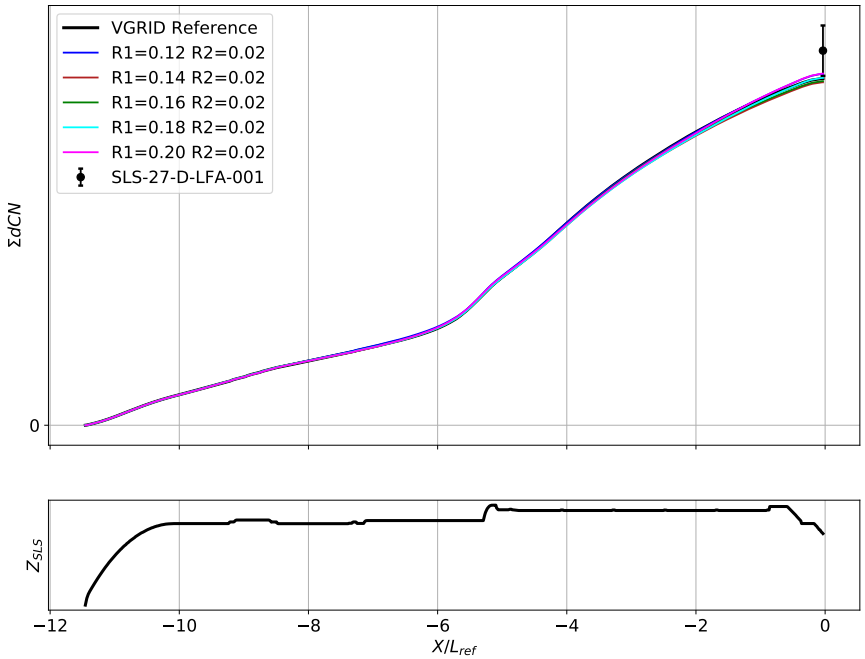

Fig. 24 Body-axis cumulative normal force lineload for the centerbody, for various values of volume grid growth rate 1 at α_P = 50° and ϕ_P = 0°.

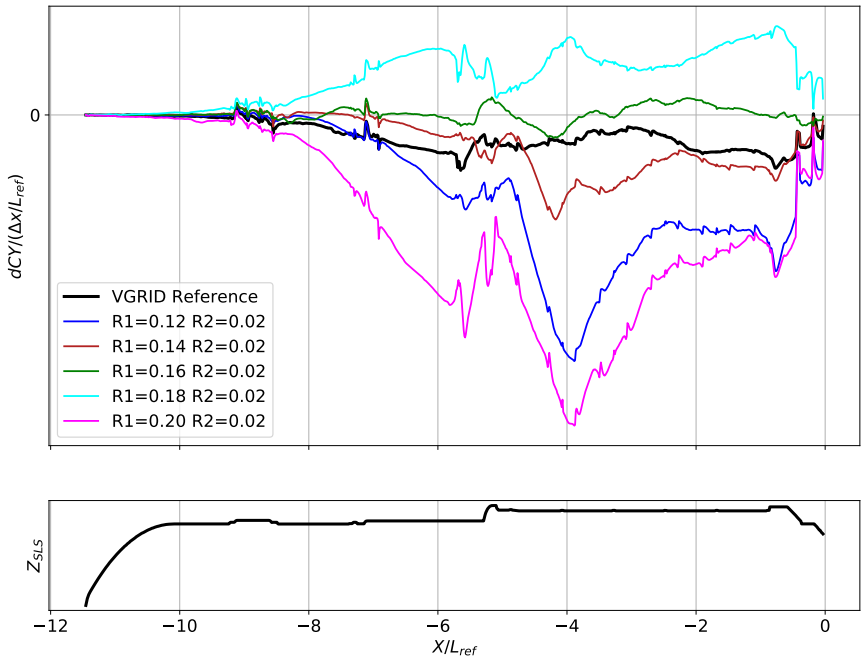

Fig. 25 Body-axis sectional side force lineload for the centerbody, for various values of volume grid growth rate 1 at α_P = 50° and ϕ_P = 0°.

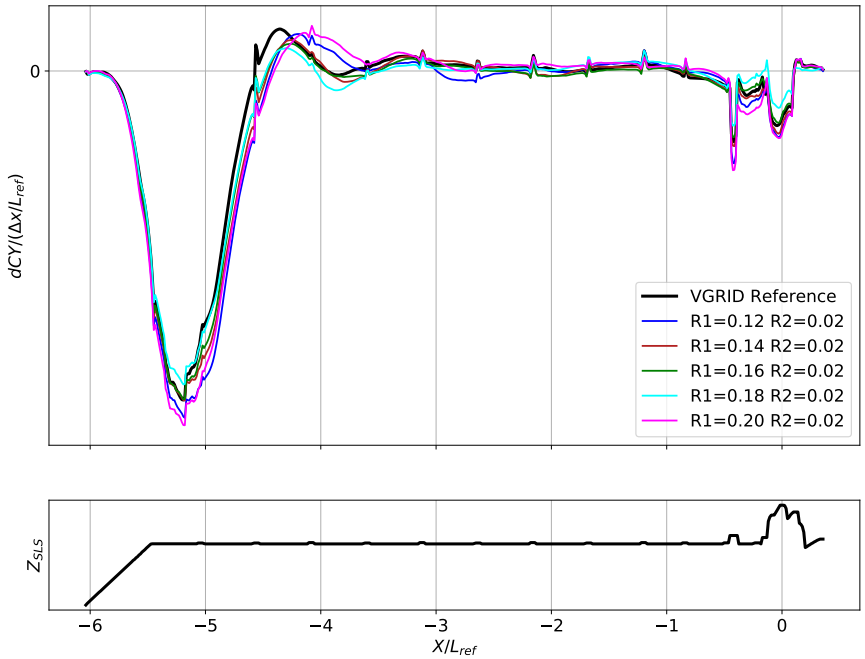

Fig. 26 Body-axis sectional side force lineload for the port SRB, for various values of volume grid growth rate 1 at α_P = 50° and ϕ_P = 0°.

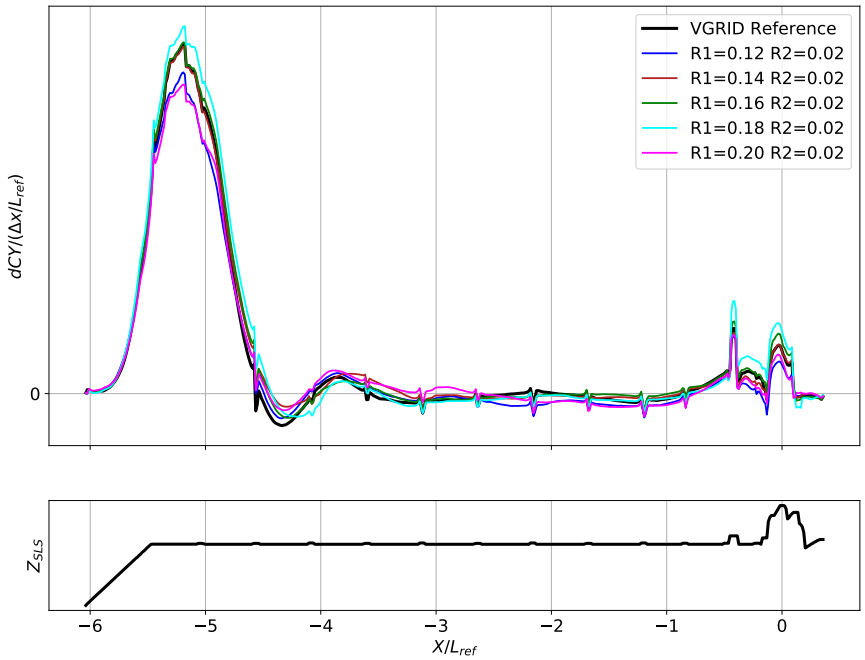

Fig. 27 Body-axis sectional side force lineload for the starboard SRB, for various values of volume grid growth rate 1 at α_P = 50° and ϕ_P = 0°.

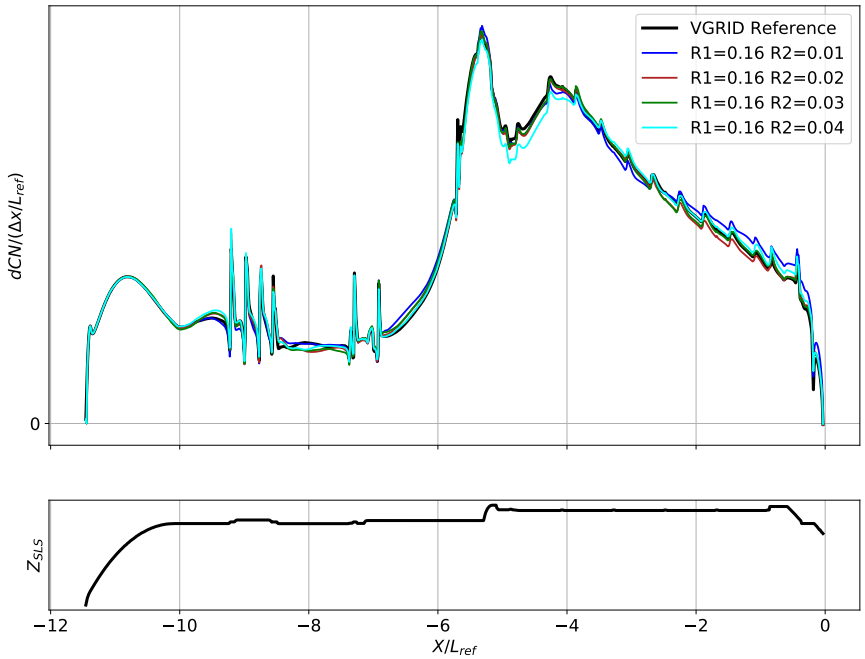

Fig. 28 Body-axis sectional normal force lineload for the centerbody, for various values of volume grid growth rate 2 at α_P = 50° and ϕ_P = 0°.

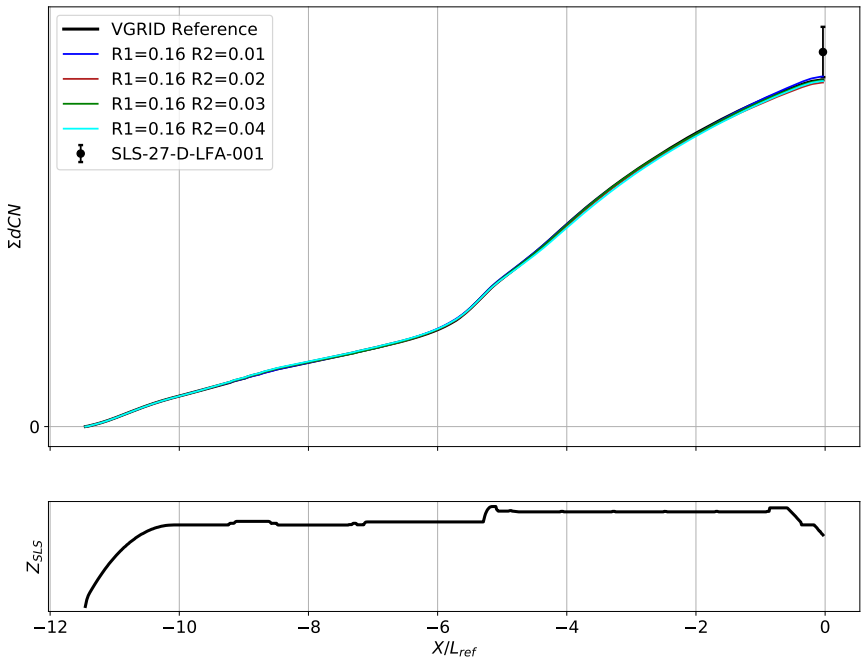

Fig. 29 Body-axis cumulative normal force lineload for the centerbody, for various values of volume grid growth rate 2 at α_P = 50° and ϕ_P = 0°.

Table 3 Resource requirements for various values of near-body trim distance in Heldenmesh versus the VGRID reference grid, to reach 10,000 timesteps.

Grid	NB Vol Cells (M)	Wall Time (hrs)	Nodes	Arch	SBU Factor	SBUs	Δ
VGRID Reference	128.8	82.84	50	Has	0.8	3314	0.00%
Heldenmesh Baseline	123.4	48.4	100	Ivy	0.66	3194	-3.60%
7.2% Trim	105.2	53.7	100	Ivy	0.66	3808	14.93%
9.0% Trim	114.7	46.9	100	Ivy	0.66	3095	-6.58%
12.6% Trim	131.6	55.1	100	Ivy	0.66	3637	9.75%

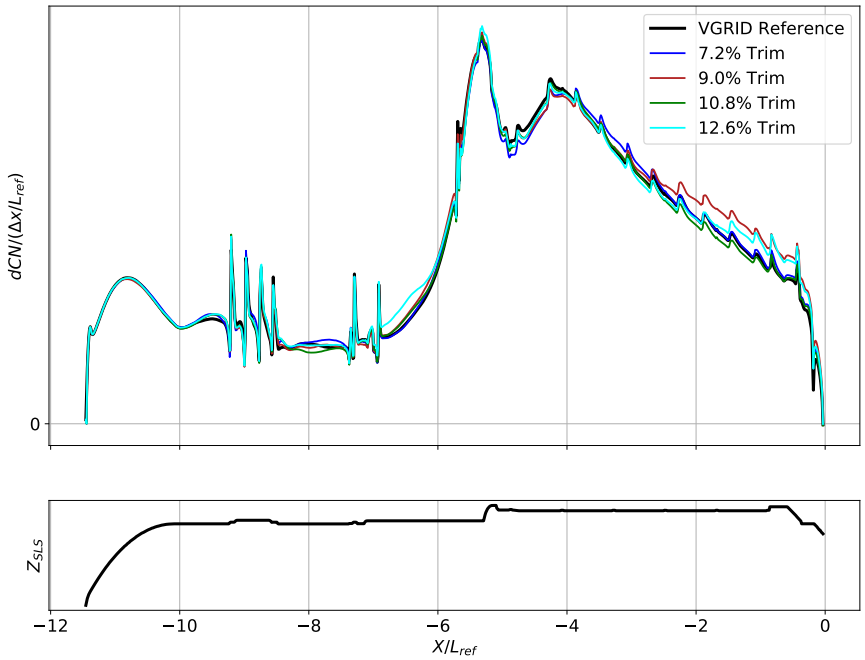

Fig. 30 Body-axis sectional normal force lineload for the centerbody, for various values of near-body trim distance at α_P = 50° and ϕ_P = 0°.

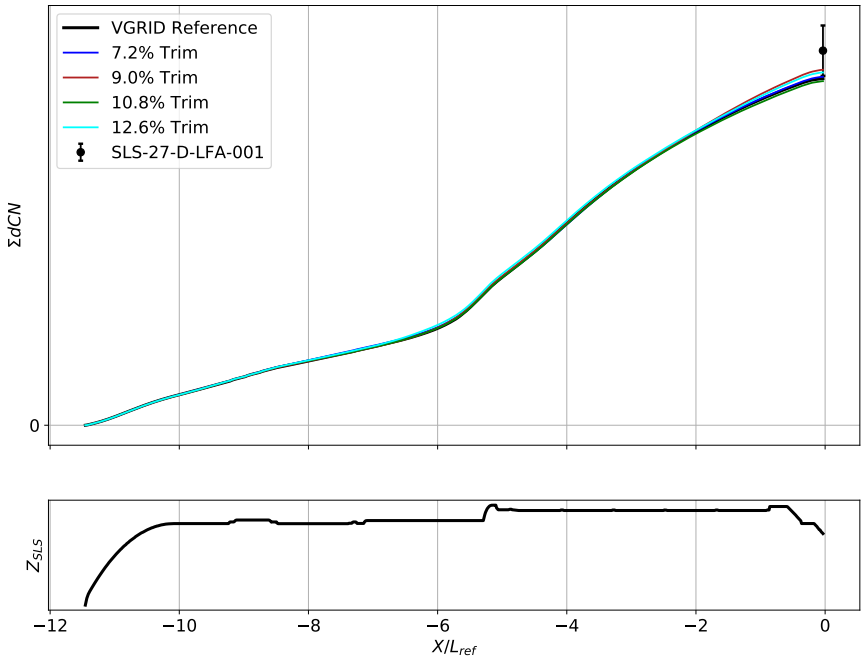

Fig. 31 Body-axis cumulative normal force lineload for the centerbody, for various values of near-body trim distance at α_P = 50° and ϕ_P = 0°.

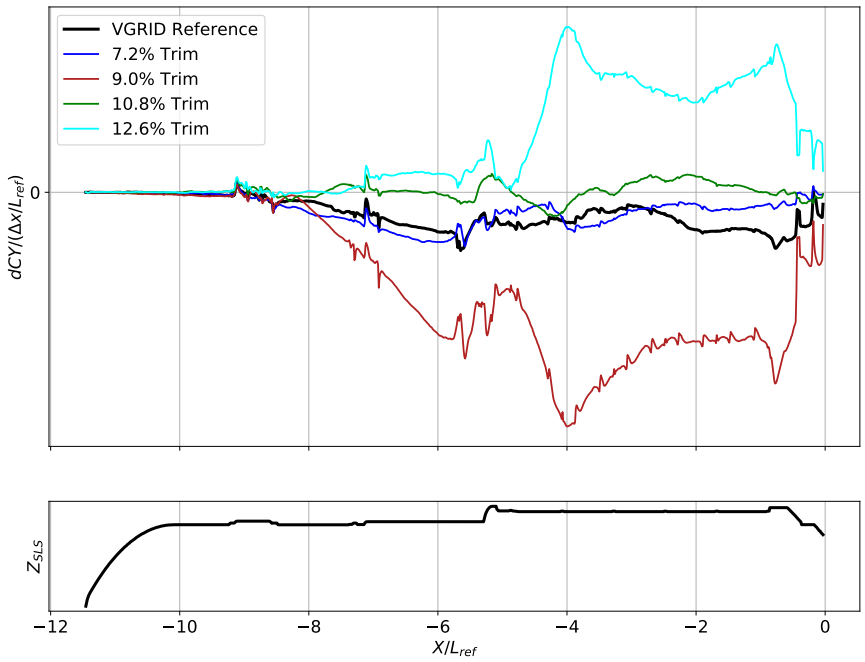

Fig. 32 Body-axis sectional side force lineload for the centerbody, for various values of near-body trim distance at $\alpha_P = 50°$ and $\phi_P = 0°$.

VI. Discussion and Conclusions

This study has demonstrated that Heldenmesh can deliver comparable grids that produce qualitatively similar results at somewhat reduced computational cost versus the team's legacy approach using VGRID/POSTGRID for a representative subset of the SLS Liftoff and Transition Lineloads problem. For the flow condition of $\alpha_P = 50°$ and $\phi_P = 0°$, which is representative of a wide range of angles of attack in the required database domain, the baseline Heldenmesh grid performed on par with the VGRID reference grid in terms of validation against experimental data for approximately 3.6% fewer computational resources required. However, when considering the fact that the Heldenmesh baseline grid had a 4.2% smaller total volume cell count than the VGRID reference, the difference in solution time does not seem to be of value.

Though exact quantitative measures of user workload were not made, the team also anecdotally noted a significant reduction in user time spent in getting from delivered geometry to first usable grid. For example, the time to generate each new test grid dropped from a few hours for VGRID to tens of minutes for Heldenmesh, enabling quicker turnaround time for inspection and adjustment. Similarly, the time required for the overall process of getting to the first grid dropped from an order of months to an order of weeks. This significant reduction in the user's time for the gridding portion of the overall workflow has made the change in tools worthwhile for the time being, and it is useful to know that no discernible degradation of solution quality or resource requirement has resulted to offset this benefit.

Aside from the choice of grid generation tool, the results of the study suggest that current best practice approaches for grid generation are more than adequate and could generally be significantly coarsened, with the exception of the surface and volume region just forward of the aft end of the centerbody. It is also likely that other localized regions of flow that experience high gradients, such as the regions immediately surrounding the SRB attach hardware, should not be coarsened and may indeed require further refinement. These local grid convergence questions will need to be resolved in more specific studies. Considering that the coarsest surface mesh studied, the 1.60-scale Heldenmesh grid, still performed roughly as well as the VGRID reference case, but with a roughly 60% reduction in resources required, there exists high potential for solution time and resource savings through more careful study of this factor. No discernible effect on either the solution quality or resources required was observed due to the range of volume grid growth settings or near-body trim distance values studied.

The limitations to this study include the fact that no angle-of-attack ranges of interest besides the midrange were explored. In particular, 90° angle of attack is a flow condition of key importance to the end users of the Liftoff and

Transition lineloads, as these cases represent the scenario in which the SLS is on the mobile launch pad prior to launch, and subject to ground wind loading from crosswinds. Further, no attempt was made in this study to determine the effect of varying off-body refinement parameters, which would likely also have a significant effect on the results given the dominance of separation and wake phenomena in the flow regime of interest.

Another limitation is that in many cases, the off-body Cartesian adaptive mesh refinement was constrained by memory limits imposed by the number and type of nodes requested for the jobs. This was in part a side-effect of the need to run many cases in a limited amount of time. As the Cartesian mesh settings may have as large or an even larger effect than the near-body grid on the solution quality and resources required, further investigations into these parameters may reveal additional improvements to the team's best practices.

References

[1] Chan, D. T., Paulson, J. W., Shea, P. R., Toro, K. G., Parker, P. A., and Commo, S. A., "Aerodynamic Characterization and Improved Testing Methods for the Space Launch System Liftoff and Transition Environment," *AIAA 2019-3398*, 2019.

[2] Abdol-Hamid, K. S., and Ghaffari, F., "Error Estimates of the Ares I Computed Turbulent Ascent Longitudinal Aerodynamic Analysis," *Journal of Spacecraft and Rockets*, Vol. 49, No. 4, 2012, pp. 609–616.

[3] Krist, S. E., and Ghaffari, F., "SLS Liftoff CFD Lessons Learned," internal presentation, SLS Aerodynamics Technical Interchange Meeting, September 2014.

[4] Morton, S. A., McDaniel, D. R., Sears, D. R., Tillman, B., and Tuckey, T. R., "Kestrel – A Fixed Wing Virtual Aircraft Product of the CREATE Program," *AIAA 2009-338*, 2009.

[5] *Kestrel Users Guide v7.1.2*, Computational Research and Engineering Acquisition Tools And Environments (CREATE), 2017.

[6] Chan, W. M., Gomez III, R. J., Rogers, S. E., and Buning, P. G., "Best Practices in Overset Grid Generation," *AIAA 2002-3191*, 2002.

[7] Ratnayake, N. A., Krist, S. E., Ghaffari, F., and Deere, K., "Computational Fluid Dynamics Methods Used in the Development of the Space Launch System Liftoff and Transition Lineloads Databases," *2019 AIAA Aviation Forum*, 2019.

[8] Frink, N. T., Pirzadeh, S. Z., Parikh, P. C., Pandaya, M. J., and Bhat, M., "The NASA Tetrahedral Unstructured Software System (TetrUSS)," *Aeronautical Journal*, Vol. 104, No. 1040, 2000, pp. 491–499.

[9] NASA High-End Computing Program, "STANDARD BILLING UNITS," https://www.hec.nasa.gov/user/policies/sbus.html, October 2019.

[10] Fidler, J., and Bateman., M., "Asymmetric Vortex Effects on Missile Configurations," *Journal of Spacecraft and Rockets*, Vol. 12, No. 11, 1975, pp. 674–681.

[11] Sohail, M., Chao, Y., Ullah, R., and Yamin, M., "Computational challenges in high angle of attack flow," *World Academy of Science, Engineering and Technology*, Vol. 80, 2011, pp. 1148–1155.

[12] Keener, E. R., Chapman, G. T., Cohen, L., and Taleghani, J., "Side Forces on a Tangent Ogive Forebody With a Fineness Ratio of 3.5 at High Angles of Attack and Mach Numbers From 0.1 to 0.7," *NASA-TM-X-3437*, 1977.

[13] Hartwich, P., Hall, R., and Hemsch, M., "Navier-Stokes Computations of Vortex Asymmetries Controlled by Small Surface Imperfections," *AIAA 90-0395*, 1990.