

# Kestrel Results at Liftoff Conditions for a Space Launch System Configuration in Proximity to the Launch Tower

Steven E. Krist, Ph.D. <sup>\*</sup>, Nalin A. Ratnayake<sup>†</sup>, and Farhad Ghaffari <sup>‡</sup>  
*NASA Langley Research Center, Hampton, VA, 23681-2199*

Aerodynamic data books for Space Launch System vehicles require databases for the integrated forces and moments and section loads during liftoff and transition to the ascent phase of flight. While the force and moment database can be generated from wind tunnel results, computational analyses are necessary to provide the extensive surface information required to generate proper line loads. Of the two flight regimes, the liftoff problem is the more costly and complex situation to simulate, as it requires modeling of the vehicle in proximity to the launch tower. The effects of massive separation on the leeward pressure fields of both the tower and vehicle are not well captured with RANS methods, necessitating the use of more advanced methods, such as Delayed Detached Eddy Simulation, in conjunction with computational grids sufficiently refined to resolve the wakes. Details on the computational setup for employing the Kestrel flow solver to address the liftoff problem are presented. The methodology involves the use of independent unstructured near-body grids for the vehicle and the tower, overset by a solution adaptive Cartesian off-body grid. Results from the simulations are compared to experimental results from a test in the NASA Langley Research Center 14- by 22-Foot Subsonic Tunnel.

## I. Nomenclature

<i>AMR</i>	=	Adaptive Mesh Refinement
<i>CFD</i>	=	Computational Fluid Dynamics
<i>DDES</i>	=	Delayed Detached Eddy Simulation
<i>DoD HPCMP</i>	=	Department of Defense High Performance Computing Modernization Program
<i>GNC</i>	=	Guidance, Navigation, and Control
<i>LaRC</i>	=	Langley Research Center
<i>LAS</i>	=	Launch Abort System
<i>LES</i>	=	Large Eddy Simulation
<i>LOT</i>	=	liftoff and transition (phase of flight)
<i>MSFC</i>	=	Marshall Space Flight Center
<i>MPCV</i>	=	Multi-Purpose Crew Vehicle
<i>NASA</i>	=	National Aeronautics and Space Administration
<i>NBOB</i>	=	Near-Body, Off-Body (dual grid strategy)
<i>OML</i>	=	Outer Mold Line
<i>RANS</i>	=	Reynolds-Averaged Navier-Stokes
<i>SA</i>	=	Spalart-Allmaras (turbulence model)
<i>SLS</i>	=	Space Launch System
<i>SRB</i>	=	Solid Rocket Booster
<i>14x22</i>	=	NASA LaRC 14- by 22-Foot Subsonic Tunnel
$\alpha_P$	=	total angle of attack (missile axis system), deg
$\phi_P$	=	aerodynamic roll angle (missile axis system), deg

---

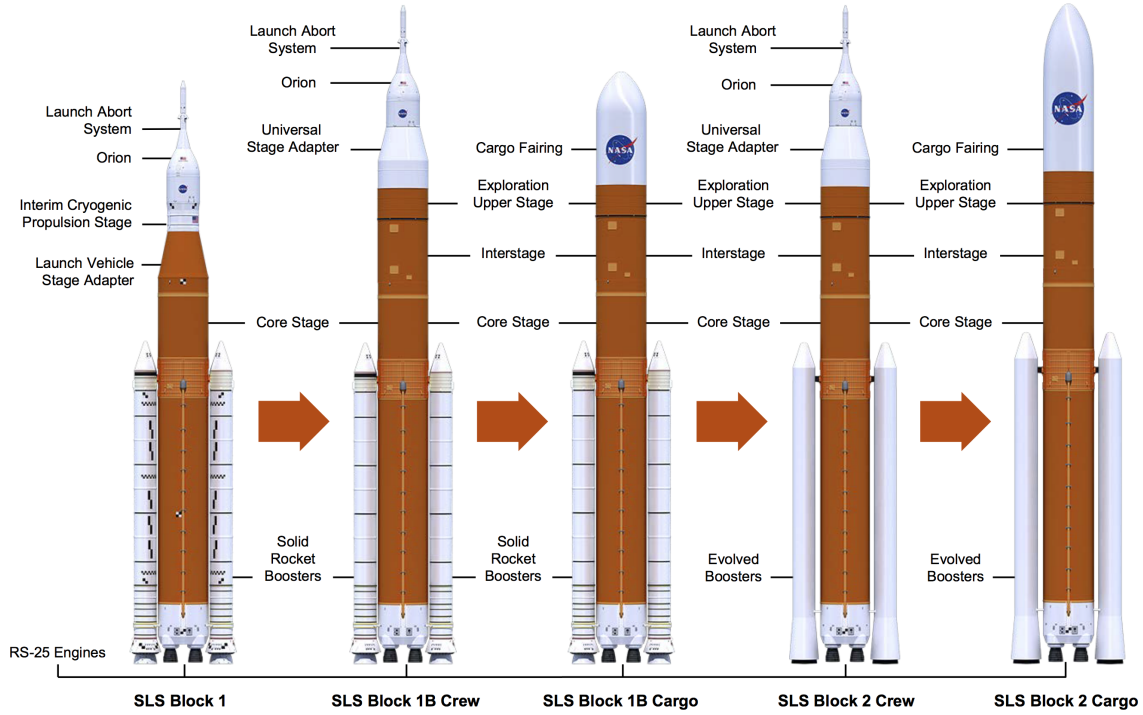
<sup>\*</sup>Research Aerospace Engineer, Configuration Aerodynamics Branch, AIAA Senior Member

<sup>†</sup>Research Aerospace Engineer, Configuration Aerodynamics Branch, AIAA Senior Member

<sup>‡</sup>Research Aerospace Engineer, Configuration Aerodynamics Branch, AIAA Associate Fellow

## II. Introduction

The NASA Space Launch System (SLS) is comprised of a family of launch vehicles designed to provide a new capability for deep-space exploration. Variants of the SLS launch vehicle embody a common, multistage stack that is comprised of the core containing the liquid-fueled main rocket engines, upon which two five-segment Solid Rocket Boosters (SRBs) are attached to the sides. For the SLS-10000 (Block 1 Crew) and SLS-28000 (Block 1B Crew) series, the Orion Multi-Purpose Crew Vehicle (MPCV) sits atop the stack and the Launch Abort System (LAS) shroud encloses the MPCV. For the SLS-27000 (Block 1B Cargo) variant, the MPCV/LAS is replaced with a payload enclosed within a shroud. Figure 1 shows the three configurations of interest in the present study; it also depicts the Block 2 Crew and Cargo variants, which incorporate the next generation SRBs.




**Fig. 1 SLS Vehicle Configurations.**

The SLS liftoff and transition (LOT) lineloads databases provide the distributed aerodynamic forces and moments acting along the length of the vehicle, from vehicle on-pad, through liftoff, then through transition to the nominal ascent trajectory. These distributed loads are provided as a function of axial location with individual data tables supplied for the central core stage and each of the two SRBs. The lineloads database for a configuration is comprised of two separate data sets, one for liftoff, where significant aerodynamic interactions between the vehicle and launch tower occur, and one for transition, which covers the vehicle in isolation after it has cleared the launch tower, but prior to reaching the nominal ascent phase of flight which starts at Mach 0.3. The primary end users of the LOT lineloads databases are the Loads group and the Guidance, Navigation, and Control (GNC) team at the Marshall Space Flight Center (MSFC). The Loads group primarily uses the databases for analysis of bending and shear on the vehicle. The GNC group primarily uses them for sensitivity analysis related to the vehicle trajectory and control system during the LOT phase of flight.

While the overall vehicle force and moment coefficients can be efficiently obtained through wind tunnel testing, computational analyses have become indispensable in obtaining the extensive amount of surface information (i.e., distributed pressure and skin friction coefficients) required to generate accurate lineloads. This paper provides a summary of the investigations conducted by the NASA Langley Research Center (LaRC) SLS Computational Fluid Dynamics (CFD) team in support of aerodynamic database development for the liftoff phase of flight. Particular challenges and decisions associated with the development of CFD for the SLS vehicles in the transition phase are described in the companion paper by Ratnayake, et al.[1].

### A. The Liftoff Problem

The pre-launch arrangement of an SLS-10000 configuration on the mobile launch platform is shown in Figure 2; an artist's rendition of the Block 1 Crew vehicle on the existing launch pad is shown on the left; an engineer's rendition of the Outer Mold Line (OML) for the corresponding CFD model is shown on the right. The CFD model does not render every feature of the actual vehicle, but it does retain many of the smaller features, such as the cameras, latches, and the brackets on the feed lines. On the other hand, the CFD model of the launch tower is simplified to an extreme degree; the skeleton of the support structure and major features like the elevator shaft and the four equipment rooms are modeled, but that is about it.




**Fig. 2** Artist's rendition of the Block 1 Crew vehicle on the existing launch pad (left) and the computational OMLs for the corresponding CFD models (right).

The liftoff phase of flight begins as the vehicle leaves the pad and ends after the vehicle clears the launch tower, entering the transition phase of flight. Over the liftoff flight segment, the total angle of attack experienced by the vehicle,  $\alpha_P$ , starts at  $90^\circ$  and decreases rapidly to around  $10^\circ$ . As depicted in Figure 3, the position of the vehicle during liftoff is measured by the vehicle liftoff ratio  $h/L$ , where  $h$  is the elevation of the vehicle above the launch platform and  $L$  is the total height of the tower. By definition, the vehicle clears the tower at  $h/L = 1$ , but interference effects from the tower are still measurable as the vehicle continues to rise; the upper limit of the liftoff regime is defined to expire at  $h/L = 1.5$ , at which point launch tower interference effects are miniscule.


For transition testing, an SLS model is mounted on a sting and suspended in free air, oriented from  $\alpha_P = 0^\circ$  to  $90^\circ$  through roll angles ( $\phi_P$ ) from  $0^\circ$  to  $360^\circ$ . For liftoff testing, an SLS model is mounted on a vertical sting that can freely traverse up and down in the proximity of the launch tower, maintaining a constant horizontal distance away from the tower. Hence, the total angle of attack for liftoff testing is necessarily fixed

at  $90^\circ$ . However, the intent of liftoff testing is to develop a database that captures the interference effects of the tower on the vehicle forces and moments at all angles of attack. In order to resolve the incompatibility between what can be tested and what happens in flight, it is simply assumed that the tower effects do not depend on the vehicle angle of attack. Under that assumption, the liftoff force and moment database can be constructed as an incremental database in the  $(h/L, \phi_P)$  parameter space; at a given roll condition, the transition force and moment values at  $\alpha_P = 90^\circ$  are subtracted from the liftoff force and moment values at the appropriate value of  $h/L$  to provide the increment. Likewise, an incremental liftoff line loads database can be constructed by subtracting off the line loads from CFD simulations for the transition problem at  $\alpha_P = 90^\circ$  from the CFD simulation-generated line loads for the liftoff problem.

Figure 4 shows the orientation of the launch vehicle with respect to the tower and defines the azimuth angle for ground winds, which is the same as the definition of the roll angle,  $\phi_P$ , for both the liftoff and transition flight regimes. As shown, the tower lies north of the vehicle. Consequently, a wind blowing out of the north,  $\phi_P = 0^\circ$ , blows through the tower prior to impacting the bottom of the vehicle and imparts a positive normal force onto the vehicle. Tower interference effects are required at all roll angles and at various vehicle liftoff ratios from  $h/L = 0$  to 1.2.



**Fig. 4 Orientation of the vehicle with respect to the launch tower and the definition of the roll angle.**



**Fig. 3 Definition of the vehicle liftoff ratio  $h/L$ .**

Experimental testing of the SLS configurations in the LOT regime was conducted in the NASA LaRC 14- by 22-Foot Subsonic Tunnel (14x22). For the full scale vehicle in a 30 knot wind, the Mach number at liftoff conditions is around 0.05 and Reynolds number is close to 10 million. However, the flow in the 14x22 is unsteady at dynamic pressures below 30 psf, which does not provide an ideal situation for force and moment testing. Moreover, grit is spread abundantly over the vehicle surfaces to ensure that the flow over most of the model is turbulent; exploratory tests of the SLS models indicated that the grit was most effective at dynamic pressures above 50 psf, which corresponds to a Mach number of 0.183. Hence, most of the test was run at Mach 0.183; at that speed the Reynolds number is 615,000 based on the nominal core diameter of the model. Wind tunnel and force and moment database results shown in the plots below are from Test 633, which was the second LOT test for the SLS vehicles. The models used in Test 633 were much more detailed than those used in Test 609, which was the first entry into the 14x22 for SLS LOT testing.

## B. Computational Considerations for the Liftoff Problem

The fundamental difficulty in numerical simulations for the liftoff flight regime is accurately capturing the effects of large wakes on the pressure field about the leeward surface of the vehicle. It is common knowledge in the CFD community that Reynolds-Averaged Navier-Stokes (RANS) methods are inadequate to the task [2]. Since the application of Large Eddy Simulation (LES) methods to liftoff problems is still a decade or more away from being feasible, it is necessary to resort to hybrid RANS/LES methods, such as Detached Eddy Simulation methods and its many variants, e.g., Delayed Detached Eddy Simulations (DDES). As described by Spalart [3], DDES methods essentially involve using RANS methods in the boundary layer region and LES-like methods for the extended wake. DES methods can be used with any turbulence model that makes use of a properly defined length scale. However, the computational grid must be sufficiently refined on the leeward side of the body in order to adequately capture the small eddy content of the wake, making DDES methods much more resource intensive than RANS methods.

The first LOT line loads database developed in the SLS program was a transition database for the SLS-10003 configuration. The computations were generated with the USM3D flow solver [4], results from which are compared



to data from 14x22 Test 609 in the paper by Krist et al.[5]. A subsequent in depth investigation into the application of hybrid RANS/LES methods for the transition problem was conducted by Ratnayake, et al.[1], in which three flow solvers, USM3D, FUN3D [6] and Kestrel [7] were evaluated. Results from a series of simulations conducted for the Block 1 Crew configuration led to the conclusion that the Kestrel flow solver was the most appropriate tool for the transition problem, not because solution quality was any better than that from USM3D or FUN3D, but because it required significantly less computer resources. Moreover, while simulations of the transition problem demand the use of significant computational resources, the liftoff problem raises the resource requirements to a whole new level. Consequently, the Kestrel flow solver was selected for simulations of the liftoff problem.

### III. Configurations and Computational Approach

Simulations of the liftoff problem using the Kestrel flow solver have been conducted for the Block 1 Crew and Block 1B Crew and Cargo configurations. Most of the following discussion will be in regard to the Block 1B Crew vehicle, also referred to as the SLS-28005 configuration, since that was the first vehicle on which the CFD workflow for the liftoff problem was developed.

#### A. Vehicle and Launch Tower Configurations

The wind tunnel and CFD models for the SLS-28005 configuration are both based on the same definition. While the CFD model incorporated a few minor simplifications to the OML in order to round off some features that would otherwise cause grid quality issues, the OML of the CFD model is nearly the same as the as-built OML of the wind tunnel model. A rendering of the CFD model is shown in Figure 5. The configuration consists of the core stage with port and starboard SRBs, both of which are attached to the core with a set of forward and aft brackets. The MPCV is mounted on top of the core stage, but does not actually contribute to the OML of the vehicle, since the shroud at the end of the LAS fully encapsulates the MPCV. The CFD model incorporates all of the protuberances that were included on the wind tunnel model. The metric portion of the vehicle, upon which integrated forces and moments are measured, includes the core stage from the tip of the LAS to termination of the nozzle aerodynamic fairings (excludes the black base surfaces, the cavity, and the sting) and the SRBs up to the cap on the SRB nozzles (purple and magenta surfaces). Note that separate line loads are required for the core stage and the port and starboard SRBs, and these are to be computed over the entire length of the respective metric surfaces.





Fig. 5 The Block 1B Crew (SLS-28005) CFD OML.

The models for the Block 1B configuration are quite complex. Though hard to see in the figure, there are eleven additional structural rings on each of the SRBs, four additional rings at the top of the core body, and significantly more detail in the representations of the attach brackets, feed lines, and system tunnels than that included on any of the models used in previous analyses. An indication of how meticulous the expectation for CFD modeling has become can be gleaned from the inset figure, which highlights the detail with which the feed line brackets are represented. Given that there are nine brackets on each of the two feed lines, the increase in grid size due to modeling of the brackets is substantial. Unfortunately, there was not time or resources to construct and run a CFD model without the brackets to see if the influence of the brackets can be detected.

A photograph of the Block 1B Cargo wind tunnel test article in proximity to the launch tower, installed in the 14x22 during Test 633, is shown on the left side of Figure 6; a rendering of the corresponding CFD models for the case with  $h/L = 0.5$  is shown on the right side of the figure. Since the wind tunnel model already represents a significant simplification of the actual launch tower, it was not necessary to make many changes to that definition in creating the CFD OML. In fact, the only significant differences between the two are that the floors at each level of the tower have been removed in the CFD OML (except around the equipment rooms), the CFD model of the platform is missing some features, and the sting is shorter in the vehicle CFD model. The motivation for excluding the floors at most of the levels is that their inclusion requires viscous grids on both the upper and lower surface of each floor, causing a substantial increase in grid size. The reason for using such a small sting on the vehicle model is that, with the length shown, the same grid can be used for all  $h/L$  values greater than 0.1. For cases with  $h/L = 0$ , the vehicle model is modified in that the sting is reduced to a mere stub, while the launch tower model is modified to include the Tail Service Masts (the "tombstones" that sit on the south side of the platform) and the hole in the platform through which the sting protrudes.



**Fig. 6 The Block 1 Cargo vehicle installed in the wind tunnel (left) and the corresponding CFD models (right).**

Another feature that is barely discernible in the photograph is that the launch tower apparatus is attached to a turntable so that testing can proceed through a  $330^\circ$  range in roll angle; limitations on the turntable preclude accessing full range of  $360^\circ$ . Since the SLS configurations are fairly close to being symmetric, it is assumed that the interference effects of the tower at a roll angle of  $90^\circ$  will be similar to that at  $270^\circ$ . Consequently, in order to accommodate the limitation in the turntable, the tower is mounted so that the restriction occurs for roll angles from  $75^\circ$  to  $105^\circ$  and the force and moment database values over that range are filled in by mirroring the data from  $255^\circ$  to  $285^\circ$ .

A significant modification of the test procedures was also implemented between Test 609 (for the previous design cycle) and Test 633; the decision was made to gain component-level force and moment breakdowns at the cost of pressure

information on the SRBs. By installing load cells in both the forward and aft attach brackets, forces and moments on the core and each SRB could be distinguished. However, the installation precluded the ability to instrument the pressure ports on the SRBs, taking away the ability to discern how well CFD pressure distributions on the SRBs correlated with wind tunnel pressure port values. Hence, it became easier to discern whether the CFD was capturing the correct forces and moments on the three major components, but more difficult to tell what was causing the differences that arose.

## **B. Kestrel Flow Solver Methodology**

Kestrel is a virtual fixed wing aircraft simulation tool under constant development within the Computational Research and Engineering Acquisition Tools and Environment, Air Vehicle (CREATE-AV) project, run under the auspices of the Department of Defense High Performance Computing Modernization Program (DoD HPCMP). The Kestrel software was developed in response to a DoD need for quick turnaround, high-fidelity, multidisciplinary analyses of complex configurations and scenarios in order to support acquisition efforts [7]. Kestrel utilizes a Common Scalable Infrastructure (CSI) to implement a modular approach for connecting the computational tools required to conduct multidisciplinary assessments of fixed-wing aircraft. The CSI architecture allows for the incorporation of a variety of components for assessment of the aerodynamics, structural dynamics, kinematics, and kinetics involved in maneuvers, store separation, and other complex aerodynamic scenarios at subsonic, transonic, and supersonic flight conditions. Of particular interest here, the architecture enables a dual mesh flow solver methodology, referred to as the near-body, off-body (NBOB) methodology. In the NBOB methodology, the near-body flow is computed on an unstructured grid that encapsulates the solid body, and the unstructured grid is immersed in a Cartesian mesh that extends to the far field. Within the overset domain, where the near-body and off-body grids overlap, hole cutting, connectivity operations, and interpolation between the two domains are handled with a python-based infrastructure that makes use of the Parallel Unsteady Domain Information Technology software toolkit.

The dual mesh approach is particularly well suited for situations with regions of massively separated flow. In such cases, the extensive wake on the leeward side of the body can be resolved in the off-body Cartesian mesh, for which operations are significantly cheaper to implement than operations on the near-body unstructured grid. Moreover, the Cartesian mesh is well suited for the use of higher-order spatial discretizations and is amenable to the implementation of solution Adaptive Mesh Refinement (AMR). Both of these capabilities aid in obtaining the greatest possible accuracy while running on the smallest possible grid. For Kestrel, the most efficient way to run simulations is to solve as much of the problem as possible on the off-body Cartesian mesh, with as little as possible on the near-body unstructured grid.

The Kestrel flow solver for the near-body domain is KCFD, a cell-centered, finite volume method for solving the Euler and RANS equations on both standard tetrahedral grids and mixed element unstructured meshes, which can have arbitrary combinations of tetrahedra, prisms, and pyramids. Available turbulence models include a number of the Spalart-Almaris (SA) variants and the Mentor two-equation model with and without Shear Stress Transport. In addition to the traditional RANS capability, KCFD provides an implementation of the DDES method for use in situations involving massively separated flow. The CFD code employed on the off-body domain is SAMAir, a finite difference method for solving the Euler and RANS equations on Cartesian meshes that is derived from an earlier code, ARC3DC, developed at NASA Ames Research Center. SAMAir is capable of implementing all of the models available in KCFD, allowing for a uniform implementation of the governing equations in both solvers.

Generation of the initial Cartesian off-body grid, as well as implementation of the solution adaptive grid refinement, is quite simple for the Kestrel user, as both processes are automated within the Kestrel MeshManager component. In order to generate the initial Cartesian off-body grid, the only detail to be specified by the user is the extent of the computational domain. The MeshManager is then able to grab enough information from the unstructured grid (or grids) to compute a characteristic length for the unstructured grid fringe points, which is then used to size the finest level of the Cartesian mesh. At that point, increasingly coarser layers are defined and the actual extents of the outer boundaries are backed out. If the user is dissatisfied with the compatibility of the unstructured and Cartesian cells in the overset fringe, an additional input provides the option for adding additional layers or removing layers from the Cartesian mesh.

Implementation of the AMR is a bit more complex than generation of the initial Cartesian mesh, but follows the same sort of process once cells for refinement and derefinement have been flagged. In this case, the user may specify multiple zones in which the AMR is to be applied, supplying the maximum and minimum values of the corner points of the refinement box. There are a number of options for specification of the refinement variable, which is used to trigger the refinement or derefinement of a cell based on whether it is larger or smaller than a specified threshold value.

There are also a number of useful tools that come with the Kestrel suite of codes. The two tools used in this work were the KUI (Kestrel User Interface) and a grid manipulation program named Carpenter. KUI is a GUI for setting up

the input deck. It not only helps with organization of the problem topology, but sets default values for many parameters and prevents the selection of incompatible options. Carpenter is a grid manipulation tool that transforms a number of grid formats into the Kestrel preferred .avm format and provides several other useful functions. Perhaps the most widely used option is the trim procedure, which allows one to trim the outer boundary of an unstructured grid down to a specific distance from the solid surface, thereby providing a simple means to take just about any unstructured grid and turn it into an appropriate near-body grid for use with Kestrel. Another useful option is the conversion of tetrahedral unstructured grids into mixed element unstructured grids. The advantage of mixed element grids is that most of the tetrahedra in the boundary layer of an unstructured grid can be converted into prisms, for which the computational procedures tend to be much more efficient.

### C. Unstructured Grid Development

The strategy used in constructing unstructured grids for the SLS vehicles was based upon five years of work in developing best practice procedures for evaluating the SLS configurations in the transition regime of flight, as discussed in the work by Ratnayake, et al.[1]. The grids were generated using the Tetrahedral Unstructured Software System (TetrUSS), which is developed and maintained at LaRC. Three tools from the TetrUSS suite of utilities are used in the grid generation process; GridTool, VGRID, and POSTGRID. The three components have been applied to and validated on a wide array of aerodynamic problems over the course of roughly two decades [8].

The grid generation process begins with importation of the CAD geometry into GridTool, where surface patches covering the entire OML are defined, then grouped to form components. Of primary importance in generating accurate solutions is the source specification. Since the grid spacing and local density of the cells are related to the strengths and locations of the sources, appropriate placement of the sources by the user refines the grid where it is needed, namely, in regions of high flow field gradients. Line sources with anisotropic stretching are typically employed where components intersect or around features that embody sharp discontinuities in the surface. Volume sources in the shape of a cylinder, cone, or sphere are employed within the perimeter of the larger components of the vehicle. A screen shot highlighting the source specification for the Block 1B tower used in  $h/L = 0.0$  liftoff cases is presented in Figure 7; yellow lines indicate a line source and pink cylinders indicate a volume source. In configurations like the tower, for which most beams feature four discontinuities, a line source ends up being specified for just about every edge on the OML.

Output from GridTool is imported to VGRID, which employs a three-step process to generate an unstructured volume mesh composed of tetrahedra. The first step generates a surface triangulation, the local resolution of which is controlled by the source specification. The second step employs an advancing layer method to grow a boundary layer grid normal to the surface. The third step employs an advancing front method to discretize the inviscid flow region. The final step is to use the tool POSTGRID to close holes in the volume grid left by the VGRID process.

In developing grids for the USM3D and FUN3D flow solvers, volume sources must be extended well out from the body in order to adequately resolve the eddy content of the wakes. However, the volume sourcing must be restricted to the wake region as much as possible, since refining regions of the flow where no action is taking place increases the resource requirements of the simulation without providing any benefit. Hence, a separate grid must be generated for each roll angle so that the refined grid region can be limited to the vicinity of the wake. No such problem exists with grids for the Kestrel flow solver. Since the unstructured grids are to be converted to near-body grids by trimming the outer boundary back to a small distance from the solid surface, generating refined wake grids around the entire vehicle does not cause the final trimmed grid to be that much larger than a grid for which only the wake region for a specific roll angle is refined. A big benefit of using the NBOB strategy in combination with AMR on the Cartesian grid, is that a single near-body grid will be valid for all roll angles.

Using this sourcing strategy for Kestrel grids, previous input files developed for USM3D grids were slightly tweaked so that the volume sources were centered along the centerlines of the core and the two SRBs, rather than being offset to better enclose the wake region. In addition, separate grids were developed for the  $h/L = 0.0$  cases and the nonzero cases; a rather short sting, suitable for runs down to  $h/L = 0.1$ , was used for the nonzero cases, whereas a stubby sting that only went down a short distance into the hole in the platform was used for the  $h/L = 0.0$  case. The input file to GridTool for the SLS-28005 configuration consisted of 2328 patches and 6579 line and volume sources. Upon generating the grid, Carpenter was used to convert the grid to a mixed element grid and trim it down to 24 inches. The resulting near-body grid contained 121 million cells, of which 56 million were tetrahedrons, 61 million were prisms, and 4 million were pyramids.


The initial plan for generating the tower grid was also to use the grid generation tools from TetrUSS. The tower model was dissected into 1457 patches, and source sizing was kept consistent with the source sizing on the launch

vehicle. Eventually a total of 6700 sources were distributed on the tower OML. Unfortunately, the TetrUSS software was unable to generate a valid grid for the launch tower. While VGRID would frequently proceed through its paces, POSTGRID was unable to close all the holes in the volume grid left by the VGRID process. This type of issue can usually be overcome by refining the specification of the sourcing, but while numerous attempts were made to prescribe a suitable refinement, none were successful in generating a grid that could be "closed" with POSTGRID.

Fortunately, the Heldenmesh software suite became available at about the time that the VGRID/POSTGRID process was found to be unworkable. Heldenmesh is a software suite developed by the Helden Aerospace Corporation for the rapid generation of unstructured grids about complex configurations; it was designed to be compatible with, but a robust extension of, the TetrUSS grid generation tools. The specific code used in this effort is Heldenmesh version 1.1. Heldenmesh progresses in the same manner as VGRID, generating a surface triangulation, using an advancing layer method to grow the boundary layer, then an advancing front method to generate the inviscid grid. Moreover, Heldenmesh is compatible with the GridTool input deck, so the effort that was put into patching and sourcing the tower was not wasted. One of the distinctions between the two codes is that Heldenmesh checks the validity and quality of each tetrahedral cell as it is generated, while ensuring that no holes are left in the grid as the procedure extends out into the far field. Heldenmesh proved to be quite robust and was able to generate the tower grid with the same input deck for which the VGRID/POSTGRID process failed.

Upon generating the unstructured mesh for the tower, the next step was to generate an appropriate near-body grid for use with Kestrel. Once again Carpenter was used, first to convert the tetrahedral grid into a mixed element grid, then to trim the tower. At the time, the only guidance available for picking the trim distance was from transition simulations conducted for various SLS launch vehicle configurations in free air. A modest investigation of the most appropriate trim distance was conducted in working the transition problem, so that the SLS-28005 configuration was trimmed to 24 inches. However, when issues arose for the SLS-27005 configuration, it was found that a trim distance of 36 inches was more appropriate. Since both grids were to be used in the liftoff simulations, it was decided to stick to the more conservative value and trim the launch tower to 36 inches. The final trimmed tower grid contained roughly 102 million cells, of which 55 million were prisms, 39 million were tetrahedrons, and 8 million were pyramids.

While the unstructured grids for the vehicle and tower were quite large at 121 and 102 million cells, respectively, they are quite small relative to the Cartesian grid. In the Liftoff simulations completed to date, the Cartesian mesh varied from about 1.5 billion points for cases at  $h/L = 0.915$  to about 3 billion points for cases at  $h/L = 0.0$ .



**Fig. 7 Distribution of line (yellow) and volume (pink) sources on the Block 1B launch tower.**



#### D. A Glitch in the Topology

The intent in setting up the liftoff problem was to mimic the wind tunnel installation. Since the launch tower sits on the wind tunnel floor, the problem was set up with the Xmax Cartesian grid boundary coincident with the bottom of the launch tower platform. Viscous gridding was then utilized on the five surfaces of the platform and the outer boundary patch that encloses the bottom of the platform was made coincident with the wind tunnel floor and gridded as an inviscid surface.

That set up did not work, as the the runs would bomb after the initial Cartesian grid had been defined. So, the problem was simplified by building a grid for the platform alone and running that simple single near-body grid on the tunnel floor. Those runs did not work either and eventually it was determined that the problem involved the MeshManager trying to grid the inside of the platform. Since a correction for the issue proved to be much less straight forward than one might expect, a fix to the topology was devised that involved raising the launch platform off the tunnel floor. This caused a bit of an increase in the size of the tower grid, and care had to be taken to specify the location of the Xmax boundary of the Cartesian grid. However, the fix worked.

The situation for the fix is depicted in Figure 8, where cross sections of the three grids at the  $y=0$  plane are shown. The unstructured grids for the vehicle and tower are shown in blue and purple with the Cartesian grid shown in black. This case is for  $h/L = 0$ , as indicated by the presence of the Tail Service Masts on the platform and the stubby vehicle sting projecting down into the hole in the platform floor. Note that the unstructured grid on the bottom of the platform is terminated prior to reaching the side of the Cartesian mesh that is coincident with the floor of the wind tunnel.




Fig. 8 Cross sections of the vehicle and tower grids imbedded in the Cartesian grid at the start of a run.

### IV. Kestrel Results for SLS-28005 Liftoff Simulations

#### A. Run Matrix and Resource Requirements

Liftoff simulations of the SLS-28005 configuration were conducted at wind tunnel conditions, specifically, at Mach 0.183. The corresponding Reynolds number is 615,000 based on the nominal core diameter. The matrix of flow conditions at which simulations were run, for which the independent variables are the vehicle liftoff ratio  $h/L$  and roll angle  $\phi_P$ , is shown in Table 1.

The run matrix was developed in conjunction with the end-users of the liftoff line loads databases, namely, the Loads and GNC groups at MSFC. The Loads group was most concerned about flow conditions with the vehicle still on the launch pad, for which  $h/L = 0.0$ . The GNC group was primarily concerned about the potential of a tower impact with winds from the South pushing the vehicle into the tower, for which  $\phi_P = 180^\circ$ ; they also showed interest in tower interference effects at  $h/L = 0.915$ , at which point the vehicle is nearly clear of the tower. An additional run at  $h/L = 0.2$  and  $\phi_P = 0^\circ$  was performed during debugging of the CFD process, prior to establishing the run matrix.

All the CFD simulations were run on the large Pleiades Linux cluster of the NASA Advanced Supercomputing

**Table 1 Run matrix for the liftoff CFD simulations with the end-user identified.**

$h/L$	$\phi_P$ , deg	End-User	$h/L$	$\phi_P$ , deg	End-User	$h/L$	$\phi_P$ , deg	End-User
0.0	0	Loads	0.2	0	Devel	0.915	60	GNC
0.0	60	Loads	0.2	180	GNC	0.915	180	GNC
0.0	90	Loads	0.5	180	GNC	0.915	270	GNC
0.0	180	Loads/GNC	0.8	180	GNC	0.915	300	GNC
0.0	270	Loads						

(NAS) facility located at the Ames Research Center. The simulations were run almost exclusively on Haswell processors, using 1200 CPUs for transition simulations and 2400 CPUs for liftoff simulations and taking anywhere from five to ten days. As expected, the shortest runs occurred for roll angles of 180° and the longest runs took place at roll angles of 0°.

All of the simulations were conducted with version 7.1.2 of Kestrel, rather than version 8, which was available on DoD computer clusters at the time. While the faster speed of version 8 would have made a huge improvement on throughput of the simulations, difficulties arose in implementing the NBOB methodology on the Pleiades cluster, even though cases that were restricted to using the single body approach worked just fine. It is evident that the DoD and NAS clusters are not entirely compatible.

## B. Kestrel Input File

The simulations were run with the SA+DDES model invoked for both the near-body and off-body grids. Inviscid fluxes were discretized with the HLLE++ method, convective fluxes were handled with the Van Leer flux splitting scheme, and the 2nd-order temporal discretization was handled with Gauss-Seidel. The most significant differences between inputs for the two Kestrel flow solvers are that the spatial discretization was 2nd order and the advective temporal damping was set to 0.01 for KCFD, whereas for SAMAir the spatial discretization was 3rd order and the advective temporal damping was set to 0.02. The dimensional time step was set at  $\Delta t = 0.0005$  sec. with 3 subiterations per time step and a maximum of 32 sweeps per subiteration. The refinement variable for the solution adaption procedure was the vorticity magnitude with the threshold value set to 100. Refinements were begun after the first 500 time steps and continued every 250 time steps throughout the rest of the simulation.

Although time dependent solutions never reach a strict definition of convergence, they do reach a point where time histories of the integrated forces and moments flatten out and the subsequent trends can be characterized as having reached a statistically steady state, as the values oscillate about some pseudomean. Once a statistically steady state was reached, time-averaging procedures within the code were invoked and the solutions were run out additional time steps, thereby defining the time-averaging window over which time-averaged values of the force and moment coefficients and surface and volume flow properties were obtained. The nominal procedure was to run the simulations out 10000 time steps, then run the time-averaging procedure another 4000 time steps. However, the time-averaging procedure was occasionally initiated earlier than that if the solution was a bit early in reaching a statistically steady state, or restarted at 14000 time steps and run another 4000 times steps with time-averaging if the solution was a bit slow to settle out.

While there are numerous Kestrel input parameters to be set up for LOT simulations, the LaRC SLS CFD team was fortunate enough to have the initial input deck set up by the lead authors of Ref. [7], namely, Scott Morton and David McDaniels. Consequently, over the last two years of developing best practice procedures for the liftoff problem, most of the effort in exploring options for the Kestrel input deck have been in regard to refinements of the unstructured near-body grid for the vehicle and the resolution and extent of the solution adaption region for the wake. In order to illustrate some of the finer points of the set up for the AMR procedure, a cross section of a Kestrel grid generated at the end of a liftoff simulation is shown in Figure 9. The direction of flow is from the left to the right of the page, with the tower beams on the left and the vehicle on the right. So the flow is passing through the tower, then impinging upon the vehicle. The contours are of vorticity magnitude with increasing vorticity going from blue to green, yellow, red, and white. The unstructured grids about the vehicle and the beams of the tower are not shown in order to clearly delineate the extent of the Cartesian grid in the vicinity of the solid bodies.

The extent to which the Cartesian grid is refined by the solution adaption procedure is designated in the input deck by specifying the 8 corners of a rectangular box; here, the values are set at 4 times the radius of the core, in the plus and minus directions for both the y and z axes, and large enough in the x direction to enclose the vehicle. Clearly, most of




Fig. 9 A cross section of the Cartesian grid demonstrating results from the Kestrel AMR implementation.

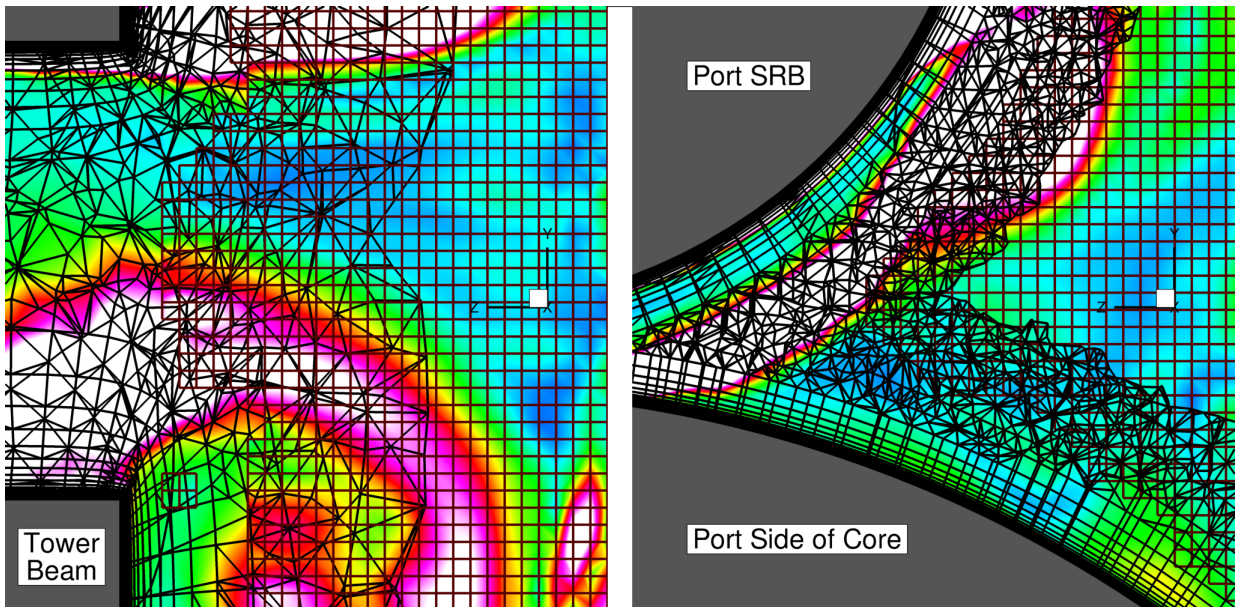


Fig. 10 Cross sections of the Cartesian grid around a tower beam (left) and around the gap between the port SRB and the core (right).



the region within the tower was not refined with the solution adaption procedure. Instead, the refinement region begins just prior to the last row of beams on the side of the tower closest to the vehicle; the grid around those beams is refined and preserves the small eddy content as it flows toward the vehicle. The beauty of the solution adaption procedure is that the Cartesian grid only gets refined in regions where gradients of the specified flow field variable are large. For this particular case, the procedure was based on the magnitude of the vorticity with a threshold set at 100; the threshold value corresponds to the yellow contours in the figure. In areas where the vorticity is low, particularly behind the core and SRBs but also in a thin strip emanating from one of the gaps between the tower beams, the grid is not refined nearly as much as it could have been.


A closer look at the extent to which the unstructured grids and Cartesian mesh overlap is shown in Figure 10, where zoomed in cross sections of a region on the port side of the tower and the gap between the port SRB and core are depicted. The position of the overlap region is dictated by the trim distance used in trimming down the unstructured grid to a specified distance from solid walls, as well as the size of the grid cells at the outer edge of the unstructured grid. In the vehicle grid the trim distance is set at 14% of the nominal radius of the core, whereas the trim distance for the tower grid is 50% larger. For the gap between the port SRB and the core, it is evident that the size of the cells at the outer boundary of the unstructured grid are smaller than those of the Cartesian mesh. Conversely, in the region of the tower beams, the size of the cells at the outer boundary of the unstructured grid are larger than those of the Cartesian mesh. The reason for the discrepancy is that in generating the grid for the vehicle, volume sources were used to ensure that outside the boundary layer, a constant cell size was kept for a large distance away from the solid walls; this ensures that the grid resolution will be sufficient to meet the DDES requirements. For the tower, no such restriction is invoked, allowing the grid size outside the boundary layer to expand in the usual manner for RANS simulations.

### C. Kestrel Results at $h/L = 0.0$ and $\phi_P = 0^\circ$

The very first case to be simulated with the Kestrel setup for the liftoff problem was for the SLS-28005 configuration at  $h/L = 0.2$ ,  $\phi_P = 0^\circ$ . This is one of the most severe conditions in the run matrix, as the flow passes through the tower prior to rushing around the vehicle. The situation is illuminated in Figure 11, where pressure coefficient contours are provided over a panoramic view of the vehicle and launch tower. At this condition, the LAS shroud has barely cleared the launch tower. The wind is from the front and left of the figure, blowing around and through the launch tower, then into and around the vehicle. The purple colors indicate regions of high pressure with decreasing pressure values for red then orange, yellow, and green. It is no surprise to see that the highest pressures occur on the windward surface of the tower beams and equipment rooms. The effect on the vehicle is a reduction in pressure on the windward surface that is not very uniformly distributed.


The pressure coefficient distributions on the windward surface of the vehicle for both the liftoff case with the tower and the transition case without the tower are shown in Figure 12. The view for the liftoff case includes a translucent outline of the tower structure, as well as black lines to indicate the stations at which rows of pressure taps were installed on the wind tunnel model. The most positive pressures show up in the contours as white with decreasing pressure values moving from red through orange, then yellow, to green. The view through the tower is oriented to see the bottom of the vehicle, which puts the port SRB on the top of the two figures and the starboard SRB on the bottom.

The most striking difference between the two results is how the stagnation pressure is suppressed for the liftoff case, as the contours along the center of the core and the insides of the SRBs are red and purple rather than bright white as in the transition case. As a consequence, the acceleration of the flow around the sides of the core and the SRBs is reduced, so the pressure doesn't drop as much, as evidenced by the red contours on the outsides of the core and SRBs for the liftoff case, as opposed to yellow and green contours for the transition case. Note that the pressures are higher on the outside of the starboard SRB than on the outside of the port SRB, reflecting the influence of the elevator shaft which



**Fig. 11 Pressure coefficient contours on the vehicle and tower at  $h/L = 0.2$ ,  $\phi_P = 0^\circ$ .**

is in line with the port side of the vehicle. Amazingly enough, some fairly small features of the tower interference effect are preserved in the simulation, like the white patches located on the center of the core that line up with gaps between the tower beams. Meanwhile, the influence of the tower dwindles rather quickly above the tower, as the pressure distributions on the LAS are quite similar for the two cases.



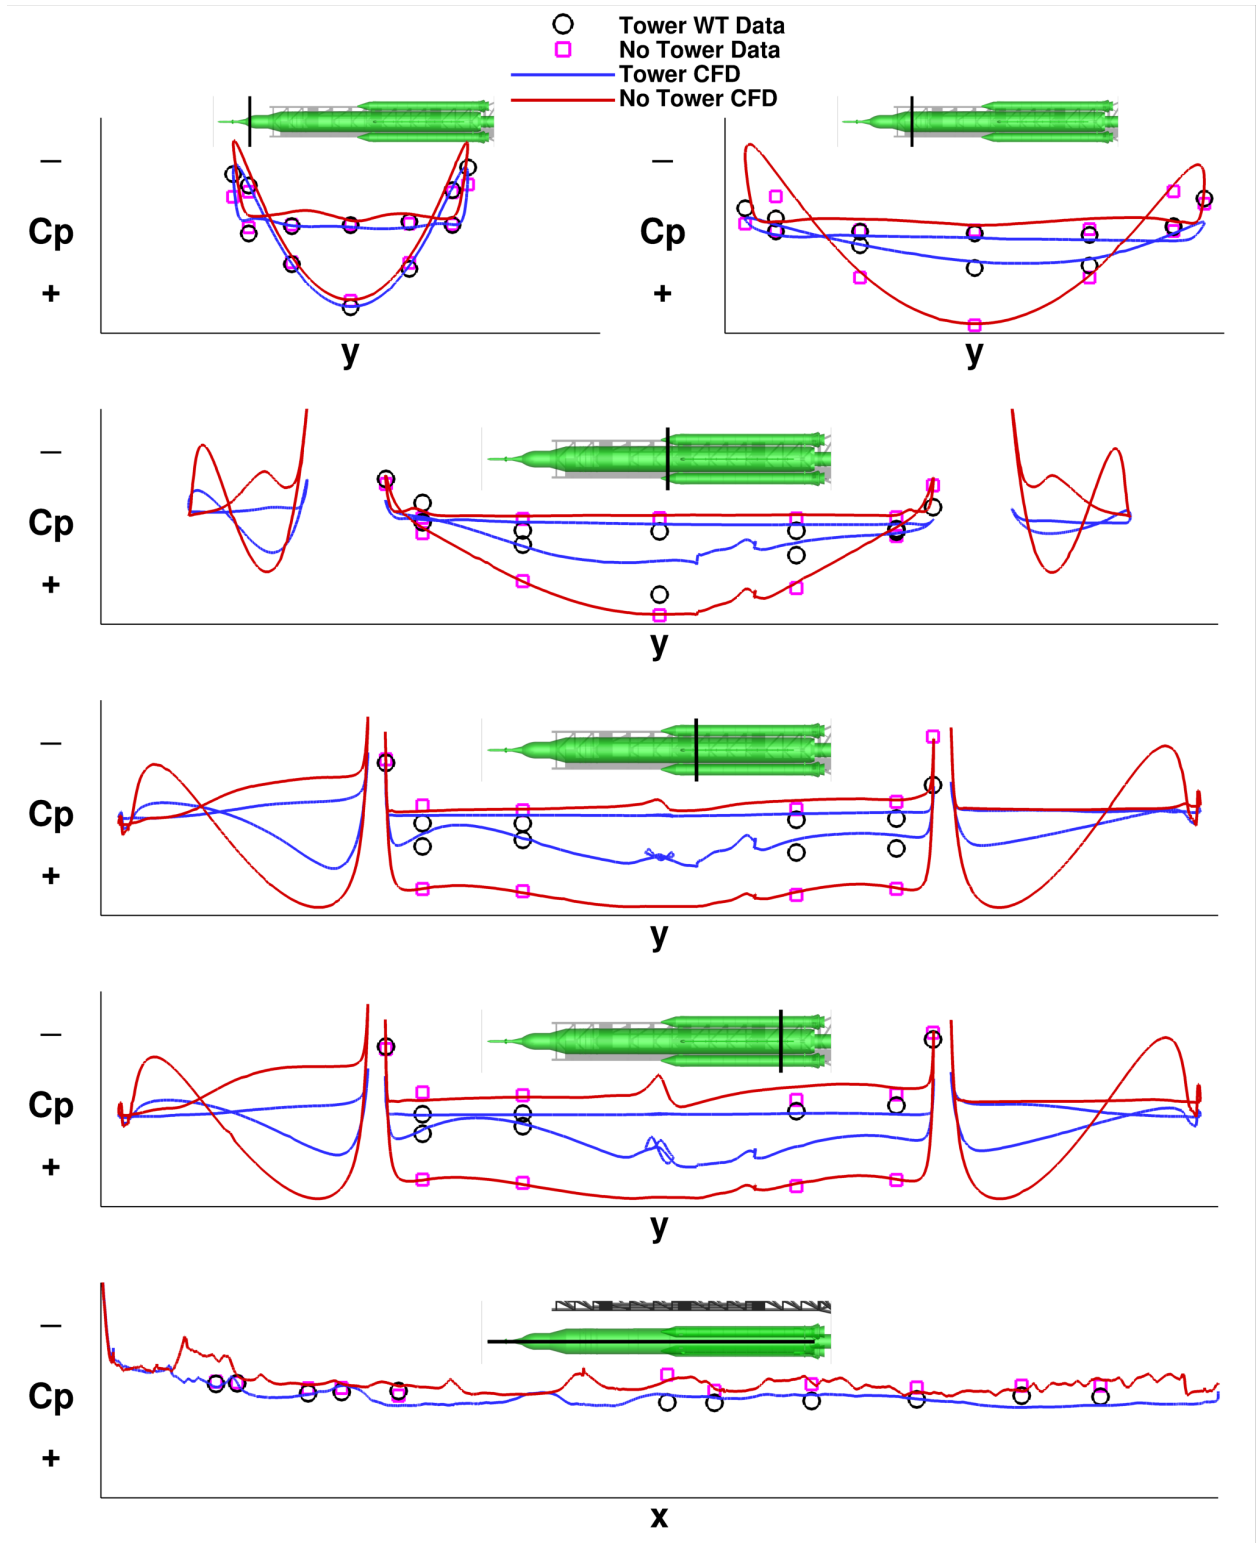
**Fig. 12 Pressure coefficient distributions on the windward surface for the  $h/L = 0.2$ ,  $\phi_P = 0^\circ$  liftoff case (top) and the  $\alpha_P = 90.0^\circ$ ,  $\phi_P = 0^\circ$  transition case (bottom).**

The other set of wind tunnel data that can be used to assess the veracity of the CFD results is the static pressure data, which are taken at 62 pressure ports located on the core; unfortunately, there are no pressure ports on the SRBs, since the installation of load cells on the brackets precluded the ability to instrument the components. In addition to the rows of pressure ports at the stations shown in Figure 12, an additional longitudinal line of pressure ports was placed along the core at an angle of  $30^\circ$  from the y axis; for this particular view, the line of ports would appear as running down the starboard side of the core, on the far side of the vehicle.

Wind tunnel and time-averaged CFD pressure coefficient data for all of the pressure port stations are shown in Figure 13. In the orientation of this plot, pressures for the port SRB are on the left, core pressure values are in the middle, and values for the starboard SRB are on the right. In each plot, the inset icon of the vehicle indicates the station at which the pressures are taken. The aerodynamic convention for plotting the pressure coefficient is employed, where the vertical axis direction is reversed from the usual orientation; due to the sensitive nature of the data, magnitudes of the coefficients are not included. For the liftoff data with the tower effects, the wind tunnel data are shown with black open circles and the CFD data are shown with blue lines. For the transition data with no tower, the wind tunnel data are shown with purple open squares and the CFD data are shown with red lines. As discussed earlier, due to the installation of load cells on the SRB fore and aft brackets, there is no wind tunnel pressure data available for the SRBs.

The biggest trend from the pressure plots is how much less positive the pressure coefficient values on the windward side of the vehicle are for the liftoff case in comparison to the transition case (no tower), except at the forward most station. It almost goes without notice that the leeward pressure coefficient values for the liftoff case are also consistently less negative than the values for the transition case. It is evident that the CFD results are very well correlated with the wind tunnel values for both the transition case and the liftoff case. Although the pressure coefficient values for the wind tunnel and CFD data are better correlated for the transition results than for the liftoff cases, the correlation between the wind tunnel and CFD data for the liftoff case is quite remarkable seeing as how the influence of the tower had to be conveyed two core diameters downstream before it interacted with the vehicle. In fact, there is only one wind tunnel data point that lies well away from the CFD value; it is located at  $y=0$  on station 3. Nonetheless, the trends between the wind tunnel and CFD values, and between the liftoff and transition data sets, suggest that the CFD results provide a valid increment to the lineloads database. Of course there is no experimental pressure data on the SRBs to confirm the veracity of the CFD results, but the trends look good.






**Fig. 13** Pressure coefficient values from Test 633 and the Kestrel DDES+SA simulation for liftoff at  $h/L = 0.2$ ,  $\phi_P = 0^\circ$  and transition at  $\alpha_P = 90^\circ$ ,  $\phi_P = 0^\circ$ , at all five stations and along a longitudinal cut at  $\phi = 30^\circ$ .

An interesting feature seen on the SRBs is that on the starboard side, the trend in pressure values on the leeward side is very flat, whereas on the port SRB, the trend is very flat for the liftoff case but exhibits much more negative pressure coefficient values for the transition case. As discussed in some detail by Ratnayake, et al.[1], this peculiar behavior is driven by the Coanda effect exhibited in the regions just downstream of the gaps between the SRBs and the core. Considering just the transition case on the port side, the flow through the gap attaches to the port SRB, which causes more negative pressure coefficient values behind the SRB; on the starboard side, the flow through the gap wrapped around the core, which causes more negative pressure coefficient values on the starboard side of the core but yields flat, less negative values on the starboard SRB. For the liftoff case, the leeward pressure values are very flat for both SRBs, suggesting that the flows through both gaps attach to the core; however, the blockage from the tower appears to have suppressed the Coanda effect on both sides.

The influence of the tower on the line loads can be gleaned from Figure 14, which shows the side force (top) and normal force (middle) line loads on the core; the outline of the core and SRB OMLs are shown at the bottom. The liftoff data are shown with blue lines and the transition data are shown with red lines. Once again, the magnitudes of the coefficients are not included due to the sensitive nature of the data. However, as an aid in understanding the data trends, a horizontal line is provided at the zero value of the coefficient and the ranges of the vertical axes for the normal and side force are the same. Given the orientation of the vehicle for this case, the normal force is the dominant coefficient, at least for the transition case. For the liftoff case, the magnitudes of the normal force values are much reduced from those for the transition case. The other notable feature in the plot is how much more the line load trend for the liftoff case undulates along the length of the vehicle, reflecting how the flow gets disrupted as it passes through the tower framework.



**Fig. 14 Normal and side force line loads on the stack for liftoff at  $h/L = 0.2$ ,  $\phi_P = 0^\circ$  and transition.**

Side force and normal force line loads on the SRBs are shown in Figure 15; the ranges of the scales are kept the same as those in the core line load plots to provide some perspective. Based on the previous discussion about the influence of the Coanda effect and that the pressures on the leeward surface of the port SRB were more negative for the transition case (no tower), it is no surprise that the transition case normal force line load values for the port SRB are of greater magnitude than for the starboard SRB. Moreover, the liftoff normal force line load values for both SRBs are of smaller magnitude than those for the transition line loads. Conversely, it is the starboard SRB for the transition case that exhibits the largest magnitude side force line load values. Looking back at both the surface and cross sectional pressure plots, it is evident that the flow accelerates much more around the outside of the starboard SRB for the transition case than it does for any of the other cases, yielding an imbalance in the integrated side force. The side force line loads for the port SRB in both the transition and liftoff cases, and the line load for the starboard SRB in the liftoff case, are near zero for most of the length of the vehicle.




Fig. 15 Normal and side force line loads on the SRBs for liftoff at  $h/L = 0.2$ ,  $\phi_P = 0^\circ$  and transition.

#### D. Kestrel results for a roll sweep at $\phi_P = 180^\circ$

Due to the GNC community's particular interest in the cases with the wind blowing out of the South, an  $h/L$  sweep at  $\phi_P = 180^\circ$  provides the best opportunity to assess the trends from the liftoff CFD against the trends in the force and moment database. This is also one of the more difficult orientations to assess, since tower interference effects on the vehicle are primarily due to blockage effects, which can be quite subtle when the blockage is two core diameters downstream of the vehicle.

Plots of the normal force coefficient for a roll sweep at  $\phi_P = 180^\circ$  are shown for the vehicle and each of the three components in Figure 16. The database values are from the force and moment database and the CFD values are from Kestrel with DDES+SA. The top plot shows total vehicle coefficient values, while the three plots at the bottom show the values for the port SRB, core, and starboard SRB, respectively; the same range is used in all four plots to provide some perspective as to the differences between the CFD and database values. The convention in this plot is reversed from that for the line plots of the pressure coefficient; here the force and moment database values are plotted with black lines while the CFD values are plotted with red circles. Both the experimental and CFD normal force coefficient values located at  $h/L = 1.5$  are the values from the transition simulation, as that is the vehicle liftoff ratio at which the liftoff regime ends and the transition region begins.

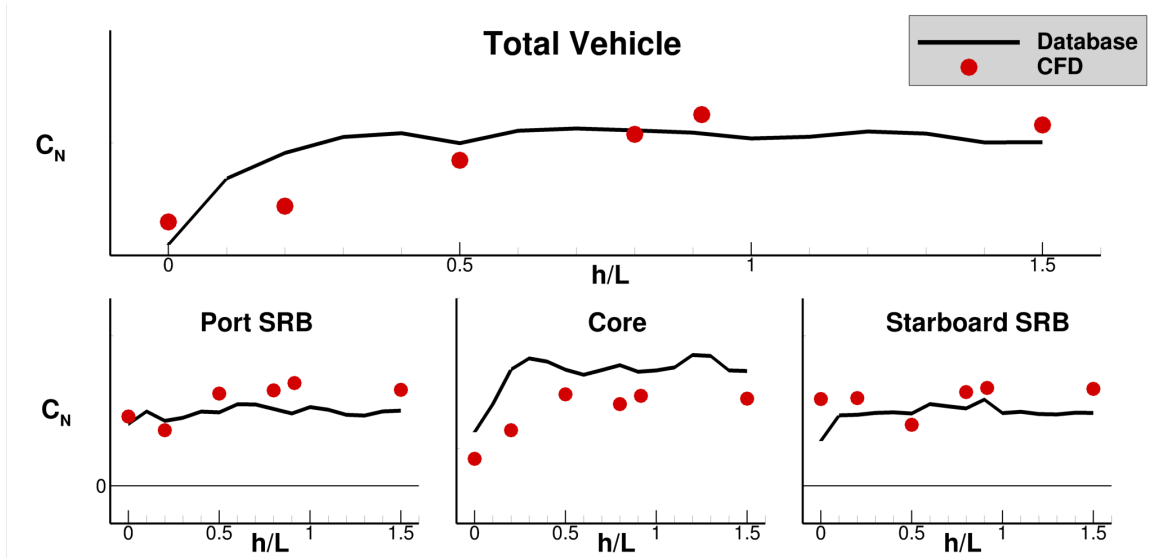



Fig. 16 Normal force coefficients for the vehicle (top) and port SRB, core, and starboard SRB (bottom) from the force and moment database and DDES+SA results for a pitch sweep at  $\phi_P = 180^\circ$ .

For the total vehicle, starting at  $h/L = 1.5$  and moving toward decreasing vehicle liftoff ratio values, the trend of the force and moment database normal force is quite flat until about  $h/L = 0.3$ , at which point it begins decreasing at an ever faster rate. The CFD actually does a good job of predicting that trend except at  $h/L = 0.2$ , where the CFD underpredicts the normal force value by a significant amount. Looking at the component values, it is apparent that the largest discrepancies between the CFD and database values arise on the core and just about all of the discrepancy at  $h/L = 0.2$  is due to the underprediction of the normal force on the core. The most significant discrepancy on the other components is for the starboard SRB at  $h/L = 0.0$ , where the database value drops off noticeably but the CFD value does not. It is conjectured that the discrepancy arises due to the CFD not capturing the interference effects of the platform and Tail Service Masts. Despite some noticeable discrepancies between the database and CFD values, the trends in the database values are fairly well captured by the CFD.

As shown in Figure 17, the situation for the side force coefficient is not nearly as sanguine as that for the normal force; once again the plots use the same range of values as used in the normal force plot. For the liftoff cases, the CFD vastly overpredicts the total vehicle side force at  $h/L = 0.0$  and  $0.2$  and underpredicts it at  $h/L = 0.5$ , with a smaller underprediction at  $h/L = 0.915$ . Moreover, even at the points where the CFD and force and moment database values for the total vehicle are in good agreement, it is only because the values on the port SRB are underpredicted by about the same amount as they are overpredicted on the starboard SRB. It is evident that the CFD did not predict the force and moment database trend at all.



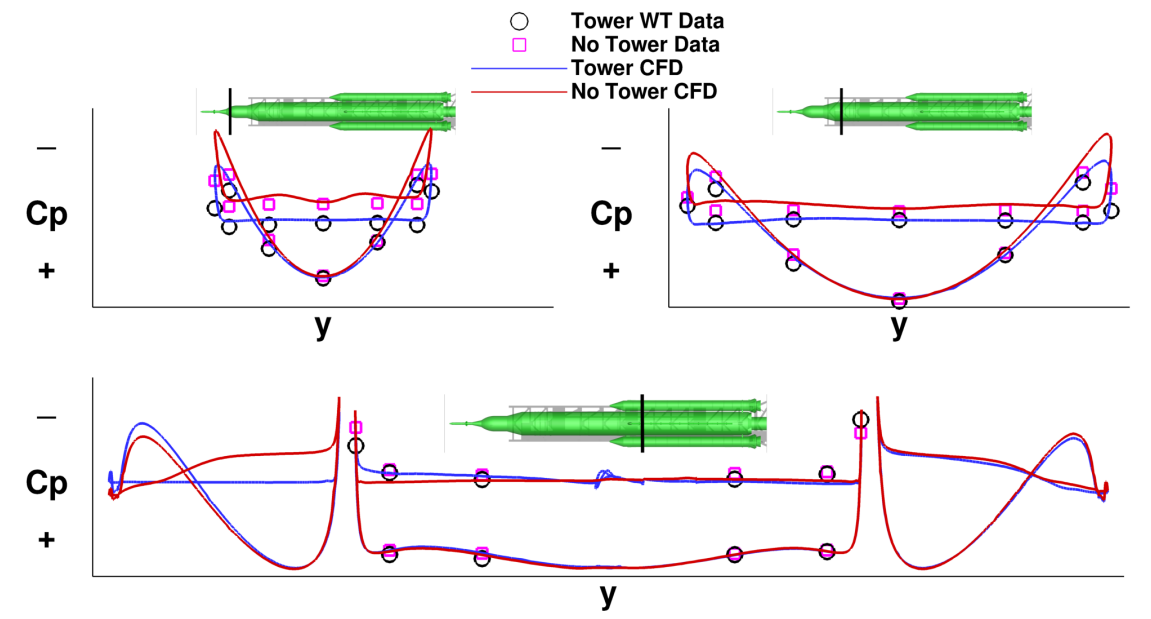
**Fig. 17** Side force coefficients for the vehicle (top) and port SRB, core, and starboard SRB (bottom) from the force and moment database and DDES+SA results for a pitch sweep at  $\phi_P = 180^\circ$ .

As a start toward explaining the situation, consider the component values of the side force coefficient for the transition case ( $h/L = 1.5$ ), for both the force and moment database and the CFD. Whereas the force and moment database and CFD total vehicle side force values are in very good agreement, the CFD values for all three components are close to zero, but the database value for the SRBs are quite large in magnitude, with the port SRB being positive and the starboard SRB being negative. It appears that the transition wind tunnel test captured the Coanda mode where the flow through both gaps attaches to the sides of the core, whereas the CFD has captured the mode where the flow through both gaps attaches to the SRBs.

Next, consider the results for the starboard SRB. It is evident in the liftoff force and moment database values that the wind tunnel test captured a Coanda mode where the flow attaches to the starboard side of the core at most conditions, yielding large negative values of the side force on the starboard SRB that are of the same magnitude as in the transition case. However, the mode flipped at  $h/L = 0.915$  to attach to the SRB, yielding a near zero side force value that is in good agreement with the CFD. For the CFD, at most conditions, the simulations caught the Coanda mode where the flow attaches to the starboard SRB, yielding a nonzero side force. However, the mode flipped for the CFD at  $h/L = 0.5$  to attaching to the starboard side of the core, putting the side force value in very good agreement with the database value. The only discrepancy left to explain in the starboard SRB database side force values is at  $h/L = 0.0$ , where the database

value plunges to an even more negative value; a similar interference effect is not captured in the CFD. It is conjectured that this trend arises due to the interference of the Tail Service Masts, which do not have much of an influence on the vehicle at higher  $h/L$  values.

The situation can be further clarified by considering the wind tunnel and time-averaged CFD pressure coefficient data for the  $h/L = 0.2$  case; the two most forward pressure port stations and the fourth station are shown in Figure 18. Note that this is the case for which the CFD predicted vehicle normal force is underpredicted by the largest magnitude. As with the pressure comparisons for  $h/L = 0.2$ ,  $\phi_P = 0^\circ$ , the correlation between the wind tunnel data and CFD values is quite good. From the results at the first two stations, it is seen that the tower has a significant blockage effect on the top of the vehicle, which effectively increases the back pressure in the wake, even though the MPCV has just cleared the tower. From the results for the fourth station (which is representative of the results at the third and fifth stations), it is seen that the tower has little effect on the core pressures but does have an impact on which Coanda mode is expressed. Although comparison of the wind tunnel pressure values to the CFD values provides little indication as to why the normal force is underpredicted in the CFD, comparison of the CFD liftoff and transition pressures on the SRBs elucidates some of the side force trends. On the starboard SRB, the pressures for the transition and liftoff results are nearly identical, indicating that the CFD caught the same Coanda mode in both simulations; consequently, in comparing the side force coefficient values in Figure 17, the CFD values at  $h/L = 0.2$  and  $1.5$  are quite comparable. Conversely, on the port SRB the liftoff and transition CFD caught different Coanda modes, the result being that the CFD side force coefficient value at  $h/L = 0.2$  is quite a bit larger than the CFD value at  $h/L = 1.5$ , but it is comparable to the force and moment database value at  $h/L = 1.5$ .




**Fig. 18** Pressure coefficient values from Test 633 and the Kestrel DDES+SA simulation for liftoff at  $h/L = 0.2$ ,  $\phi_P = 180^\circ$  and transition at  $\alpha_P = 90^\circ$ ,  $\phi_P = 180^\circ$ .

The influence of the tower on the line loads at  $h/L = 0.2$  is depicted in Figure 19, where the side force (top) and normal force (middle) line loads on the core are shown. The normal force for the liftoff case appears to be suppressed all along the vehicle length, but a little more so over the front of the vehicle. The comparison of the database and CFD integrated normal force values suggests that the liftoff normal force line load should be a bit less suppressed than it is.

The discussion above leads one to the awkward conclusion that even though the CFD results do not agree all that well with some of the wind tunnel measurements, the CFD does in fact capture valid states for the liftoff and transition problems. Since there are multiple modes of the Coanda effect, there is a big impact on side force that is dependent on which particular mode is expressed, and the phenomena is extremely sensitive to geometry, flow conditions, and, in numerical simulations, the mesh resolution and flow solver parameters, it is likely that wind tunnel measurements and CFD computed values will not always correlate well, even when both the wind tunnel test and the CFD computations have been properly implemented. The conclusion in this investigation is that the CFD solutions have captured viable





**Fig. 19 Lineloads on the stack for liftoff at  $h/L = 0.2$ ,  $\phi_P = 180^\circ$  and transition at  $\alpha_P = 90^\circ$ ,  $\phi_P = 180^\circ$ .**

states of the flow, just not necessarily the same states captured in the experiments. As detailed in the companion paper by Wignall [9], such discrepancies play havoc with the methods employed by database developers.

## V. Conclusions

The Kestrel flow solver is very effective in simulations of the liftoff and transition problems for SLS launch vehicle configurations. Over the many simulations run, the Kestrel results usually correlated very well with the wind tunnel data and captured many of the trends exhibited in the force and moment database; this is particularly true in regard to the missile axis normal force, which is the most dominant force acting on the vehicle. The Kestrel results also captured some trends that were not exhibited in the force and moment database; this is particularly true in regard to the missile axis side force, for which trends in both the database and the CFD results are somewhat erratic. Whether or not the mercurial behavior of the side force trends can properly be attributed to the vagaries of the Coanda effect, then reconciled in a database, is a topic for future investigations.

## References

- [1] Ratnayake, N. A., Krist, S. E., Ghaffari, F., and Deere, K. A., "Computational Fluid Dynamics Methods Used in the Development of the Space Launch System Liftoff and Transition Lineloads Databases," *AIAA AVIATION Forum*, 2019.
- [2] Strelets, M., "Detached Eddy Simulation of Massively Separated Flows," *AIAA 2001-0879*, 2001.
- [3] Spalart, P., Deck, S., Shur, M., Squires, K., Strelets, M. K., and Tavin, A., "A new version of detached-eddy simulation, resistant to ambiguous grid densities," *Theoretical Computational Fluid Dynamics*, Vol. 20, No. 181, 2006.
- [4] 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.
- [5] Krist, S. E., and Ghaffari, F., "Detached Eddy Simulation Results for a Space Launch System Configuration at Liftoff Conditions and Comparison with Experiment," *AIAA 2015-0776*, 2015.
- [6] Biedron, R. T., Carlson, J.-R., Derlaga, J. M., Gnoffo, P. A., Hammond, D. P., Jones, W. T., Kleb, B., Lee-Rausch, E. M., Nielsen, E. J., Park, M. A., Rumsey, C. L., Thomas, J. L., and Wood, W. A., "FUN3D Manual: 13.2," *NASA/TM-2017-219661*, 2017.
- [7] 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.
- [8] NASA Langley Research Center, "TetrUSS: Tetrahedral Unstructured Software System, Publications," , 2006. URL <https://tetruss.larc.nasa.gov/tetruss-related-publications/>.
- [9] Wignall, T., "Physics-Based Reduced Order Model Development for Launch Vehicle Aerodynamics," *AIAA AVIATION Forum*, 2019.