

Output-Based Adaptive Meshing Applied to Space Launch System Booster Separation Analysis

Derek J. Dalle*

Science and Technology Corp., Moffett Field, CA 94035

Stuart E. Rogers†

NASA Ames Research Center, Moffett Field, CA 94035

This paper presents details of Computational Fluid Dynamic (CFD) simulations of the Space Launch System during solid-rocket booster separation using the Cart3D inviscid code with comparisons to Overflow viscous CFD results and a wind tunnel test performed at NASA Langley Research Center’s Unitary Plan Wind Tunnel. The Space Launch System (SLS) launch vehicle includes two solid-rocket boosters that burn out before the primary core stage and thus must be discarded during the ascent trajectory. The main challenges for creating an aerodynamic database for this separation event are the large number of basis variables (including orientation of the core, relative position and orientation of the boosters, and rocket thrust levels) and the complex flow caused by the booster separation motors. The solid-rocket boosters are modified from their form when used with the Space Shuttle Launch Vehicle, which has a rich flight history. However, the differences between the SLS core and the Space Shuttle External Tank result in the boosters separating with much narrower clearances, and so reducing aerodynamic uncertainty is necessary to clear the integrated system for flight. This paper discusses an approach that has been developed to analyze about 6000 wind tunnel simulations and 5000 flight vehicle simulations using Cart3D in adaptive-meshing mode. In addition, a discussion is presented of Overflow viscous CFD runs used for uncertainty quantification. Finally, the article presents lessons learned and improvements that will be implemented in future separation databases.

I. Introduction

The Space Launch System (SLS) is a new exploration-class launch vehicle currently in development by NASA and a team of contractors and will be the first exploration-class launch vehicle since Saturn V. To achieve its mission, SLS is using Solid Rocket Boosters (SRBs) that are only slightly modified hardware from the Space Shuttle Launch Vehicle. These SRBs expend their fuel long before the SLS core first stage, and thus there is a separation event that must occur during a phase of flight where aerodynamic loads are nontrivial. A drawing of SLS with the two SRBs attached on the left and right sides is shown in Fig. 1.

Because gravity is not acting in a way that pulls the expended boosters away from the core, and aerodynamic loads also do not act in a way to substantially encourage separation, small rockets called Booster Separation Motors (BSMs) are used to actively push the boosters away from the core. There are 16 of these BSMs—four forward BSMs and four aft BSMs on each booster. The forward BSMs fire in a partially forward (upstream-facing) direction, which results in complex plumes and interactions of those plumes with the high-speed freestream flow. In addition, the low dynamic pressure during the separation event combined with the relatively high-thrust separation motors means that the plume affects a very large volume of flow—on the scale of the SLS vehicle itself.

The SRBs are very similar to Space Shuttle hardware, which have a rich history of successful flights. However, separating from the SLS core turns out to be more challenging than separating from the Space Shuttle External Tank. The reason for this is simple: whereas the External Tank essentially ended at the

*Research Scientist/Engineer, Applied Modeling & Simulation Branch, Member AIAA, derek.j.dalle@nasa.gov

†Aerospace Eng., Applied Modeling & Simulation Branch, Associate Fellow AIAA, stuart.e.rogers@nasa.gov




Figure 1. Artist rendition of the SLS initial configuration

point where the SRBs were attached, SLS has a substantial amount of hardware, including its main engines, directly aft of the SRB attachment point. As a result, clearances between the aft portion of the SRB and the aft portion of the SLS core stage can be small. To have sufficient confidence that no recontact will occur, reducing aerodynamic uncertainty is essential.

The complex behavior caused primarily by the forward-facing Booster Separation Motors, but also affected by the close proximity (at least immediately after separation), means that reducing this uncertainty is a nontrivial problem. In addition, there are many degrees of freedom in this separation event, which means that the aerodynamic database requires thousands of individual flow solutions. In the most liberal run matrix, there could be independent variables for the relative positions of the boosters (6 variables), separate attitudes for each booster (6 variables), attitude of the core (2 variables), thrust levels for each rocket boundary condition (4 main engines, 2 booster engines, and 16 BSMs), altitude, and Mach number. This would give a database with 38 lookup variables, so some simplifying assumptions must be made. Even after simplifications, the run matrix used to support the actual flight database for this work contains 2730 individual simulations, and thus efficiency and robustness are essential.

To complete this number of simulations at a reasonable expense (both CPU hours and calendar time), an inviscid solver with adaptive meshing was selected as the primary analysis tool. Cart3D [1,2] is a solver developed at NASA Ames with exactly these capabilities and uses an embedded-boundary Cartesian grid approach that makes volume mesh generation automatic. The solver also has a history of use with powered rocket boundary conditions, including nozzles that are small relative to the scale of the vehicle, from its use in support of the Orion program [3,4]. The inviscid assumption is tentatively acceptable due to the high supersonic Mach number (greater than 4.0) at separation and low dynamic pressure. However, complex flow and large boundary layers may prove this assumption to be less than ideal, and viscous Overflow [5] simulations have been used to evaluate this assumption and provide an increment to the flight database uncertainty quantification that accounts for viscous effects.

Figures 2 and 3 demonstrate how the flow solution can vary quite dramatically as a function of the mesh used to compute it, demonstrating the degree of sensitivity of the solution to discretization error at these conditions. Figure 2 shows the static pressure on the surface and two planar slices of the flow for a Cart3D solution with the boosters displaced axially (Δx), laterally away from the core (Δy), and down in the crew's reference frame ($\Delta z < 0$). The displacements are on the scale of a few feet, and thus the boosters are still in close proximity to the core. The orientation of the core is zero angle of attack ($\alpha = 0$) and zero sideslip ($\beta = 0$), and the boosters are symmetrically deflected 1 degree down ($\Delta\alpha = -1^\circ$) and 0 degrees




Figure 2. Static pressure for grid created with manually specified grid resolution




Figure 3. Static pressure using 7 cycles of adaptive meshing

outward ($\Delta\beta = 0$). The manually-specified grid resolution in Fig. 2 is based on experience from SLS Design and Analysis Cycle 2 (DAC-2) for the booster separation. Notably, the flow solution in Fig. 2 does not appear to be converged, and the solution is not symmetric to the level one would expect. Figure 3 shows a result calculated on the same DAC-2 geometry using an automated initial mesh and 7 adaptation cycles. The output function used for the mesh adaptation was the sum of axial forces on the boosters. Using this adaptive procedure leads to a solution that is visibly smoother and more symmetric, but this does not provide evidence that it is a more accurate solution. In fact, the following sections will indicate that the solution in Fig. 2 is closer to reality.

The two figures use the same color scale, so the deeper red in Fig. 3 indicates that the further-aft strong shock results in higher pressures on the core, but over less area. The adaptive meshing capability is helpful in reducing sensitivity to the initial mesh, but for complex cases such as this one with a high degree of sensitivity to discretization error, there remains some dependence on the initial mesh unless the user is willing to run conservatively on very large meshes.

To support the SLS project, a new aerodynamic database was produced for the DAC-3 geometry, which is the configuration for Critical Design Review (CDR). Comparisons to viscous Overflow results are shown, which provides insights that are hard to predict for a database with complex flow and many degrees of freedom. The procedures for this work also took advantage of lessons learned to support first-stage separation analysis that supported Ares I [6, 7] and especially Ares V [8, 9], which was a similar design to the SLS. In addition, a wind tunnel study performed at NASA Langley Research Center's Unitary Plan Wind Tunnel in September and October of 2014 was used to provide validation for the computational aerodynamic database. Because of the range of geometric scales, from booster separation motors with throat planes several inches in diameter to the overall vehicle length of more than 300 ft, completing a wind tunnel test to support the aerodynamic database was judged to be prohibitively expensive. As a result, the Computational Fluid Dynamics (CFD) database generated in this work is a high-priority analysis for the mission success of SLS.

II. CFD Inputs

This section discusses procedures that were used to create the Cart3D simulations of the wind tunnel, Cart3D simulations for the CDR separation database, and Overflow simulations used to complete the uncertainty quantification. Differences from previous run procedures, particularly those used for DAC-2 separation analysis, are highlighted and discussed. An additional aspect to this work was cases that were rerun with slightly perturbed inputs when solutions with inadequate quality were found. The quality assurance procedure is discussed in Sec. IV.C, but the slightly different inputs used for rerun cases are discussed here. Finally, potential areas of inputs that could be used to improve both performance and efficiency are also discussed in this section.

A. Input Sensitivity Study

Cart3D has several tools for creating volume meshes. For many flow geometries, the differences can be considered quite subtle because the overall solution will not vary qualitatively if even a roughly adequate initial mesh is selected. For this SLS configuration, as indicated in Figs. 2 and 3, the solution is highly dependent on the volume mesh generation parameters, and so some understanding of this dependence is essential.

The procedure for creating volume meshes in Cart3D is to provide a triangulated surface, an initial flow domain size, and some guidance on how to create the volume mesh. The program `cubes` is then used to create the initial Cartesian mesh, including cut cells for cubes that intersect the triangulated surface. If Cart3D is run with adaptive meshing, the initial mesh produced by `cubes` is simply a starting point. After the flow solution is computed on that initial mesh, Cart3D also calculates an adjoint (based on a user-defined scalar output function) and an error map, which is used to refine some fraction of the cells (cubes or cut cells) in the volume mesh. The user also specifies how many such adaptation cycles to compute




Figure 4. Surface and *xy*-plane cuts of static pressure for several different meshing approaches using DAC-2 geometry

and the mesh growth factor for each adaptation cycle. Standard procedure is to use a quite coarse initial mesh and then use mesh adaptation until mesh convergence is observed and error estimates are below a target level. For difficult configurations such as the SLS booster separation, that procedure would require an impractically large mesh. As a secondary practice, the approach here is to use as many adaptation cycles as practical, which may be limited either by mesh size or failure to converge the adjoint.

Based on early indications and experience with the DAC-2 database, the decision was made to quantify the effect of the initial mesh generated by `cubes` on the solution. A comparison along these lines in Fig. 4 shows that the dependence is considerable. Figure 4a shows a slice through the same solution as Fig. 2 but using a manually specified mesh. The forward BSMs are visible on the two booster nosecones. Figure 4b shows a second solution without using adaptive meshing but using a Cart3D program called `autoInputs`, which generates meshes with resolution that depends on distance from the surface and local surface curvature. The solution in Fig. 4b looks smoother than that of Fig. 4a, but the local high-pressure

regions are missing.

Figures 4c and 4d both employ adaptive meshing and initial meshes created using `autoInputs`, but with differing levels of initial refinement. The solution in Fig. 4c has a shock that is further aft and much stronger, and that of Fig. 4d has a different feature that is described in more detail in Sec. IV. The resulting surface pressures, and thus forces and moments on the components, are notably different. Approximately 20 different initial meshes and mesh growth procedures were investigated for this condition, and the initial conclusion was that there are two distinct flow-solution categories to which the adaptive meshing can converge. Subsequent analysis after running the full set of cases for the DAC-3 database showed that it is more accurate to say that there are actually three such categories, and these categories are discussed in more detail in Sec. IV. Data from the aforementioned wind tunnel test at NASA Langley’s Unitary Plan Wind Tunnel also showed two qualitative solution categories, and there are some conditions for which the wind tunnel results show either bifurcation or switching between the two categories. One of these wind tunnel categories corresponds relatively well to the forward-shock solutions shown in Fig. 4, but the other wind tunnel solution category has a strong shock all the way forward to the tip of the launch abort tower (which is not even pictured in Fig. 4). Using all of the sources of available data, it eventually became clear that the solutions with shocks farther forward are closer to the likely physical reality. However, confidence in that statement came after running the Cart3D database runs, and so there are data points in the database that show the aft shock location.

B. Cart3D Database Inputs

The input procedure used for the bulk of the CDR SLS booster separation database is an adaptive procedure with approximately 20 million cells on the final mesh. Analysis from the SLS guidance, navigation, & control team indicated that the most important force and moment coefficients affecting booster clearance are booster side force and booster yawing moment. As a result, the objective function used for the database runs emphasizes the booster side force, but the other seven force coefficients are also included with lower weights. Previous experience operating Cart3D in adaptive mode on complex configurations that are unlikely to have true mesh convergence led to the decision to ensure that all nine force coefficients are included in the objective function. Finally, a small weight for one component of the BSM thrust was also added to help ensure convergence of BSM thrust. The full objective function used to drive mesh adaptation for the Cart3D database runs is

$$\begin{aligned}
 J = & C_{Y,LSRB} + C_{Y,RSRB} + 0.1C_{N,LSRB} + 0.1C_{N,RSRB} + 0.005C_{A,LSRB} + 0.005C_{A,RSRB} \\
 & + 0.001C_{A,CORE} + 0.002C_{Y,CORE} + 0.0001C_{N,CORE} \\
 & + 0.002C_{Y,LBSMs} + 0.001C_{N,LBSMs} + 0.002C_{Y,RBSMs} + 0.002C_{N,RBSMs}
 \end{aligned} \tag{1}$$

where *LBSMs* refer to the left BSMs, and the other force coefficients do not include the thrust conditions, nozzles, or base regions. The asymmetry between the weights for $C_{N,LBSMs}$ and $C_{N,RBSMs}$ is an error that had a minimal effect on the database. Using an adjoint-based mesh adaptation procedure, the Cart3D software then attempts to refine the mesh in the regions that have the most impact on this objective function.

With this objective function, Cart3D was run for five adaptive mesh cycles with 300 iterations per cycle and 500 iterations on the final mesh. This adaptive portion of the run procedure used first-order spatial discretization, which was judged to be necessary to obtain adequately robust adjoint convergence. Once the final mesh was obtained, an additional 3000 iterations using second-order inputs were run, resulting in 5000 total iterations. To reduce the wall time to obtain these solutions, the second-order iterations were run using the newly-available MPI version of Cart3D on 240 cores. Simulations were run on the *Pleiades* supercomputer at the NASA Advanced Supercomputing facility at the NASA Ames Research Center, and both Ivy Bridge nodes with 20 cores and Haswell nodes with 24 cores were utilized. The database coefficients were calculated by averaging over at least 2000 iterations. When a visual inspection of the iterative histories indicated that the mean was not converged over the last 2000 iterations, additional sets of 3000 second-order

iterations were added until at least 2000 iterations with similar appearance were obtained. In addition, the standard deviation of the force and moment coefficients over the averaging window were also measured and recorded, and that data was used in the subsequent uncertainty analysis.

Finally, some aspects of the initial volume mesh were modified from the defaults to enable better robustness and quality for the CDR database. The first issue is the wide range of scales for the SLS vehicle. In order to start with a fairly coarse grid and enable the adaptation to create an efficient mesh, considerations must be made to ensure adequate resolution of the BSM nozzles on the initial mesh. Without doing so, an entire BSM may lie within a single volume cell on the initial mesh, leading to slow mesh convergence or even inconsistent results if the error indicator does not detect the presence of the thrust boundary condition. To counter that, Cart3D's *XLev* option was used to make five additional levels of refinement on the BSM boundary condition planes and nozzle interiors. Similarly, two levels of additional refinement were added to the core engine and booster main engine surfaces.

The sensitivity study described in the preceding subsection also indicated that increasing the initial resolution near the forward portion of the core is an advantageous procedure, so a refinement box was placed around the Orion service module. This box ensures at least 9 levels of refinement in a box that contains the service module with a padding region of 50 inches in all dimensions. Cart3D defines refinement boxes with *BBoxes*, which specify a minimum absolute refinement level, which is different from an *XLev* specification. An *XLev* specification is applied after the first pass at the volume mesh and commands the volume mesh to apply a number of additional refinements on cut cells. Finally, two levels of *BBox* are also added to improve the solution quality for the main engine plumes. The reason for these last two refinement boxes is that the objective function is not directly affected by main engine thrust, but some action needs to be taken to ensure adequate resolution of the plumes to determine when they might begin to have some effect on the aerodynamic forces.

C. Rerun Procedures for Cart3D Database

Following the full quality assurance procedure described in Sec. IV.C, some initial Cart3D solutions were determined to be unsatisfactory when using the input procedures described in the preceding subsection. The most prominent cause leading to such a determination was a large asymmetry between the force coefficients on the left and right booster. Although the boosters are not perfectly symmetric, the force coefficients should be roughly symmetric, and some simulations showed very different results on the two boosters. The strategy used to improve the quality for such cases was simple; provide an alternate set of inputs with only minor changes so that better symmetry might be attained without using a highly changed procedure that might lead to obvious jumps in the database.

Toward this goal, the minor change selected for the rerun cases was to alter the objective function used for mesh adaptation. The original objective function, shown in Eq. (1), was changed to

$$\begin{aligned}
 J = & -C_{Y,LSRB} + C_{Y,RSRB} + 0.1C_{N,LSRB} + 0.1C_{N,RSRB} + 0.005C_{A,LSRB} + 0.005C_{A,RSRB} \\
 & + 0.001C_{A,CORE} + 0.002C_{Y,CORE} + 0.0001C_{N,CORE} \\
 & + 0.0005C_{Y,LBSMs} + 0.0005C_{N,LBSMs} + 0.0005C_{Y,RBSMs} + 0.0005C_{N,RBSMs}
 \end{aligned} \quad (2)$$

The main difference here is the replacement of $C_{Y,LSRB}$ with $-C_{Y,LSRB}$. The impact of such a change is not fully understood, but results indicate that there is an appreciable effect on the mesh, and the resulting meshes may be slightly more symmetric.

In addition to this change in the objective function, the rerun cases also used more iterations at each stage of the run, with 7000 total iterations instead of the 5000 used in the initial runs. The effect of the change in iterations appeared to be very minor, but it led to fewer necessary restarts to get a 2000-iteration window with stable history.

D. Inputs for Cart3D Wind Tunnel Comparisons

The inputs used for the wind tunnel comparison runs are similar to those used for the Cart3D-based database, but several differences were necessary. It should be noted that the wind tunnel comparisons were actually run before the database cases, so the lessons learned from the wind tunnel simulations influenced the inputs for the database runs rather than the other way around. The first difference is simply one of scale; the wind tunnel model is a 0.9%-scale model, and so parameters tied to physical lengths, such as reference lengths and time steps used for time-accurate simulations, differ correspondingly. Another difference was that more iterations were run for a nominal case. When moving from the wind tunnel comparison runs to the flight geometry used for the database, it was noticed that iterative variation was decreased for the flight geometry, and so the number of iterations was decreased from 10,000 to 5000.

The objective function used to drive mesh adaptation was

$$\begin{aligned} J = & C_{Y,LSRB} + C_{Y,RSRB} + 0.5C_{N,LSRB} + 0.5C_{N,RSRB} + 0.5C_{A,LSRB} + 0.5C_{A,RSRB} \\ & + 0.05C_{A,CORE} + 0.01C_{Y,CORE} + 0.01C_{N,CORE} \\ & + 0.001C_{Y,LBSMs} + 0.001C_{N,LBSMs} + 0.001C_{Y,RBSMs} + 0.001C_{N,RBSMs} \end{aligned} \quad (3)$$

Finally, there was no refinement boxes for the core engines because the engines were not part of the wind tunnel model.

Another aspect of the wind tunnel comparison was a significant number of time-accurate Cart3D runs. Two notable aspects of this analysis were that this was one of the first uses of the time-accurate Cart3D MPI version of the code, which was not publicly available at the time, and the fact that it indicated a high degree of consistency with the steady-state iterative average. Running in time-accurate mode was roughly ten times as expensive, and the conclusion that the same results could be obtained using a less physical but also less expensive procedure was a positive outcome.

E. Inputs for Overflow Comparisons

The Overflow code was run to simulate the booster-separation flow using the same geometry used by the Cart3D simulations of the flight vehicle. The Overset grid system was a derivative of that used in recent work to assess the protuberance loads during ascent [10]. For the current simulations, this grid system was modified by removing all of the smaller protuberances—a process made relatively easy by the overset grid approach. Additional volume grids were added to the BSM, core-stage engine, and SRB nozzles to enable the plume simulations. Also added were off-body refinement grids to resolve all 22 plumes, the largest of which were added the volume where the forward BSM plumes interact with the freestream flow. The resulting grid system consists of 324 million grid points and 641 zones.

The grids that were used in each of the nozzles covered the nozzle interior from the nozzle exit face to a location upstream of the throat where the nozzle area is four times the throat area. This provided a coplanar inlet grid face where boundary conditions were applied to match the desired thrust. Application of the boundary conditions required the generation of separate solution files. The boundary-condition flow properties were derived for a given total pressure and total temperature using isentropic relations and the given area ratio to ensure sonic flow at the throat. A series of calibration runs were used to derive the relation between the prescribed total pressure and the resulting thrust, which gave the ability to specify any desired thrust level at each nozzle. The fluid properties for each type of nozzle used different values for the ratio of specific heats (γ) and the molecular weight in an effort to simulate the gas properties of each plume. Overflow was run using a variable- γ and multiple-gas-species approach. This was different from the Cart3D simulations which used air as the working fluid due to the limitation of constant γ imposed by the software.

The Overflow simulations were computed with the standard release version 2.2k of the Overflow code, using the Message-Passing Interface (MPI) parallel version. The computations were performed on the SGI Altix system known as *Pleiades* at the NASA Advanced Supercomputing facility. The computations were

all performed on the Haswell nodes of this machine. All of the runs used 34 Haswell nodes, which equates to 816 cores.

The simulations were run using a third-order HLLC (Harten-Lax-van Leer-Contact) upwind differencing scheme, together with an SSOR (symmetric successive overrelaxation) line-relaxation implicit scheme. All viscous terms, including cross terms, were enabled in the code. The SST (shear stress transport) turbulence model [11,12] was used, and the flow was assumed to be turbulent everywhere. For the initial start-up phase and for steady-state iterations, first-order time integration was used together with locally-scaled time-step sizes based on a constant Courant-Friedrichs-Lewy (CFL) number. Since these flow-fields never converged to a steady state, the simulations were completed with time-accurate integration using a dual time-stepping approach in which the outer-level integration was second-order in time.

III. Run Matrix and Database

As discussed in the introduction, there are a large number of potential degrees of freedom that could be incorporated into a three-body separation database. Because the clearance for booster separation had become a critical item for the entire SLS program, some additional limitations were placed on this database in order to enable a relatively quick turnaround for both the aerodynamics (i.e., the work discussed here) and the guidance, navigation, & control simulations. Mainly, these limitations came in the form of requirements to have only minor changes in the structure of the aerodynamic database.

As a result, some degrees of freedom that may affect the forces and moments on the core and boosters could not be incorporated into the database. Specifically, freestream Mach number and dynamic pressure were held constant throughout the database (despite the fact that the conditions at separation naturally have some dispersion associated mostly with the weight of the vehicle and the weather conditions). The database used for DAC-2 also did not include a lookup variable for BSM thrust level, and so the database created for this work has BSM thrust level tied directly to axial displacement. The thrust level from each booster's primary engine also has a wide range of possible values during separation. Like the BSM thrust, the booster thrust level was tied to booster axial displacement for this work (axial displacement is used as a proxy variable for time since separation), but this variable is judged to be less likely to significantly affect separation aerodynamics. Finally, the booster roll angle (relative to the core) was held at zero, although small changes in this angle may change where the BSM plumes impinge on the core and thus significantly affect the forces and moments on the core.

Table 1. Input variables for DAC-3/CDR separation database

Variable	Description
<i>dx</i>	Booster axial translation [ft]
<i>dy</i>	Booster lateral translation, away from core [ft]
<i>dz</i>	Booster translation in +z-direction [ft]
<i>dpsi</i>	Booster yaw angle relative to core body axes [deg]
<i>dtheta</i>	Booster pitch angle relative to core body axes [deg]
<i>alpha</i>	Angle of attack of the core [deg]
<i>beta</i>	Angle of sideslip of the core [deg]
<i>CT_BSM</i>	Booster separation motor thrust coefficient

The specific input variables used for the database being discussed in this work are summarized in Table 1. As with any high-dimensional database, running a fully rectangular matrix for this case is prohibitively expensive. For this work, the approach was taken to use a block-rectangular matrix; at each value of axial booster displacement (dx), there is a rectangular matrix. The smaller matrices at each value of dx select from a fixed list of values (often called break points), and so the overall matrix is analogous to a block diagonal matrix. A set of 4–10 break points for each variable were selected for the global run matrix, but only 2–6 of these break points were considered at each value of dx . Axial displacements for the matrix varied from 0 ft

(attached position) to 10 ft, and the other break points were selected based on dispersed separation simulations that were run using the DAC-2 aerodynamic database. The run matrix fully covers the set of dispersed separation trajectories using the DAC-2 aerodynamics. Since the booster aerodynamic uncertainties became much smaller as a result of the CFD work presently discussed, this matrix is conservative.

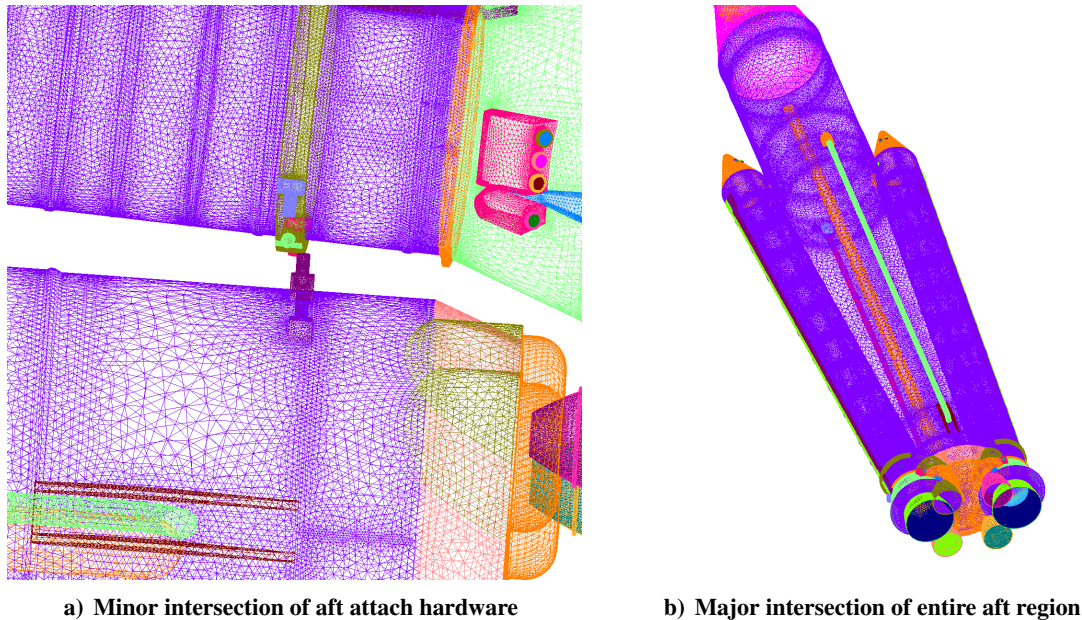


Figure 5. Examples of geometries with booster intersections

Multiplying the number of break points in each variable results in a matrix with roughly 4000 cases. However, the constraints on the database infrastructure, which required a fixed set of break points for the global matrix, resulted in many of these cases having intersections between the boosters and the core. Some of these intersections were relatively minor, usually involving only the aft attach hardware such as the example shown in Fig. 5a, and these were run after applying Cart3D's *intersect* tool. Other points in the initial run matrix included major intersections that have no realistic physical interpretation, such as the example shown in Fig. 5b, and no attempt was made to run Cart3D for such cases. After filtering for these major intersections, the database was reduced to 2732 data points.

The cases run for wind tunnel comparisons used a slightly different set of input variables. One significant difference is that the two boosters have slightly asymmetric orientations, facilitating more input variables in the run matrix file but not resulting in an increased number of cases. The booster roll angle relative to the core was not intentionally varied during the wind tunnel test, but its value was slightly nonzero, and so roll angle was also an input variable for Cart3D comparison runs. Another significant change is that the BSM thrust levels are controlled by a stagnation pressure measured in the stings that supported the SRB models, and this resulted in four different BSM thrust levels (one for each pod of four motors) controlled by three input variables (left and right sting pressures and a common stagnation temperature). Finally, the wind tunnel experiment utilized stings attached to the base, and so there were no core engines or booster main engines in the simulation.

IV. Results

The results for this effort include a diverse set of CFD simulations. Three main sets of results are used to support the booster separation database: Cart3D runs of the flight geometry used for the actual data book, Overflow runs of the flight geometry used for code-to-code comparison, and Cart3D runs of a wind tunnel geometry for comparison between CFD and reality. The wind tunnel results come from a dedicated separation test performed at the highest-speed test section of the NASA Langley Research Center Unitary

Plan Wind Tunnel. With these three sets, it is possible to estimate the uncertainty on the Cart3D-based data book using a root-sum-squared approach. An obvious weakness of this approach is the lack of direct comparison between Overflow and the wind tunnel test. Constraints on time and resources meant that it was not possible to create the necessary surface and volume grids to run Overflow on the wind tunnel geometry, which is not quite a geometric scaling of the flight geometry. However, progress on Overflow comparisons to the wind tunnel test is being made as resources allow. In addition, all three of these data sets contain solutions in which the Booster Separation Motors (BSMs) are firing and solutions in which they are not.

The most important data set, namely the Cart3D flight vehicle simulations, contain results that can qualitatively be grouped into three categories. The categories are shock position 3, which consists of solutions with a plume-induced pressure rise on the surface near the aft end of the conical interstage region; shock position 2 with a shock at the aft edge of the service module; and shock position 1 with plume-induced pressure rises somewhere on the service module. Figure 6 shows examples of these three solutions. Based on the results of the wind tunnel test and Overflow viscous CFD, there is good evidence that all of the conditions in the database should physically have shock position 1. The Cart3D results, however, have a solutions in each category. Fortunately, the category into which the particular solution falls does not greatly affect the forces and moments on the boosters, which is the most critical aspect of the separation aerodynamic database. The forces and moments on the core, especially the pitching moment and normal force, are strongly affected. However, the core is much heavier and has an active control system, so although the uncertainties on the core increased as a result of this new database, the decreases in uncertainty on the booster aerodynamics have led to a net improvement in separation performance.

The pressure rise that is often denoted a “shock” in this work is actually a much more complicated structure. The structure is induced by the partially-forward-facing BSMs and incorporates aspects of stagnation or opposed flow, oblique and/or normal shocks, and in some cases flow separation. The color scale in Fig. 6 is mostly saturated for shock position 3, which prevents a full examination of the flow topology, but it is the easiest case to describe. Such solutions basically consist of two opposed supersonic flows, namely the freestream flow and the exhaust of the forward BSMs. Both opposed flows consist of a streamline that separates the two flows and a normal shock on both sides. The strong pressure gradient visible in the shock 3 flow sample images is actually the normal shock of the freestream flow; the separation streamline is downstream (i.e., further in the $+x$ direction) and not directly visible.

The flow topology for solutions in the shock position 1 and 2 categories are similar to each other but quite different from that of the shock position 3 category. The additional aspect that separates this topology from the one previously described is a region of separated flow adjacent to the structure. In these cases (which appears to be the more physical solution), exhaust from the forward BSMs makes its way forward into a roughly conical region of subsonic and circulating flow. The freestream flow then forms an oblique shock based on the angle of this cone. Because the samples in Fig. 6 show pressure, these separated regions cannot be directly seen. This roughly conical geometry in the solution applies to a relatively confined region in the flow; the edge of this domain can be seen as a change in slope of the pressure increases in Fig. 6. Outside of this conical region dominated by separated flow, the flow topology is similar to that of the shock position 3 description. A related flow configuration resulting from supersonic retropropulsion also shows a range of possible solutions, and this has been noted in inviscid simulations, viscous simulations, and experimental data [13].

For the present configuration, solutions with shock positions that are further forward have smaller pressure increases. One indirect consequence for this discussion is that the pressure scale is less saturated for the shock position 1 images in Fig. 6. That enables further insight into the flow topology by focusing on Fig. 6e. At first the repeated increases and decreases in pressure along a single constant- x line are difficult to understand—to the point where a reasonable first impression is that it must be some non-physical artifact of CFD. It turns out that the complicated topology is basically supersonic opposed flow. First, consider a line starting in the upper left and going downstream (from left to right). The first pressure increase is either an oblique shock or normal shock that is relatively easy to understand. Then, continuing through the white

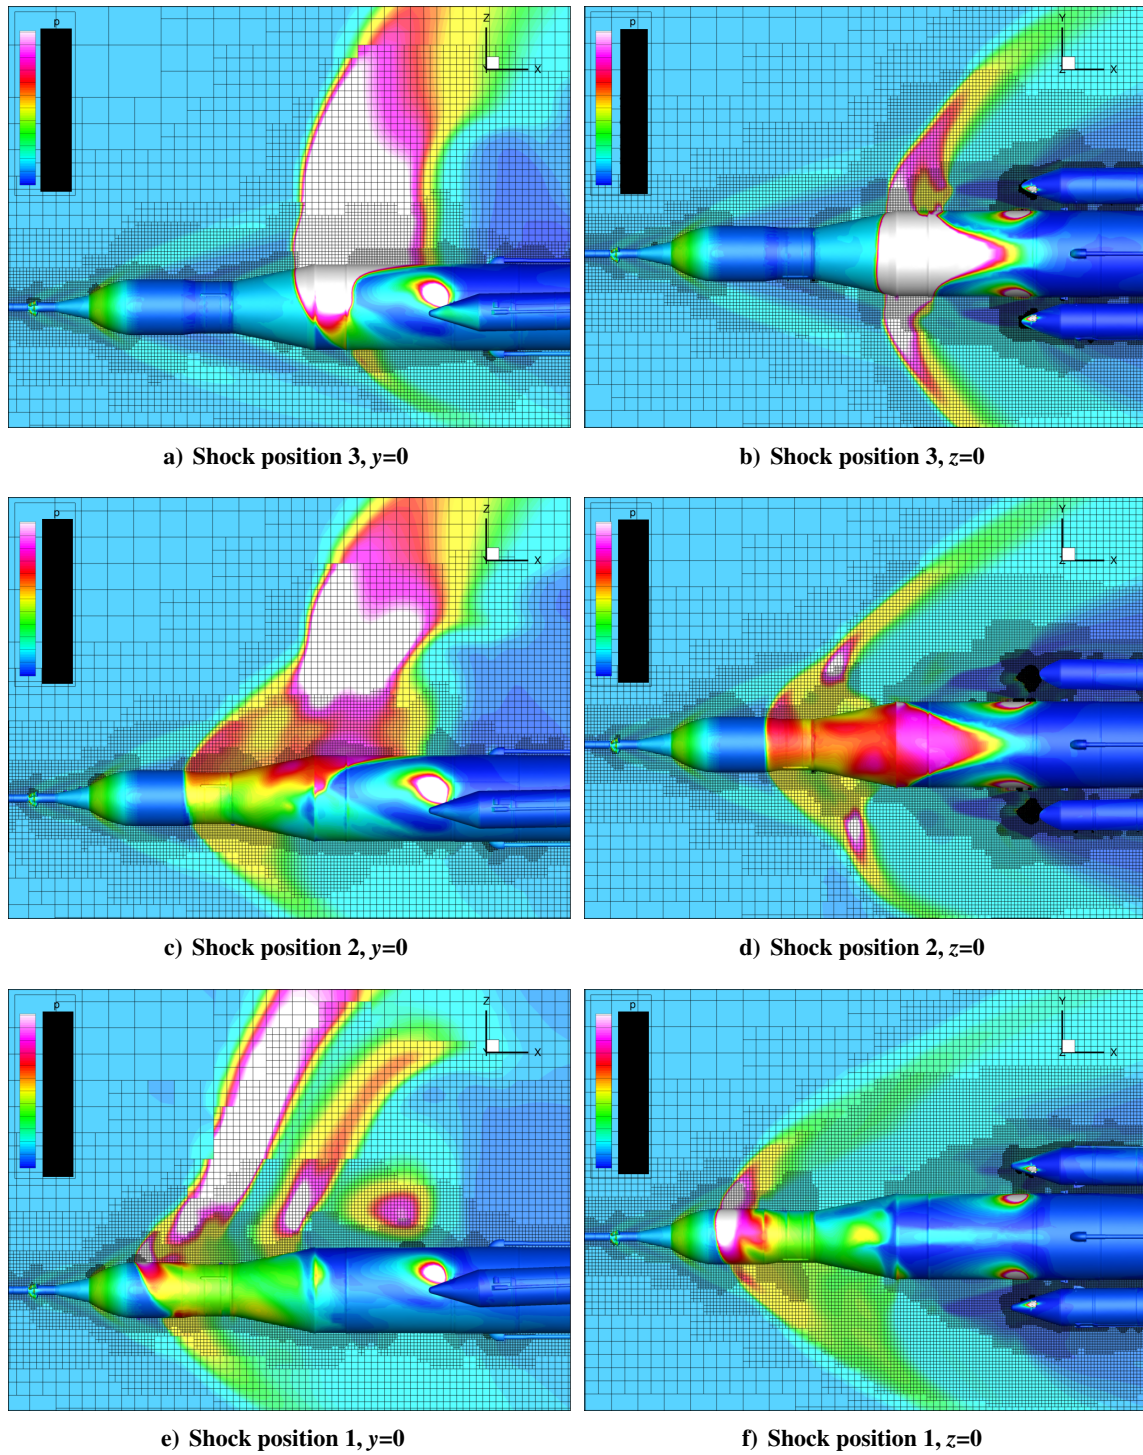


Figure 6. Three main Cart3D shock position categories with BSMs firing

region, there is a somewhat surprising decrease in pressure that is marked by a shift from white to roughly green color in Fig. 6e. This pressure decrease comes from the freestream flow turning away from the rocket and returning to supersonic flow. This greenish region is basically a subsonic stagnation flow. Staying with Fig. 6e but starting from the forward BSMs near the nose of the booster, and tracing upward and to the left, one sees several features. The first is the white oval region on the core surface, which is the result of BSM




Figure 7. Mach number slices and surface pressure coefficient with sample streamlines

exhaust flow directly impinging on the core body. Next, one can see an oval pinkish region just above the core body, which is simply the exhaust plume from the BSMs of the left booster (pictured) interacting with the plumes from the right booster. Continuing up and to the left, there is a sharp pressure increase; this is the supersonic stagnation flow normal shock on the exhaust flow side. Following this direction further, there is a pressure decrease, and this is the flow accelerating as it approaches the separation streamline, much like the pressure decrease described above for the freestream side. The interesting aspect of the flow that is not at first obvious is that the streamline separating the freestream flow from the BSM exhaust flow is roughly in the middle of the low-pressure stripe between the two high-pressure regions. Figure 7 shows sample streamlines from an Overflow solution, and the streamlines are seeded from either the freestream flow or a BSM. The separation streamline is visible in the upper portion of the figures as the region where the other streamlines coalesce.

Another aspect of the solutions is that even cases with shock positions 1 and 2 tend to have a further-aft shock in earlier iterations. As a consequence, much of the mesh adaptation occurs on a solution with strong gradients in locations that differ from the gradient locations on the final mesh. This can be seen especially clearly in Fig. 6e, which has higher resolution just downstream of the shock than the resolution across the shock itself. At first glance at residual, force & moment, and mesh convergence, the Cart3D simulations with the shock position 3 appeared to be the most converged solutions. Since increasing from first-order inputs to second-order inputs is effectively a uniform refinement, and there were many cases in which the shift from first order to second order caused a shift in the shock location, it may be the case that all mesh-converged Cart3D solutions would fall into the shock position 1 category. It may be the case that there are multiple valid inviscid solutions while the additional consideration of viscosity biases the flow toward one of these solutions. Certain wind tunnel run conditions demonstrated two shock positions depending on how the boosters were moved into place (the two positions shock position 1 and another topology with the shock even further forward); thus the potential for multiple solutions cannot be completely ignored. The wind tunnel cases exhibiting this behavior were not within the expected range of booster orientations expected during separation, but such multiple-solution cases were close enough to the expected conditions to be interesting. Alternatively, there is also evidence that given a sufficiently fine mesh, Cart3D will also select shock position 1 for all of the conditions in the BSM-on database. Although it is very difficult to determine the correct solution with this combination of information, the consistency of Overflow results and wind tunnel data within the nominal booster orientation envelope have led us to the operating assumption




Figure 8. Surface geometries for Cart3D analysis of wind tunnel test

that the relevant BSM-on configurations should all be in the category of shock position 1. For future runs, we have determined inputs that lessen the probability that Cart3D will select aft shock positions in the early adaptation cycles.

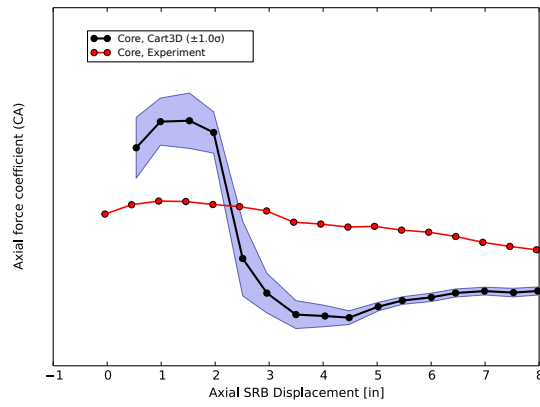
The subsections that follow describe the three sets of solutions and give further details on some aspects of the computation. In addition, comparisons between the different sources of data are shown, including the wind tunnel test. Finally, some aspects of database creation and uncertainty quantification, which allows for the creation of a valid aerodynamic database despite the issues with the Cart3D solutions, are briefly summarized.

A. Wind Tunnel Comparison

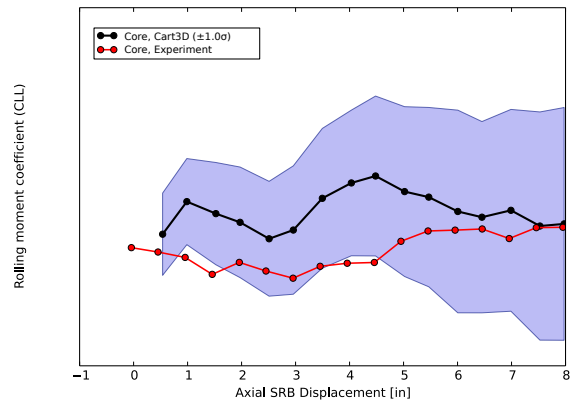
The dedicated SLS booster separation wind tunnel test, which was performed at the NASA Langley Research Center Unitary Plan Wind Tunnel in 2014, consisted of runs with three separately moving bodies and cold-gas (air) BSM thrusters. Data was collected in sweeps in which the boosters were moved in the $+x$ direction (downstream) while holding constant values for the other position variables. Collected data included all six force and moments on the core. The booster data consists of two forces and two moments, with the axial force and rolling moment not measured. The core is relatively fixed, and the position of the boosters could be controlled in all six degrees of freedom except for roll.

Figure 8 shows the surface geometry of the model used to simulate the wind tunnel test in Cart3D. The image in Fig. 8a with only stings was used for the database of Cart3D comparisons to the experimental data, which in turn contributed to the SLS booster separation uncertainty quantification. The support hardware used to hold the three separate models in place during the test are shown in Fig. 8b. Some of the BSM-off data showed unexpectedly poor agreement between Cart3D and experimental data, which appeared to be caused by interference from the support hardware (which was not modeled in the Cart3D simulations). Some additional Cart3D simulations were run with the support hardware also included, but results from that investigation are not presented here.


Figures 9 and 10 show example comparisons between Cart3D and experiment for a so-called dx sweep. The fixed SRB position variables are listed at the top of each image, and the horizontal axis of the data is the axial displacement, Δx or dx . Cart3D simulations were performed for every half inch of model-scale axial displacement. These results are shown as black dots, and a similar sampling of experimental data is shown



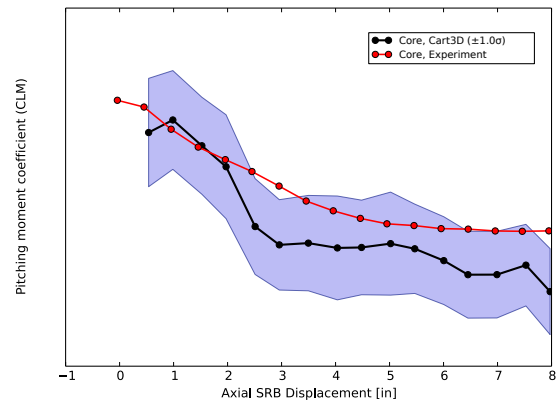
a) Core axial force, C_A




b) Core rolling moment, C_ℓ



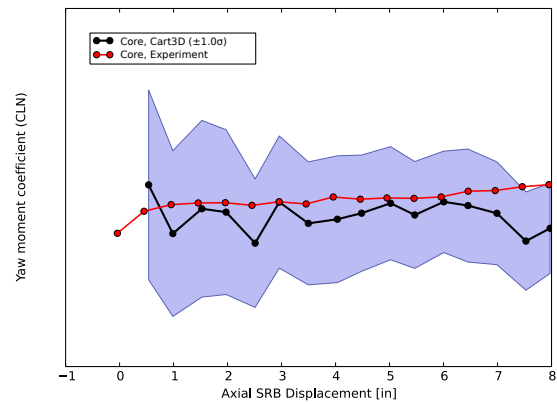
c) Core side force, C_Y



d) Core pitching moment, C_m



e) Core normal force, C_N



f) Core yawing moment, C_n

Figure 9. Sample Cart3D comparison to UPWT experimental data on core component

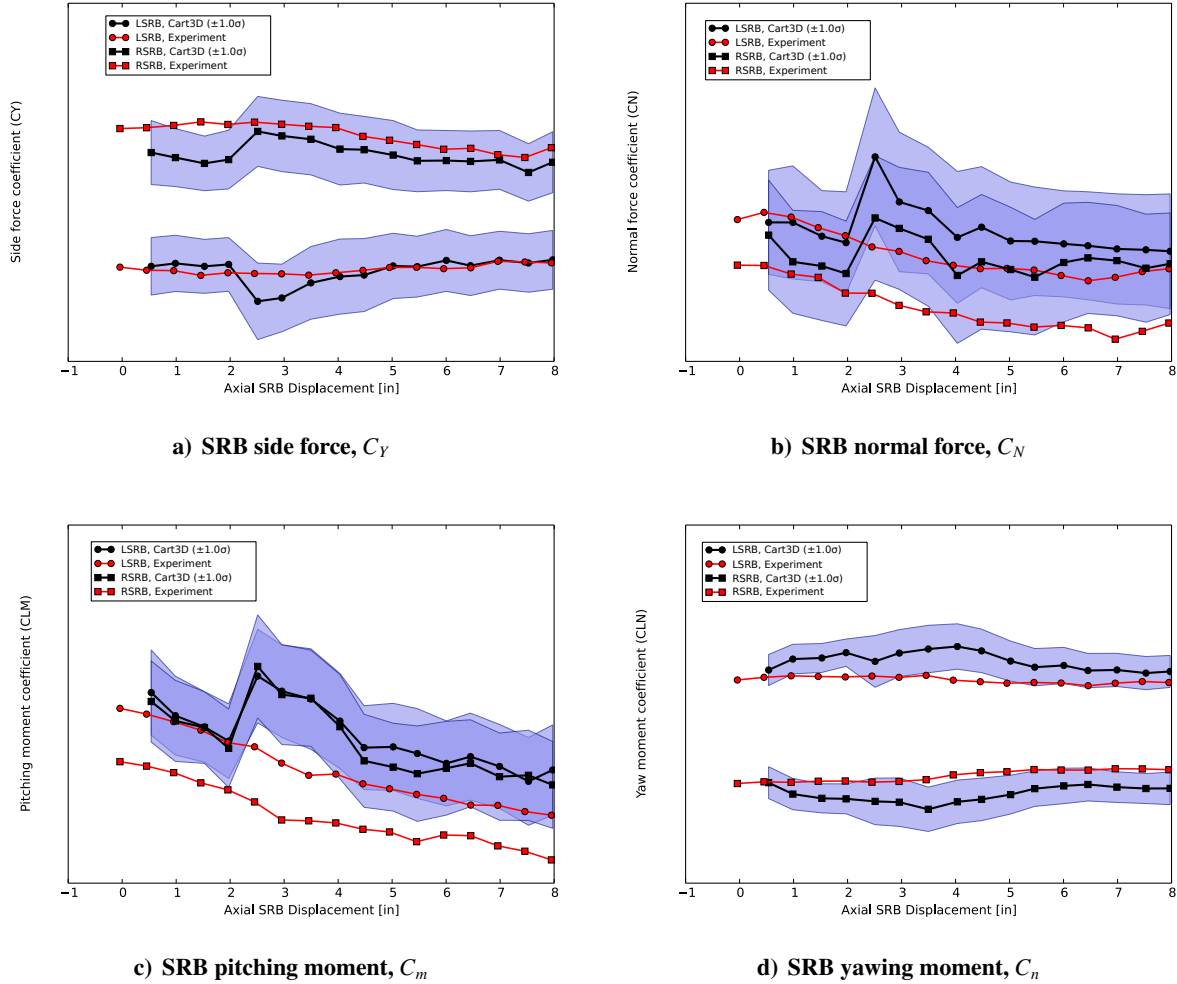


Figure 10. Sample Cart3D comparisons to UPWT experimental data on SRBs

with red dots, although the actual experimental data is much denser in Δx . The blue regions in Figs. 9-11 show the range of values within one iterative standard deviation from the black Cart3D result. The mean value (black dots) and standard deviation are calculated using the last 5000 iterations of the simulation for statistics. Finally, regions of slightly darker blue in Figs. 10 and 11 indicate values that are within one standard deviation of both the left and right booster Cart3D solutions.

The agreement between the experiment and Cart3D, especially for the boosters, is fairly good. When experimental uncertainty and Cart3D unsteadiness uncertainty are considered, the difference between the two data sets is not statistically significant for the boosters. In many cases (though not including Fig. 9), the differences in core C_N and C_m is large compared to those uncertainties. Note that the booster comparisons shown in Figs. 10 and 11 do not include C_A and C_ℓ because there is no experimental data for comparison.

Figure 11 shows the comparison of booster forces and moments, which is the more important part of the data, for another sweep. The quality of the comparison between Cart3D and experiment is similar to that of Fig. 10. Although it is not practical to show all 117 such sweeps, the quality of the results shown here is characteristic of all the sweeps. There are some cases that show worse agreement in core C_N and C_m , but the booster results are consistently good.

In addition to the BSM-on results that have been highlighted in the preceding paragraphs, there are more than 250 such sweeps of comparisons in which the BSMs are off. The BSM-off results are basically simple

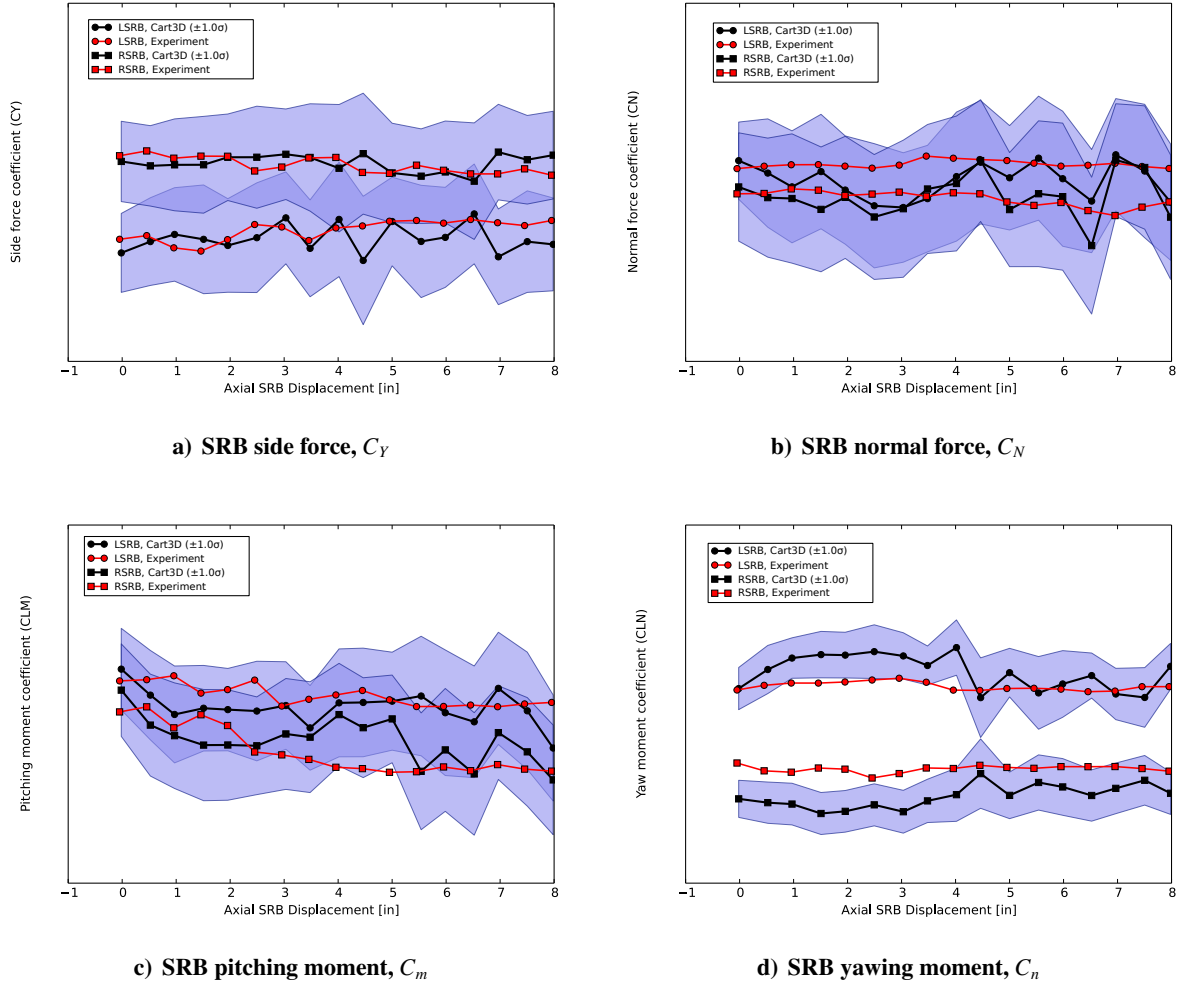



Figure 11. Sample Cart3D comparisons to UPWT experimental data for a different Δx sweep

simulations of three supersonic bodies. Such cases are not particularly challenging, although the analysis did assist in identifying the potential influence of the support hardware. Supersonic flow without further complications is also a particularly good candidate for the type of output-based mesh adaptation used in Cart3D, and the agreement between Cart3D and Overflow was consistently very good for these cases. Since these BSM-off cases are relatively straightforward, no further analysis is presented in this paper.


Although the wind tunnel test was naturally designed to be as close of a simulation of flight as possible, there are several interesting challenges that affect how Cart3D is run. As with most wind tunnel tests, the Reynolds number was not matched, but that has no effect on Cart3D, which is an inviscid solver. Further, at the low dynamic pressures where booster separation occurs, it should not be surprising that an Overflow study showed that the effect of Reynolds number is negligible for these cases. The more noteworthy differences address the BSMs. First, the wind tunnel was run with air as the BSM exhaust gas, which differs from actual BSM gas. From the perspective of Cart3D, which does not have support currently for differing ratios of specific heat, this makes things much simpler and makes the Cart3D simulation more realistic. However, a consequence of trying to match the BSM thrust and exit pressure with a different working fluid was an increase in the flow Mach number. Physical size constraints on the 0.9% model also resulted in a BSM nozzle and plenum geometry that was different from the flight geometry. Another complication to note is that the support hardware put the boosters into slightly asymmetric positions, as would be expected with any

real hardware. The decision was made to match the wind tunnel configuration as closely as possible, with the minor consequence that the Cart3D run matrix included roughly twice as many input variables.



a) Left SRB side force coefficient iterative history

b) Right SRB normal force coefficient iterative history



c) Static pressure, $y=0$

d) Static pressure, $z=0$

Figure 12. Sample of wind tunnel comparison with iterative history and static pressure slices at the final iteration

Figure 12 shows an example of the analysis applied to each individual case and at the same time shows a result typical of the wind tunnel comparison runs. Many numbers are removed for sensitivity reasons, but the example shows the type of information that is available with very little effort for each case. More so than the database runs described in the following section, the wind-tunnel simulation iterative histories show a large amount of “unsteadiness.” Since the cases were run with steady-state inputs, this is not true unsteadiness, but it appears to be similar in nature. Since no Overflow solutions with BSM thrust were available at the time that these cases were run, the initial analysis was that the BSMs caused unsteadiness in the flow. Although this is probably true to some small degree, later analysis indicates that the BSM-on flow is essentially steady, and the variation shown in Fig. 12 is mostly due to misplaced mesh resolution that is caused primarily by the shocks moving slightly after each refinement.

B. Cart3D-Based Database Runs

The introductory paragraphs to this section, Sec. IV, describes the primary features of the database runs. An interesting variety of solutions was obtained, including the three shock positions and various degrees of flow convergence. There were many cases for which the 3000 second-order iterations on the final mesh were completely steady with good convergence of the global L_2 norm. However, this convergence did not correspond well with solution quality as measured by force symmetry.

The other aspects of the database Cart3D runs are described in other sections. Sample results similar to those shown in Fig. 12 were monitored for all cases, and a more detailed description is provided in the following Sec. IV.C. The nature of the run matrix described in Sec. III makes plotting relevant trends in the Cart3D data difficult, and no such plots are included here. However, the comparison to Overflow data in Sec. IV.D does show the approximate quality of a family of solutions. The overall quality of the database solutions is discussed in the uncertainty discussion of Sec. IV.E, but it is worth noting that the uncertainties on the boosters decreased by roughly an order of magnitude compared to DAC-2 as a result of this work. However, uncertainty of the core normal force and pitching moment actually increased compared to DAC-2. The primary driver of uncertainty for most of the force and moment coefficients was the iterative “unsteadiness” in the Cart3D solution, although for some forces and moments on the core, the driver was the Overflow-to-Cart3D difference. All of this discussion has been used to lead to improved performance should it be needed for subsequent separation analysis, as discussed in more detail in Sec. IV.F.

C. Quality Analysis and Assurance

Of the 18 force and moment coefficients that are delivered for each case as part of the aerodynamic database for separation, at least one symmetry check is available for 15 of them when core sideslip is zero and all 18 when sideslip is nonzero. For the cases in which core sideslip is zero, symmetry between the two boosters is an obvious source, and one expects the core to have zero values of C_Y , C_ℓ , and C_n . For cases with nonzero core angle of sideslip, symmetry evaluation comes from the run with the opposite value of sideslip. Running both sideslip angles was not strictly necessary for creating the database, but it was very useful in evaluating quality. The first step of the quality assurance and assessment procedure was to check symmetry of the coefficients as available. The most important coefficient to check is booster side force, but the other coefficients were also checked as available.

The second step was to check the iterative histories for anything out of the ordinary. For instance, if a coefficient has an unconverged mean that was not caught during the original run procedure, that case would be restarted. The residual histories were also checked for unexplained behavior, but this criterion resulted in few or no reruns. To the extent possible, mesh convergence was also investigated at this point. Because of the complexity of the flow, mesh convergence was not attainable for BSM-on cases, but an investigation of how much a further mesh adaptation would change the results was used to influence the uncertainty quantification. All database runs used the same number of adaptations, had very similar dimensions of the finest cell, and had the a number of volume cells within a 10% range (16 to 20 million cells).

The third step is to visually inspect the two static pressure slices for symmetry. In cases of nonzero core sideslip angle, the additional check of making sure that the two solutions with opposite sideslip have the same overall shock position. These checks led to quite a few reruns, although the shock position symmetry is usually caught by the first criterion.

Figure 13 shows an example of the document format used to actually apply these quality assessment procedures. These documents are created as multipage PDFs using \LaTeX ; the iterative histories are plotted using Matplotlib, and the static pressure slices are plotted with TecPlot. In practice, these documents were usually made with roughly 100 pages per document, and these reports were made available to the SLS team. The reports do not include iterative histories for 12 of the 18 force-and-moment coefficients, but it is a very useful summary compressed onto a single page both for results reporting and quality assessment.

D. Comparison to Overflow Results

The Overflow code was run for 45 different booster-separation configurations with both the BSMs on and with the plumes off. Example comparisons between the right SRB aerodynamic side and normal forces, and the yawing and pitching moments predicted by Overflow and Cart3D are presented in Fig. 14. These plots show the aerodynamic coefficients plotted versus the SRB axial-displacement distance in feet for a handful of different configurations. In each plot, the Overflow data is represented with the circular symbols

Conditions

Variable	Abbr.	Value
n	n	2106
dx	dx	2.0
dy	dy	1.0
dz	dz	-2.0
dpsi	dps	0.5
dtheta	dth	0.0
dphi	dphi	0.0
alpha	a	0.0
beta	b	0.0
mach	m	N.NN
q	q	NN.N
CT_BSM	CT	N.NN
CT_SRB		N.N
CT_CSE		NN.N
conf g		poweron

Force & moment summary (Iteration 5000, nSt at s=2000)

Coefficient	CORE_No_Base	LSRB_No_Base	RSRB_No_Base
CA mean, $\mu(C_A)$	N.NNNN	N.NNNN	N.NNNN
CA standard deviation, $\sigma(C_A)$	N.NNNN	N.NNNN	N.NNNN
CY mean, $\mu(C_Y)$	-N.NNNN	-N.NNNN	N.NNNN
CY standard deviation, $\sigma(C_Y)$	N.NNNN	N.NNNN	N.NNNN
CY iterative uncertainty, $\epsilon(C_Y)$	N.NNNN	N.NNNN	N.NNNN
CN mean, $\mu(C_N)$	-N.NNNN	-N.NNNN	-N.NNNN
CN standard deviation, $\sigma(C_N)$	N.NNNN	N.NNNN	N.NNNN
CN iterative uncertainty, $\epsilon(C_N)$	N.NNNN	N.NNNN	N.NNNN
CLL mean, $\mu(C_{\cdot})$	-N.NNNN	N.NNNN	-N.NNNN
CLL standard deviation, $\sigma(C_{\cdot})$	N.NNNN	N.NNNN	N.NNNN
CLM mean, $\mu(C_m)$	-N.NNNN	N.NNNN	N.NNNN
CLM standard deviation, $\sigma(C_m)$	N.NNNN	N.NNNN	N.NNNN
CLN mean, $\mu(C_n)$	-N.NNNN	N.NNNN	-N.NNNN
CLN standard deviation, $\sigma(C_n)$	N.NNNN	N.NNNN	N.NNNN

Force, moment, and residual histories

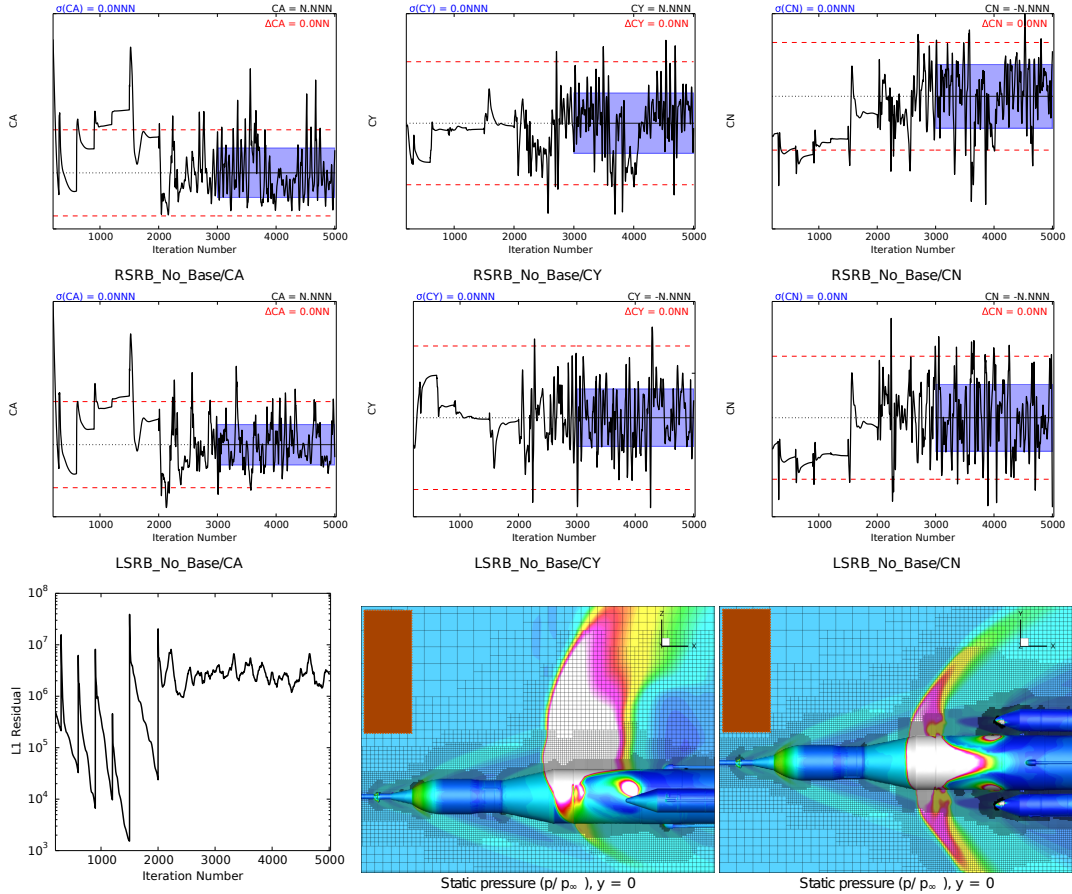


Figure 13. Sample single-page report summary for one case with BSMs on with sensitive data removed or replaced with N

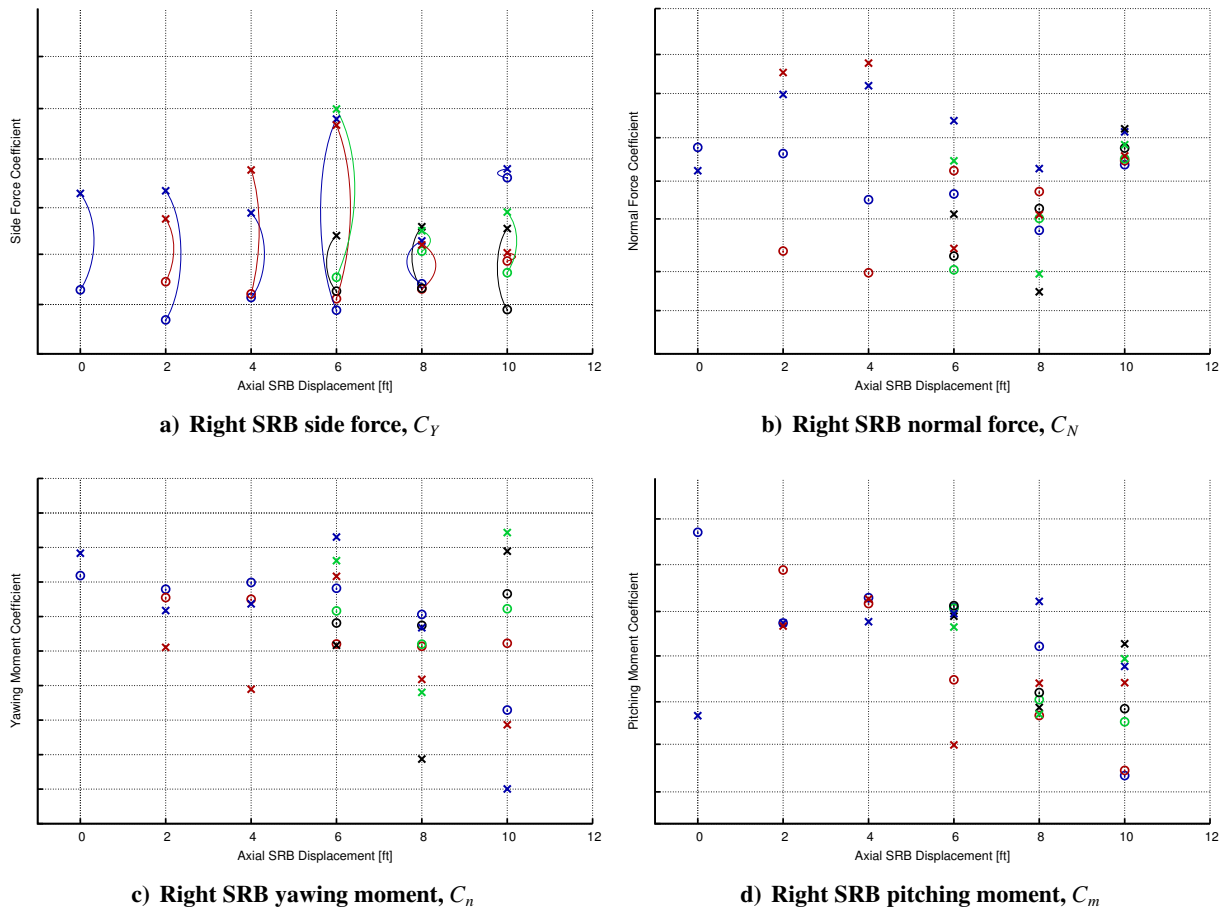


Figure 14. Cart3D-to-Overflow comparison of right SRB aerodynamics. Circles are Overflow results, and X symbols are Cart3D results. Cases at the same dx value with the same color are for matching conditions.

whereas the Cart3D data is plotted with the X symbols. The different colors represent different booster configurations. In Fig. 14a, matching conditions are connected by curves.

The differences in the aerodynamic data between the two codes contributes to the database uncertainty quantification, which is summarized in the next subsection. The differences between the two codes is generally very small, although the inability to publish the actual quantitative data makes this difficult to show here. Both the booster side-force and normal-force coefficients predicted by Overflow are consistently lower than the Cart3D values. These differences tend to be greater at the conditions with smaller axial booster displacements. The yawing-moment predicted by the two codes is generally in good agreement with similar trends predicted as the SRBs travel downstream. The pitching moment predicted by Overflow starts out higher or more positive at the beginning of the separation and tends to decrease as the SRBs move downstream whereas Cart3D predicts a lower value to start with an increase over the first few feet followed by a slight decrease.

E. Uncertainty Analysis and Database Creation

The implementation of the database was completed by the Configuration Aerodynamics branch at NASA Langley Research Center. A brief summary is included here to provide some understanding of how the data were used and how the aerodynamic results affected the rest of the SLS program. Since the run matrix as described in Sec. III is quite irregular, the first step of database creation was to interpolate the data onto a fully rectangular matrix (that is, an 8-dimensional table with fixed break points in all 8 dimensions and data

at each point). To do this, the DACE toolkit for Kriging interpolation [14] was utilized. Once this full data table was constructed, it was placed into a fast interpolation procedure based on more traditional procedures and previously developed for DAC-2 simulations.

Uncertainty analysis is perhaps the most important factor in determining clearance for this booster separation configuration. Although the primary work on this portion of the problem was performed by the same team at NASA Langley, its importance as an output of this problem dictates that at least a brief discussion should be included here.

Firstly, the individual contributors to uncertainty include difference between Cart3D and the wind tunnel experiment, difference between Cart3D and Overflow for the flight geometry, iterative variation in Cart3D results, interpolation error, and uncertainty due to asymmetry. These individual contributors are combined by taking the square root of the sum of the squares, and a further multiplier is added to cover any unmodeled uncertainties (this multiplier must be greater than or equal to 1.0). The difference between Cart3D and wind tunnel has its own contributors, which include Cart3D iterative variation, experimental uncertainty, repeatability uncertainty, and the difference between Cart3D and the experiment. In order to isolate the Cart3D-to-wind-tunnel contribution, we start with the square of the difference between the means and then subtract the square of the experimental uncertainty and the Cart3D iterative variation. After subtracting these contributors, the isolated Cart3D-to-wind-tunnel uncertainty is the square root of the resulting difference. Subtracting the Cart3D iterative variation may seem like an unfair reduction in the experimental contribution to uncertainty, but it is added back in later; this action simply avoids double-counting that aspect of the uncertainty.

After searching for significant trends in this uncertainty versus other variables such as axial booster displacement (dx), it was decided to simply use a single number for each coefficient. Thus the uncertainty database delivered to the guidance, navigation, & control team consists of 18 numbers for the BSM-on database. Part of the explanation for this decision is that the database only included cases with axial displacements of 10 ft or less. For a more complete separation aerodynamics database, the uncertainty magnitude would vary more with changes in conditions, but for this data a simpler uncertainty implementation is the most appropriate.

As mentioned before, the uncertainty magnitudes on the booster force and moments reduced by roughly an order of magnitude compared to the DAC-2 database. The reasons for this improvement include utilization of better inputs from Cart3D, upgrades to the Cart3D code over the last year or so, better availability of data (there was no relevant wind tunnel test at the time of DAC-2), and limiting the database to cases with small booster pitch and yaw angles. At the same time, uncertainties on the core remained about the same or increased slightly. Compared to the decrease in booster uncertainty, the difference here is that the BSM thrust drives core uncertainty, whereas the DAC-2 booster uncertainties were driven by BSM-off cases with larger pitch and yaw angles. With the exception of core normal force and core pitching moment, the largest individual contributor to each uncertainty in this database was the Cart3D iterative variation. For the remaining two coefficients (core normal force and pitching moment), the largest individual contributor was the difference between Cart3D and Overflow.

This improvement in booster uncertainty combined with roughly unchanged core uncertainty is a better result than it at first seems. The boosters, which have spent their usable fuel, are much lighter than the core and are not actively controlled during separation. In contrast, the core still is producing near its maximum thrust, is actively controlled, and has far more inertia than the roughly empty boosters. As a result, uncertainty on the core aerodynamics is much more tolerable. When used in a Monte Carlo trajectory simulation, the net result of the new database was an appreciable reduction in the uncertainty on booster-to-core clearance during separation, which is the most important output of this work.

F. Improvements and Extensions for Future Database Analyses

The lessons learned from quality assessment and analysis of what factors have the most impact on uncertainty made it clear which aspects of the Cart3D solutions would be most beneficial to improve. In particular,

it was not at all clear when the analysis was started that the forward pressure rise (shock position 1) was the physical solution, but after looking at multiple data sources, we decided to investigate methods to prevent the inviscid solver from selecting aft shock locations. Similarly, it was not expected that Overflow solutions would be relatively steady or that even Cart3D in some cases reported a steady iterative history.

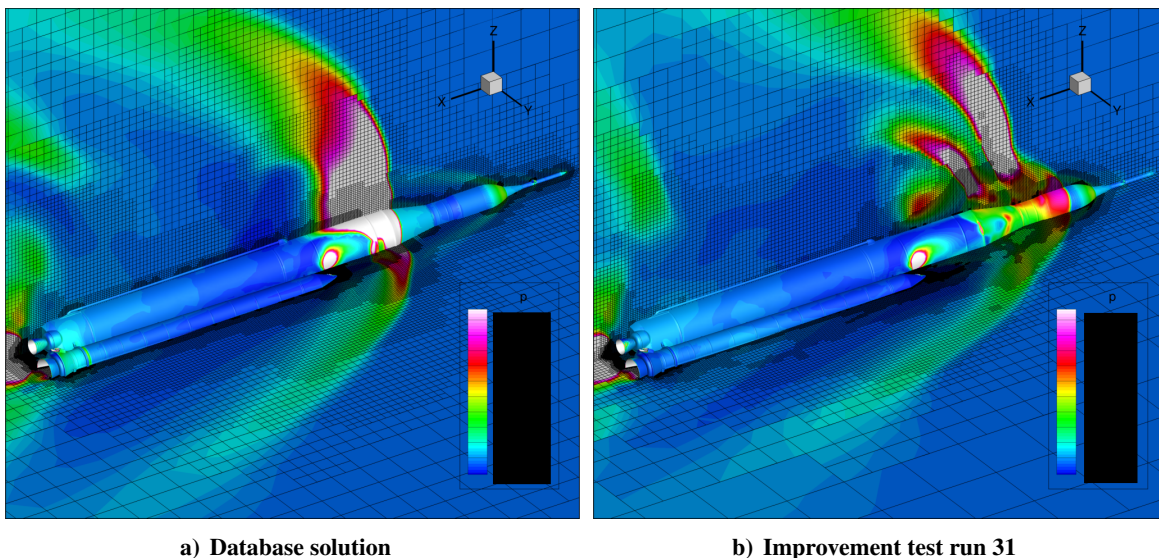


Figure 15. Comparison of static pressure between database result and one of many attempts to alter results

Based on these lessons, attempts have been made after the completion of this database to find inputs that improve the Cart3D results using a practical amount of computational resources. In particular, an effective strategy has been to use an initial mesh with coarse resolution around the aft shock position. This enables the solver to attempt the forward-shock configuration, but adaptive refinement does allow the shock to fall back to position 3 when the BSM thrust is decreased. In summary, this provides a gentle nudge favoring forward shock positions, but it is far removed from forcing Cart3D to select a certain flow topology. Figure 15 shows a comparison between the original result used in the DAC-3/CDR database in Fig. 15a and the 31st attempt to improve the results shown in Fig. 15b. At the time of writing, over 130 such test runs have been made.

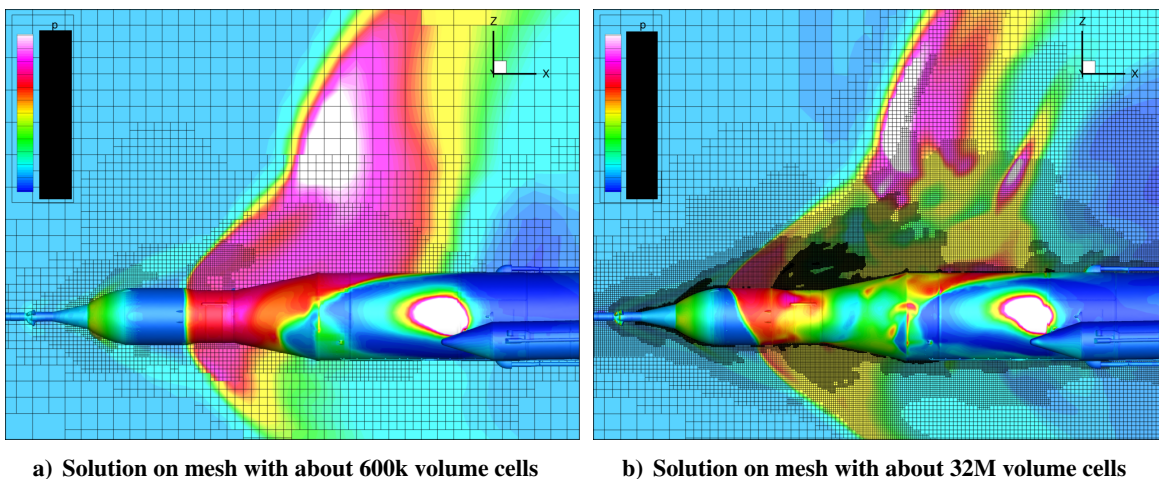


Figure 16. Effect of nearly uniform mesh refinement

Figure 16 demonstrates an instance in which three further adaptations with large growth ratios (such

that half of the volume cells are refined) causes the qualitative flow to shift toward what is believed to be a more accurate solution. In Fig. 16b, it is clear from the regions of increased mesh density that the flow demonstrated significant changes when the inputs shifted from first-order to second-order after completion of the adaptation cycles. The run conditions for the case in Fig. 16 are different from those of Fig. 15, and due to the orientation of the forward BSMs, it is a configuration in which predicting shock position 1 is more challenging. The information from this discussion is not substantially different from the flow sensitivity discussed in the introduction. It should not be surprising that mesh convergence—even of the iteratively averaged forces—could not be demonstrated for such a complex problem. However, the tool performed remarkably in the face of such a difficult flow, and the detailed uncertainty quantification ensured that the database delivered to the SLS program was aware of and accounted for these difficulties. The extensive input sensitivity study that was performed has demonstrated that a uniform refinement study (or similar) remains a useful tool in deciding the true nature of how a particular solver will handle a difficult flow—a lesson undoubtedly learned countless times by previous aerodynamicists. The important outcome of this investigation is that future databases for related configurations will come with narrower uncertainty bounds and hopefully will be developed at a lower cost.

Table 2 contains a list of additional variables that may be incorporated into a future database created before the first flight of SLS. These variables may also be used for early separation databases of future iterations of the vehicle as well. Incorporating additional independent variables into a database is simpler if conditions are selected more efficiently, and this configuration is no exception.

Table 2. Additional input variables that could affect separation aerodynamics

Variable	Description
<i>dphi</i>	Booster roll angle relative to core body axes [deg]
<i>mach</i>	Core Mach number
<i>q</i>	Dynamic pressure seen by the core [psf]
<i>h</i>	Altitude [ft]
<i>CT_SRB</i>	Booster engine thrust coefficient
<i>CT_BSM</i>	Booster separation motor thrust coefficient dispersed
<i>CT_CSE</i>	Core engine thrust coefficient

There is an additional requirement to investigate the performance when a single BSM or a single core engine fails. Checking the effect of each BSM and each core failure individually results in 20 different simulations, and considering combinations of failures would increase the size of the matrix even further. Furthermore, the time at which a core engine fails affects the Mach number and dynamic pressure at separation. Occasionally analysis is requested for other off-nominal scenarios, such as the boosters separating at different times, and we attempt to support such requests as well as possible with the resources available.

Some of these variables and considerations, for example booster thrust (*CT_SRB*), have only a small effect on aerodynamic performance and should not be incorporated into the database. Since each of these potential degrees of freedom has a range of possible conditions, it is worthwhile to investigate the effect of each and either incorporate it into the flight database or ignore it with a higher degree of confidence.

V. Conclusions

A large, multifaceted aerodynamics database for the Space Launch System booster separation event was created and discussed. The aerodynamic CFD simulations included three broad categories: inviscid Cart3D simulation of the flight geometry for the creation of the database, inviscid Cart3D of a wind tunnel configuration for comparison to experiment, and viscous Overflow simulations of the flight geometry. In addition to the complexity of the database including 22 powered rocket boundary conditions and an 8-dimensional run matrix, some new features of the popular NASA Cart3D and Overflow solvers were applied to a complex configuration. Key examples of new features used include MPI-enabled Cart3D, time-accurate Cart3D, and

mesh adaption in Overflow. A brief discussion was given of how the new aerodynamic database affected other members of the Space Launch System team. Due to both the complexity of the flow and the necessary size of the database, a Cartesian inviscid solver is the only tool that could be used to support this work with today's computational resources.

Strong sensitivity to the initial mesh was documented even when using Cart3D's adaptive meshing feature. Overflow did not demonstrate such a sensitivity, which provides one piece of evidence that viscosity plays an important role in biasing the flow toward one particular shock position. New techniques for quality assessment and assurance were discussed that enabled putting expert eyes on more than 12,000 pages worth of summary documents that were generated with minimal user interaction. In addition, work continues toward gaining higher accuracy and efficiency for this and similar problems.

Partially due to this new aerodynamic database and partially due to other factors, the uncertainty on booster force and moments was dramatically reduced. This was a critical result that had a measurable impact on confidence that the booster separation event will be successful during the first flight of SLS.

Regarding the generation of similar aerodynamic databases for problems with jet interaction, significant lessons were learned. Especially if there are any nozzles whose thrust directions has an upstream component, there may be multiple flow topologies. While a mesh convergence study is always a recommended practice, and while resource availability or time constraints often prevents engineers from properly determining the most appropriate mesh resolution, complex configurations such as the SLS booster separation benefit immensely from investigating mesh resolution and sensitivity to other inputs. When any particular solver presents multiple qualitative solutions—particularly if that solver is inviscid—any available tool should be used to help determine the physical reality of the various solution families. The description of the work carried out by NASA described in this paper demonstrates positive practices, experiences that identified possible improvements, and techniques specific to Cart3D that are useful for this type of aerodynamics problem.

Acknowledgments

Some of the insights for this work were possible due to work done for Design and Analysis Cycle 2 by Jeffrey A. Housman, Jeffrey T. Onufer, William M. Chan, and others. Database implementation and uncertainty analysis was performed by Jeremy T. Pinier, David Chan, and Floyd Wilcox at NASA Langley Research Center, and the wind tunnel test at the Unitary Plan Wind Tunnel was performed by the same team. Interactions with the Space Launch System guidance, navigation, & control team were primarily with Dane J. Childers of NASA Marshall Spaceflight Center.

Resources supporting this work were provided by the NASA High-End Computing Capability (HECC) program through the NASA Advanced Supercomputing (NAS) Division at Ames Research Center.

References

- [1] Aftosmis, M., Berger, M., and Adomavicius, G., "A Parallel Multilevel Method for Adaptively Refined Cartesian Grids with Embedded Boundaries," *38th Aerospace Sciences Meeting*, 2000, AIAA Paper 2000-0808.
- [2] Nemec, M. and Aftosmis, M., "Adjoint Error Estimation and Adaptive Refinement for Embedded-Boundary Cartesian Meshes," *18th AIAA Computational Fluid Dynamics Conference*, 2007, AIAA Paper 2007-4187.
- [3] Aftosmis, M. and Rogers, S. E., "Effects of Jet-Interaction on Pitch Control of a Launch Abort Vehicle," *46th AIAA Aerospace Sciences Meeting and Exhibit*, 2008, AIAA Paper 2008-1281.
- [4] Brauckmann, G. J., Greathouse, J. S., and White, M. E., "Rocket Plume Scaling for Orion Wind Tunnel Testing," *29th AIAA Applied Aerodynamics Conference*, 2011, AIAA Paper 2011-3341.
- [5] Nichols, R. H., Tramel, R. W., and Buning, P. G., "Solver and Turbulence Model Upgrades to OVERFLOW2 for Unsteady and High-Speed Applications," *36th AIAA Fluid Dynamics Conference*, 2006, AIAA Paper 2006-2824.
- [6] Klopfer, G. H., Onufer, J. T., Pandya, S. A., Chan, W. M., Kless, J. E., and Lee, H. C., "Analyses of the Ares I A106-Plus Launch Vehicle First Stage Separation," Tech. Rep. FS-TR-00002, NASA, 2011.

- [7] Pamadi, B. N., Pei, J., Pinier, J. T., Holland, S. D., and Covell, P. F., "Aerodynamic Analyses and Database Development for Ares I Vehicle First-Stage Separation," *Journal of Aircraft*, Vol. 49, No. 5, 2012, pp. 864–874.
- [8] Kiris, C. C., Housman, J., Gusman, M., Shauerhamer, D., Deere, K., Elmiligui, A., Khaled, A.-H., Parlette, E., Andrews, M., and Blevins, J., *49th AIAA Aerospace Science Meeting*, 2011, AIAA Paper 2011-16.
- [9] Gusman, M. R., Barad, M. F., and Kiris, C. C., "Aerodynamic Database Generation for SRB Separation from a Heavy Lift Launch Vehicle," *29th AIAA Applied Aerodynamics Conference*, 2011, AIAA Paper 2011-3651.
- [10] Rogers, S. E., Dalle, D. J., and Chan, W. M., "CFD Simulations of the Space Launch System Ascent Aerodynamics and Booster Separation," *53rd AIAA Aerospace Sciences Meeting*, January 2015, AIAA Paper 2015-0778.
- [11] Menter, F. R., "Two-Equation Eddy-Viscosity Turbulence Models for Engineering Applications," *AIAA Journal*, Vol. 32, No. 8, 1994, pp. 1598.
- [12] Menter, F. R., Kuntz, M., and Langtry, R., "Ten Years of Industrial Experience with the SST Turbulence Model," *Turbulence, Heat and Mass Transfer*, Vol. 4, 2003, pp. 625.
- [13] Bakhtian, N., M. and Aftosmis, M. J., "Analysis of Inviscid Simulations for the Study of Supersonic Retropropulsion," *29th AIAA Applied Aerodynamics Conference*, 2011, AIAA Paper 2011-3194.
- [14] Lophaven, S. N., Nielsen, H. B., and Søndergaard, J., "DACE, A Matlab Kriging Toolbox," Tech. Rep. IMM-TR-2002-12, Informatics and Mathematical Modeling, Technical University of Denmark, 2002.